CAMBRIDGE

Cambridge University Press
978-0-521-10918-5 - Language and Spatial Cognition: An Interdisciplinary Study of the Prepositions in English
Annette Herskovits
Excerpt
More information

Introduction

This book presents a fine-grained analysis of the semantics and pragmatics of locative expressions, expressions like *the teapot on the table, the man behind the counter*, and *the circle of light under the lamp*; then examines, in light of that analysis, the possibility of constructing a computational model of comprehension and production in the spatial domain.

This work did not grow out of an interest in space per se, but in the workings of language in general, its relation to cognition, and the plausibility of the computational metaphor. That led me to look for a suitable domain – simple, but neither too simple nor too narrow. Expressions predicating the position of one object relative to another satisfy this requirement. They are simple in several ways: their semantic structure involves a two-place, or at most a three-place predicate (e.g., *between, beyond*); and the position of objects is a concrete, unambiguous, and well-understood physical notion. Yet, the domain defined by these expressions turns out to be surprisingly complex. What one says does not in general correspond in a simple manner to elementary geometric predicates, and, contrary to common belief, a precise geometric definition of the place and shape of the objects in question does not suffice for choosing a preposition appropriately.

But space was compelling because it occupies a privileged place in language and in the cognitive system. Spatial metaphors pervade language, and not solely as a source of colorful expressions; they are necessary to conceptualizing various semantic domains, in particular abstract domains. In addition, much of our reasoning ability relies on visual thinking, on the manipulation of mental images.

In short, the domain of space provides a particularly clear view of most of the fundamental issues pertinent to language and to cognition. I will begin by placing the work in this larger context.

0.1 Spatial cognition and language

Underlying most work in linguistics and artificial intelligence is the belief that the world comes nicely structured along lines defined by universally accepted and objectively definable categories. The physical world seems

© Cambridge University Press www.cambridge.org

CAMBRIDGE

Cambridge University Press
978-0-521-10918-5 - Language and Spatial Cognition: An Interdisciplinary Study of the Prepositions in English
Annette Herskovits
Excerpt
More information

particularly well-delineated and unambiguously conceptualizable, made up of objects, each with a well-defined shape and position in space. Words and morphemes denote these well-drawn categories, and what people talk about (even if it is not readily apparent because people are not always careful) are these discrete objects and their relations in the world. Such an "objectivist" view of meaning (Lakoff and Johnson 1980) is a little hard to extend to abstract entities, but it seems uncontroversial for the physical world.

In fact, the spatial domain remains incomprehensible when looked at in this way. Careful examination of a range of locative sentences reveals that the spatial objects related often do not actually exist in the world, but are mental constructions – beyond and above that first mental construction produced by perception. Even objecthood is not an unalterable given.

Clearly we cannot know reality objectively, since knowledge is a product of the human organism. Yet there is a "fundamental" or "canonical" view of the physical world, which in everyday life is taken as "the world as it is". But language does not directly reflect that view. Idealizations, approximations, conceptualizations, and selections mediate between the canonical view and language. These are not strictly linguistic: they reflect spatial representational structures that are necessary for action (movement and manipulation), interactional properties (such as the function of a cup, how we hold it), often appearance rather than the actual configuration of the objects, and – most importantly – the concerns and purposes of the speaker.

The objectivist view has given rise to various mistaken hypotheses on the structure of the domain of spatial expressions. In linguistics, it has been assumed commonly that the meaning of locative expressions could be entirely specified as a proposition constructed out of a simple geometric relation applying to the objects. In artificial intelligence, there is a widespread implicit assumption that, given descriptions of the shapes of two objects, given their location (for example, by means of coordinates in some system of reference), and, in some cases, the location of an observer, one can select an appropriate preposition to describe the spatial relation of the objects. The data do not bear out either view, and this book gives abundant proof of their falsity.

Yet, our language obviously maps onto the world as it is. After all, a rose is a rose, and a cup is a cup, and every human being – whatever their individual, linguistic, and cultural differences – sees each as an object with a certain shape, color, substance. I will offer an account of the spatial domain that makes this mapping comprehensible, while avoiding the objectivist fallacy.

© Cambridge University Press www.cambridge.org

CAMBRIDGE

Cambridge University Press
978-0-521-10918-5 - Language and Spatial Cognition: An Interdisciplinary Study of the Prepositions in English
Annette Herskovits
Excerpt
More information

0.2 Linguistic issues

Central to the semantic approach taken in this book are the complementary notions of ideal and of deviations from the ideal. Anyone who has looked carefully at a semantic domain knows how elusive comprehensive and accurate semantic descriptions are. Lexical meaning has been remarkably refractory to attempts at formal definition. I will show that one can make sense of many apparent irregularities of semantic data by assuming that word meanings are defined in an ideal world – in the spatial domain, a world of lines, points, surfaces, and of definite relations of inclusion, contact, intersection, and so forth. To describe and communicate facts about the complex and imperfect world that surrounds us, we must bend and stretch these ideal concepts. Obviously, this bending and stretching cannot be done at the speaker's whim; while it may not all be the result of linguistic rules, it is constrained by the need to preserve mutual comprehension, so that our references to objects and facts are, usually, successful. For that reason, ideal concepts must be complemented by a precise account of the way in which such concepts are distorted in a linguistic and situational context.

Thus, I will propose **ideal meanings** for the spatial prepositions, and then describe two types of deviations from the ideal: those based on convention, which basically give rise to polysemy, and those that result from pragmatic processes of allowance, or tolerance. "Polysemy" may not be a fully appropriate term: if we go by the speaker's intuitions, "the meaning" of the preposition is the ideal meaning; but the ideal meaning can be conventionally exploited in various ways, which must be recorded with the preposition in the lexicon. In my description of prepositions, I have therefore devised a series of **use types** to correspond to various classes of uses distinguished by different conventions.

A specification of lexical meaning has little value without an account of how lexical meanings combine to make up the meaning of phrases and sentences – the problem of compositionality – so I will propose meaning representations for complex locative expressions, representations that are built using the semantic information attached to the preposition as well as pragmatic and world knowledge.[1]

A striking feature of the composition process is the degree to which context must be called upon for a full specification of meaning. Spatial expressions in particular present a lot of hidden indexicality – that is, context is frequently necessary to determine what proposition they express. Terms like *to the right, in front of*, and so forth, are known to have indexical uses, where the hearer must determine the relation expressed by considering the location and orientation of an implicit observer. But indexicality extends much beyond these obvious cases (I am of course

© Cambridge University Press www.cambridge.org

CAMBRIDGE

Cambridge University Press
978-0-521-10918-5 - Language and Spatial Cognition: An Interdisciplinary Study of the Prepositions in English
Annette Herskovits
Excerpt
More information

concerned only with indexicality originating in the use of the preposition not with pronouns or other pro-forms that might occur in the expression). Many contextual factors modify the relation that is being expressed, or the particular spatial objects of which the relation is true, factors implicit in the intentions and concerns of the speaker.

Context also plays a role in allowing pragmatic inferences. We use general and specific world knowledge, and knowledge about communication, to infer facts beyond those entailed by the utterance.

Pragmatic inferences raise the problem of the boundary between pragmatics and semantics. I will firmly distinguish between those aspects of meaning that are the inalienable property of the expression alone, and those elements that depend on the particular context of use – grouping the former into a complex called the **situation type**[2] associated with the expression, or its **interpretation**. Whether this distinction coincides with the semantic/pragmatic distinction is controversial. First, the situation type associated with an expression includes elements of meaning which would normally be considered pragmatic, such as some constraints on the purpose of the expression. Second, linguists do not all agree on a definition of pragmatics: some include the treatment of indexical terms (e.g. Montague 1974); others handle indexical expressions with a standard truth-conditional apparatus and restrict pragmatics to those aspects of meaning that are not truth-conditional, thus excluding indexicality (e.g. Gazdar 1979); still others draw the line between conventional aspects of meaning, and those arrived at by principles of preferred interpretation (e.g. Wilson 1975).

In effect, in the practical context of considering the construction of an artificial intelligence system, I will end up cutting the cake in a way different from any employed in traditional linguistics, using several criteria to sort elements of meaning: determined by the expression versus determined by the context, conventional versus motivated, normal versus "out of the ordinary". The **normal situation type** associated with an expression, more restricted than the situation type in that it assumes conditions to be "normal", will play a central role in the framework proposed. In an informal way, I will call "pragmatic" some principles that guide the evaluation of various indexicals in context, others that restrict the choice of expressions, and also some inferences from world knowledge.

Finally, another essential issue is idiomaticity. Linguistic theories sort idioms from regular expressions. An idiom is often defined as any expression which conventionally conveys elements of meaning other than those obtained by straightforward application of the relevant rules of composition to the meanings of the component morphemes. But we will find many conventional peculiarities of use that cannot be derived from the morphemic meanings (unless they be redefined in rather unintuitive

© Cambridge University Press www.cambridge.org

CAMBRIDGE

Cambridge University Press
978-0-521-10918-5 - Language and Spatial Cognition: An Interdisciplinary Study of the Prepositions in English
Annette Herskovits
Excerpt
More information

fashion), even with expressions not perceived as idiomatic by native speakers. Idiomaticity seems to be a matter of degree: the meaning of most expressions depends partly on the lexical items combined, and partly on conventions attached to phrase-types; those for which the latter predominates are perceived as more idiomatic. Of course, all conventions must be stated in some part of the lexicon, which motivates the two levels in the lexical information attached to each preposition: the ideal meaning, and the use types – which are phrase types centered on the preposition, each with its own particular semantic and pragmatic features. This division will allow us to account for the mixture of phrasal idiosyncrasy and morphemic motivation that seems so prevalent in language.

0.3 Computer models of language

As researchers in artificial intelligence have discovered over the last ten years, building a faithful model of linguistic activity is quite a difficult task. One need not look at metaphors or abstractions to find this; the concrete and well-defined domain of locatives is as riddled with such difficulties as are more striking feats of linguistic activity, and it yields many clues as to their source.

These difficulties can be traced to three properties of language. The first applies only in the spatial domain – although it may have some equivalent in other domains: the fact, already mentioned, that the spatial objects we relate are mental constructions that I will call **geometric descriptions.** These geometric images are neither contained in, nor directly inferable from, the canonical description.

Another problem is the context-dependence of language. In communicating, we rely on implicit shared beliefs and concerns to a degree which has not been sufficiently recognized in artificial intelligence. Characterizing precisely the context-dependent parameters and how they vary with the context is most difficult.

Vagueness is the third difficulty. Vaguely defined entities and relations occur everywhere, and no satisfactory formal solution has ever been proposed for this problem. Simply allowing degrees of truth (as in Zadeh's "fuzzy logic" (1974)) is no solution; we must understand how we integrate the interplay of many vaguely defined parameters into some definite utterance. It is not clear that this can be solved by making explicit the set of relevant parameters and attaching an appropriate weight to each, because it is not clear that a list of such parameters and weights can be specified once and for all for every expression.

All this does not mean that building computer systems is pointless. One can construct, for specific applications, useful computer systems that simplify in well-defined ways; a good informal description of the meaning

© Cambridge University Press www.cambridge.org

CAMBRIDGE

Cambridge University Press
978-0-521-10918-5 - Language and Spatial Cognition: An Interdisciplinary Study of the Prepositions in English
Annette Herskovits
Excerpt
More information

and use of locative expressions could allow one to think precisely of the simplifying assumptions needed to meet a particular purpose. More to the point, thinking about computational models of comprehension and production is a very powerful heuristic tool, forcing us to make precise our intuitions about meaning.

© Cambridge University Press www.cambridge.org

CAMBRIDGE

Cambridge University Press
978-0-521-10918-5 - Language and Spatial Cognition: An Interdisciplinary Study of the Prepositions in English
Annette Herskovits
Excerpt
More information

1 Meaning and use of locative expressions

Although trying to build computational models of comprehension and production is a good way to start thinking about meaning, there are some legitimate doubts about the soundness of the computational metaphor itself. I will address these, as they have some bearing on my general method and approach, before introducing the central questions guiding this study. One attempt to answer these questions in the domain of locative expressions, widely employed in artificial intelligence and linguistics, uses a "simple relations model", but this, I will show, fails to account for the data in many ways. Examining those deficiencies leads to an outline of the descriptive framework adopted in the book. But to begin, a precise definition of just which expressions are considered will be helpful.

1.1 Locative expressions

A **locative expression** is an expression involving a locative prepositional phrase together with whatever the phrase modifies (noun, clause, etc.).

The simplest type of locative expression is composed of three constituents: the preposition and two noun-phrases as in

the spider on the wall

The spider is the **subject** of the preposition, and *the wall* is the **object**. The subject refers to the **located entity**, the object to the **reference entity**.

A locative expression may be structured around a copulative verb:

The spider is on the wall.

or an existential quantifier:

There is a spider on the wall.

Mostly this book is about locative expressions of these first three simple types. But there are also locative expressions where the subject of the preposition is a clause:

Louis is eating snails in the kitchen.

© Cambridge University Press www.cambridge.org

CAMBRIDGE

Cambridge University Press
978-0-521-10918-5 - Language and Spatial Cognition: An Interdisciplinary Study of the Prepositions in English
Annette Herskovits
Excerpt
More information

Although most commonly that sentence is assigned the constituent structure:

[[Louis] [[is eating snails] [in the kitchen]]]

I will take the subject of the preposition to be *Louis is eating snails*, as this clause describes the event that is taking place *in the kitchen*.

A spatial prepositional phrase may also fit in the case frame of a verb:

She put the bread in the hutch.

Here the most appropriate constituent structure is:

[[She] [[put] [the bread] [in the hutch]]]

What is being located *in the hutch* is not any object, state, or event, but the "destination" of the action "to put", that is, the end place of the trajectory of the bread, which will be considered the located entity.

Prepositions fall into two categories: some are primarily static (e.g., *at, in, under*); others primarily dynamic (*to, from, via*). But static prepositions can be used in dynamic contexts (*The cat ran under the bed*), and dynamic ones in static contexts (*The lamp is two feet from the wall*). By focusing on locative expressions of the first three simple types, I exclude an in-depth study of dynamic prepositions, or of dynamic contexts for the static ones, but much of what holds true in the simple static cases carries over to dynamic contexts. Usually, some aspect of the action denoted by the verb involves a part of space, which is located with respect to the reference object by means of the preposition exactly as any static physical object would be. Consider for instance

She ran behind the fence to the road.
He writes only on blue paper.

In the first example, a directed line segment – the trajectory of the run – is located *behind the fence*, and its end is on the road. In the second example, the trajectory of the tip of the moving pen is located *on the blue paper*; so is the resulting writing, the trajectory being in fact coincident with the space it occupies. In these and many similar examples, the preposition simply relates the relevant trajectory to the chosen reference object, as it would two ordinary objects; there are, however, examples where this analysis fails, and where special, verb-dependent rules apply.

Basically, this study deals with the use of locative expressions in American English, more specifically West Coast American English. I have generally relied on my own linguistic intuitions, though I have checked with half a dozen people when those intuitions wavered, and note when conflicting opinions appeared.

For detailed study, I have selected two classes of prepositions: the three

© Cambridge University Press www.cambridge.org

CAMBRIDGE

Cambridge University Press
978-0-521-10918-5 - Language and Spatial Cognition: An Interdisciplinary Study of the Prepositions in English
Annette Herskovits
Excerpt
More information

"basic topological prepositions", which are *at*, *on*, and *in*, and the projective prepositions, which are fundamentally derived from the experience of viewing, such as *in front of*, *behind*, *to the right*, etc. The first group is treated in Chapter 9, the second in Chapter 10.

1.2 The philosophical basis of the computational metaphor

When research in artificial intelligence began, there were forecasts that computers capable of fully intelligent conversational exchanges would soon exist. Even when computer scientists had to give up these unrealistic expectations, they attributed their failure to quantitative complexity; most still held that a computer program could in principle fully perform as a language user, that some well-designed algorithm manipulating physical symbols could fully explain linguistic ability (Winograd 1976; Johnson-Laird 1977; Woods 1978). One dissenting voice was that of Dreyfus, an interpreter of Heidegger, who argued that no algorithmic procedure could account for intelligence or language, and that knowledge could not be exhaustively formalized (1979). By 1980, a number of computer scientists had become more cautious about identifying intelligence with symbol manipulation, and some, such as Winograd (1980), even rejected the idea.

Some of the principal arguments against the computational metaphor involve the notion of "background" (Searle 1979; Dreyfus 1979; Winograd 1980; Winograd and Flores 1985). Very briefly, this holds that what a sentence conveys will vary from context to context, according to the purposes and assumptions of the speaker; these constitute a background that cannot be exhaustively formalized as there is no absolute way of breaking it down. Any one of many possible ways would reflect the biases – the purposes and assumptions – of the analyst.

Another set of arguments centers on the notion of mental representations, and whether they can be said to underlie language use, and to represent the world. Winograd and Flores, for instance, claim that one cannot assume that a language user has mental representations of the regularities observed in language use; that language comprehension and production do not proceed by manipulating representations of such regularities. They argue further that cognition is not based on the manipulation of "mental models" or representations of the world; a sentence is not a description of a state of affairs in the world. Following Maturana (1978), they describe the nervous system as a closed system: "...the system can do only that which is determined by its own structure and activity – its action cannot be understood as a reflection of an external world it perceives". We tend to accept the idea that instinctual behavior is not based on a representation of the world, but on a set of external conditions that cause a certain perturbation to the nervous system.

© Cambridge University Press www.cambridge.org

CAMBRIDGE

Cambridge University Press
978-0-521-10918-5 - Language and Spatial Cognition: An Interdisciplinary Study of the Prepositions in English
Annette Herskovits
Excerpt
More information

However, when it comes to cognition, we feel certain that our mind holds a representation of the world. Yet this certainty is unfounded.

Without attempting a full discussion of these ideas, I want to place the analysis proposed in these pages in relation to them, because I find them persuasive.

I will speak of relating a locative expression to a situation type. A situation will be taken to include the current relevant purposes, concerns, and beliefs of the speaker or hearer – though "beliefs" should be interpreted broadly to include the representation of scenes provided by perception. If indefinitely many different complexes of purposes and beliefs may be seen as bearing upon the production or interpretation of a given sentence, it is impossible to give a systematic description of the relation between an utterance and a situation type, or to list necessary and sufficient conditions for the use of an expression. But one can point to some salient and commonly encountered relations between locative expressions and some articulated aspects of experience and purpose. And one can also sometimes discern extensions, processes which stretch the use of expressions beyond those central cases. I will talk of predicting the interpretation of a sentence in a given context, and of predicting the sentence appropriate to a given situation, using as instruments for such predictions the regularities that link locative expressions and situations; but strictly speaking I may only be describing the regularities that link locative expressions and the salient aspects of commonly encountered situations in which they would be used – aspects that present themselves as we imagine some of the ways in which a locative expression would be appropriately used, or on the contrary, ill-formed or misused.

Of course, any semantic description would have the same limitations. Yet, the study of regularities of this kind is worthwhile, even if those regularities remain somewhat intractable. Any student of semantics knows that any attempt to capture the meaning of a word always founders against the encounter of unforeseen uses. But language is not a free-for-all, and even if we do not now see the nature of its rules, there are constraints; without them, there would be no language and no communication. Looking at regularities, attempting to express them formally, and exploring where generalizations fail, is a necessary first step toward building a more adequate theory.

In short, while I share Winograd and Flores' disbelief in the assumption that we can model the use of knowledge by manipulations of symbols, I think such models have been revealing in patches. They also are the only rigorous means we now have of capturing some epistemological phenomena. Exploring further the computational description of language comprehension and production should enable us to more clearly under-

© Cambridge University Press www.cambridge.org