

AD-A085 750

ILLINOIS UNIV AT URBANA-CHAMPAIGN COORDINATED SCIENCE LAB F/0 S/7
COMPUTATIONAL INTERPRETATION: ENGLISH SPATIAL PREPOSITIONS. (U)

FEB 79 L C BOGESS

N00014-67-A-0305-0026


T-75

NL

UNCLASSIFIED

1 of 2

AD-A085 750



MICROCOPY RESOLUTION TEST CHART
NATIONAL BUREAU OF STANDARDS-1963-A

REPORT 1-15


FEBRUARY 1979

CSL COORDINATED SCIENCE LABORATORY

12

ADA 085750

**COMPUTATIONAL INTERPRETATION
OF ENGLISH SPATIAL PREPOSITIONS**



437

12

① COMPUTATIONAL INTERPRETATION OF ENGLISH SPATIAL PREPOSITIONS

by

Lois Carolyn Boggess

10 T-2

This work was supported in part by the Office of Naval
Research under Contract N00014-67-A-0305-0026

15
ALI 17

DTIC
ELECTE
JUN 18 1980
C D

This document has been approved
for public release and sale, its
distribution is unlimited.

09710

COMPUTATIONAL INTERPRETATION OF ENGLISH SPATIAL PREPOSITIONS

BY

LOIS CAROLYN BOGGESS

B.S., University of Mississippi, 1969

B.A., University of Mississippi, 1970

M.A., University of Mississippi, 1971

THESIS

Submitted in partial fulfillment of the requirements
for the degree of Doctor of Philosophy in Computer Science
in the Graduate College of the
University of Illinois at Urbana-Champaign, 1978

Thesis Advisor: Prof. David L. Waltz

Urbana, Illinois

COMPUTATIONAL INTERPRETATION OF ENGLISH SPATIAL PREPOSITIONS

Lois Carolyn Boggess, Ph.D.
Coordinated Science Laboratory and
Department of Computer Science
University of Illinois at Urbana-Champaign, 1978

It seems clear to anyone who pays attention to the use of prepositions in language that any one preposition, when used to describe the spatial relationship between different objects can produce strikingly different mental models for different objects. The mental model produced by the description 'a bowl on a table' seems to be somewhat different from that produced by 'a poster on a wall' which in turn is somewhat different from 'a shelf on a wall' which again is different from 'a fly on a ceiling'.

It is the contention of this paper that the preposition in conjunction with a small set of features of the objects (mostly perceptual features) can account for such variations in spatial relations.

The thesis discusses a means of taking English-language descriptions involving prepositions and their semantic subjects and objects and deriving a three-dimensional model of the spatial relationships of the subject and object.

The relationship of some of the spatial prepositions to a coordinate system is explored, as well as canonical definitions for prepositions based on analyses of descriptions using "neutral" subjects and/or objects ("whatchamacallit", "you-know-what", and so on).

→ next page

Examples are taken from a simple program which accompanies the theory. The program is supplied with approximate descriptions of the shapes of a variety of objects. Each preposition in the program has one definition (e.g., there is only one procedure for on, rather than several--ON1, ON2, ON3, and so on); in general the definition is made up of several components, each of which is responsive to a perceptual characteristic of the semantic subject or object.

✓ The program takes extended descriptions involving many objects, each of which is incorporated into the overall model. Once an object has been described, it is possible to interrogate the model about the relation of that object to any other in the model, without recourse to inference rules of the following kind: "if A is on B and B is in C then A is (probably) in C." ↗

Accession For	
NTIS GRA&I	☒
DDC TAB	☐
Unannounced	☐
Justification	
By _____	
Distribution/_____	
Availability Codes	
Dist	Availand/or special
A	

TABLE OF CONTENTS

	Page
Chapter one	
INTRODUCTION	1
THE COORDINATE SYSTEM	2
OTHER PRIMITIVE CONCEPTS	10
Chapter two	
THE PREPOSITIONS	21
Semantic subject and object	25
Prepositions and motion	31
Verbs of projection	35
Geographic objects	37
IN	41
Embedding	47
Two-dimensional interior	50
Geography and contact locatives	55
Some last thoughts	56
ON	61
Departures from the prototype	72
On without support	75
In/on: interchangeability	79
TO AND FROM	82
The "legal" to	88
To invokes scenarios	90
Some particular uses of <u>from</u>	92
AT	96
MOTIONAL IN AND ON--INTO AND ONTO	101
TOWARD	104
NEAR	105
OTHER PREPOSITIONS	108
Chapter three	
THE IMPLEMENTATION	114
Program data	114
Input sentences	117
The prepositions and primitive relations	120
Some examples	123
Questions	130

	Page
Chapter four	
WHAT CAN WE CONCLUDE FROM ALL THIS? . . .	133
Perceptual attributes	133
Location restrictions	137
Inferencing	136
Regularities	142
List of References	146
Appendix	
A WORD ABOUT PRIMITIVES	149

Chapter one

INTRODUCTION

It seems clear to anyone who pays attention to the use of prepositions in language that any one preposition, when used to describe the spatial relationship between different objects can produce strikingly different mental models for different objects. The mental model produced by the description "a bowl on a table" seems to be somewhat different from that produced by "a poster on a wall" which in turn is somewhat different from "a shelf on a wall" which again is different from "a fly on a ceiling".

It is the contention of this paper that the preposition in conjunction with a small set of features of the objects (mostly perceptual features) can account for such variations in spatial relations.


This is not the first attempt to account for the use of prepositions well enough to produce a model of the spatial relations of the objects described. A particularly elegant theoretical model was proposed by Cooper [1968], but an effort by the present author to translate the elements of Cooper's theory into a computer implementation was disappointing. Cooper did allow the preposition to make selection restrictions on the objects related (e.g., in Cooper's scheme the preposition at could not take a

geopolitical area as semantic object, so "the man at New York" would be disallowed); in retrospect it would appear that her major failure, aside from the vagueness of some of the terms, was a failure to allow the perceptual attributes of the objects themselves to affect the spatial relationships inferred.

THE COORDINATE SYSTEM

For a computer model of concepts of space, there can hardly be more fundamental primitives than the bases for the coordinate system used to record and interpret space relations. H. Clark [1973] proposes a rectangular coordinate system as being natural physiologically, psychologically and linguistically. He divides space with a horizontal plane at ground level, allowing gravity to distinguish the vertical direction, thereby providing one axis of the coordinate system; since gravity is asymmetrical, it provides a basis for a directed axis, and Clark argues in favor of the positive direction's being what humans call upward from ground level since in normal circumstances what is upward from the horizontal plane is readily perceptible (a "positive" characteristic) and what is downward from the plane at ground level is generally not readily perceptible.

Clark divides space with a second plane, also on the basis of asymmetry of perception: for a human standing on the ground at the vertical axis, he proposes a vertical plane separating that which is in front of the person from that which is behind him. The intersection of this plane with the horizontal one provides a second axis for the coordinate system, and as before, since what is in front of the plane is easily perceptible to eyes and ears, and what is behind the plane is less readily perceptible, the positive direction is defined to be what humans would call forward. In addition, he argues that forward is psychologically positive because it is the normal direction of motion, perhaps in part because it is optimally perceptible.



A third plane is defined perpendicular to the first two, dividing space (on the basis of symmetry this time) into left and right.

On the basis of Clark's recommendations, then, the model discussed in this paper was originally intended to use rectangular coordinates. However, it is not unusual for people to specify a distance relationship without a direction ("The arrow missed his head by inches." "We drove 600 miles today.") or a direction without a specific distance ("Look beyond that oak tree, a little to the left."). Of course, it would be possible to separate distance from direction by the simple expedient of defining the direction of interest to be one of the coordinate axes. But in general, in rectangular coordinates, distance and direction are not cleanly separable, and people seem to speak of one separately from the other fairly often. In the two-dimensional case, for example, if the direction is explicitly to be left undefined, but the distance d is known, both x and y coordinates remain undetermined: $x = \sqrt{d^2 - y^2}$ and $y = \sqrt{d^2 - x^2}$. If on the other hand one were using cylindrical coordinates, the distance r would be known and the direction would be explicitly undefined, a somewhat more straightforward representation of what is known and unknown.

The special status accorded left and right in rectangular coordinates can be conferred in cylindrical coordinates to the special angles $\theta = \pm 90^\circ$. "In front of" and "in back of" correspond to $\theta = 0^\circ$ and $\theta = 180^\circ$. In cylindrical coordinates "to the left" (as opposed to "to the left of...") has the interpretation of negative θ , with a default

of -90° but capable of further specification, as in "turn a little to the left" which explicitly specifies small θ . Rectangular coordinates have no such natural interpretation.

Which system do people use? It seems likely that they use both, and use both readily.

So far, a spherical coordinate system has not even been discussed. In fact, the same arguments used in favor of cylindrical coordinates over rectangular coordinates could be used to argue for spherical coordinates rather than cylindrical. Cylindrical coordinates would have slight advantages when objects with particular heights or at specific altitudes are part of the discussion; spherical coordinates would have the advantage where heights and altitudes are variable, or the sentence is to be interpreted as a search instruction ("See that star just above the moon?") where what is important is not so much literal position as eye movement.

Often it is the case that a description of position is used by a hearer to determine how he should move his eyes to see the object under discussion, so it should not be surprising that many descriptions that appear to be defining positions lend themselves more readily to a search-procedure interpretation than to a literal-position interpretation. In a situation where a child is playing with a toy mouse and a cat, the phrase "the mouse above the cat's head" signals a relationship between the literal positions of the mouse and

the head which does not hold for the referents in the star/moon example. Yet both examples work equally well for search-procedures: "locate the head (or moon) and then raise your eyes (increase the angle from the horizontal plane) until you find the mouse (or star)."

A good system, then, needs spherical coordinates as well. As it happens, a system that performs well on the toy mouse and cat example should not only be able to extract the spherical coordinate interpretation, but also record some bounds on the probable height of the toy mouse after interpreting the phrase. (Exactly what those bounds are is a subject for later discussion.) Absolute height is a cylindrical or rectangular coordinate or a spherical coordinate with special angle $\phi=0$.

Whichever coordinate system is used, it is apt to be a relativistic one, in the sense that there is no absolute origin. In the sentence "I met Anne at church," the church is treated as local origin, and where the church is relative to anywhere else is assumed to be either known or irrelevant. In fact, in many cases the relative position of what I have called the "local origin" is known to the speaker and hearer, and is probably stored in conceptual memory relative to the speaker/hearer's current location or any of several habitual locations (home, office, schoolroom [for children], eating place, regularly visited locations, frequently visited locations, and so forth, roughly in

descending order of probability, where, as usual, people have no trouble switching from one to another). The position of the local origin relative to some more permanent origin can be called on, if necessary, but very often is not required.

Prepositions signal local origins. In the sentence "A bird is in a cage," the object of the preposition in is the local origin, or point of reference, and the position of the bird is given with respect to the cage. In "The cage is above a table," the local origin, again signaled by the preposition, is the table, and a proposition is made about the location of the cage relative to this new local origin. In "Across the room from the table is a bookcase," the across-from preposition pair again signals the table as the point of reference for the position of the bookcase, and so on. Generally only one local origin is given at a time, and people can string them together to ascertain the spatial relationships between two objects whose positions with respect to one another have not been explicitly stated, as for instance the bird and the bookcase in the example above.

All three coordinate systems have two axes in common: the vertical axis defined by gravity and the axis which Clark defined on a biological basis--the "forward" direction which is the most favored direction for perception and the direction of normal movement. "In front" means in this forward direction from a human. The definition can apply to

non-humans and even inanimate objects. "In front of" an animal has the obviously parallel definition. "In front of" a car is defined by the "perceptual apparatus" of the car--headlights, windshield wipers, and so forth, and the direction of normal motion, as well as the sense of "in front" of a person in normal position in a car. "In front" of a desk is defined to coincide with "in front" for a person in normal position at the desk.

So the sentence "The ball is behind the car" has the obvious interpretation--the ball is at a position in the opposite direction from the forward direction of the car. Yet there is a second interpretation for the same sentence, as well as an appropriate meaning for "the ball is behind the tree" where trees have no marked forward direction. To explain this we appeal to what Clark calls a canonical encounter. Suppose a speaker and listener are standing facing one another but somewhat separated. Either person may serve as the origin of a coordinate system. For an object between the two persons, the speaker has the option of saying "the ball is in front of me" or "the ball is in front of you." For an object that is behind the speaker, he is far more likely to say "the ball is behind me" than "the ball is in front of you" and similarly for an object behind the listener. So "in front of," in the case of the canonical encounter, seems to be reserved primarily for positions between the two persons involved. Curiously, this usage is extended to include "canonical encounters" between

a person and another object. "The ball is in front of that tree" can mean that the ball is between the speaker and the tree, as if the tree were in canonical position with the speaker and hence "facing" him. The meaning of "the ball is behind the tree" then is the obvious counterpart, even though trees do not have "backs." Moreover, such a derived "front" and "back" for an object can override the ordinary front and back. Cars do have conventional fronts, yet "the ball is behind the car" can mean to a searcher "on the other side of the car" regardless of the car's orientation.

In sum, the coordinate handling system is the most basic primitive in a computer model of space. It provides a natural interpretation for the concepts "up", "down", "in front of", "in back of", "left" and "right". Since humans seem to be able to switch between rectangular, cylindrical and spherical coordinates as the situation demands, a proper model to handle natural language use should also be able to operate in all three modes and translate freely from one to another. Furthermore, a good system must be able to change from one local origin to another, and at times to handle more than one local origin and simultaneously different orientations of local coordinate systems.

OTHER PRIMITIVE CONCEPTS

For a number of primitive concepts this model appeals to other computer models whose primary emphasis has been on vision. For instance, in general the concept boundary as used in this model may be understood to refer to a sort of visual outline of whatever it is whose boundary is under discussion, much the same as the sorts of visual outlines of objects vision programs are producing. There are exceptions to this general rule, of course, but even the exceptions tend to be an extension of the rule. In some cases, for instance, "the boundary" of an object is defined to be the visual outline of an object as viewed from a particular vantage point, or, if you will, its boundary relative to a marked (as particularly relevant) plane. For example, the boundary of a field would be defined as the boundary relative to the horizontal plane at ground level (with a few assumptions about the ground being horizontal and so forth), which is an extension of the original concept in that one might actually fly over the field and literally see the visual outline so defined, even though in actual practice one rarely does.

In some instances boundaries will be used which are defined but not perceptible. In other words, a few objects have vague boundaries which are nevertheless able to be treated as if they are the immediately perceptible kind. For example, in geographic discussions, people can speak of

geographic areas--"out of the plains." "in the mountains." "through the Midwest"--as if they were bounded areas and indeed if the persons were asked to provide a boundary, given a map, they probably could describe an outline. (They probably would also apologize that the outline was not exact, but on the other hand the outline is generally firm enough that if an observer generates large perturbations in it, a protest is likely to result.)

A second concept borrowed from vision programs is that which I call the center of visual mass. The concept is roughly what its name implies--given a two-dimensional figure, the center of visual mass is effectively where the center of mass defined by the laws of physics would lie if the figure were cut from a sheet of homogeneous solid material. A long, narrow perturbation of a compact, squatty figure would have only a small effect in changing the center of visual mass of the figure.

Yet another primitive taken from vision programs is the concept interior. In at least one vision theory currently under development, the interior of a region is derived from considerations of relative homogeneity of texture, color, and so forth [Waltz, 1978]. On this view, boundaries are determined by loss of homogeneity--changes in texture, for instance--and are in effect derived from the notion of interior, rather than the definition of interior being dependent on boundaries. Regions can be defined as

combinations of these homogeneous regions (e.g., a face would be defined as the obvious region plus the eyes and mouth, which would otherwise be excluded).

A face is, in fact, effectively a two-dimensional phenomenon--or at least a surface, as opposed to a three-dimensional object. In the sentence "A pie hit Mark in the face," we are told that the area of contact between Mark and the pie was within the two-dimensional interior of Mark's face. In can be one-, two-, or three-dimensional, but when used in conjunction with a semantic object that is in essence a surface, the two-dimensional characterization is invoked.*

So far, all the concepts borrowed from vision programs have been by nature two-dimensional, and there is a temptation to discount two-dimensional definitions since we live in a three-dimensional (at least three-dimensional) world, but in practice two-dimensional descriptions are often preferred over three-dimensional interpretations. Most geographical references are in two dimensions (in analogy with a map). And, as has been mentioned, for a number of objects there is a plane that is marked as being

* "On" can imply the same sort of two-dimensional interior relationship between a point and an area, as in "a point on a map" or "Mark slapped Bill on the back." More about this later.

the appropriate plane of reference--in the field example above, the horizontal plane is so marked. Not only can the explicit property of an object mark a particular direction or plane as relevant, but not infrequently, the description of a scene can do the same, either through use of particular prepositions, or through a combination of properties of the objects involved.

Two two-dimensional analogies have been most productive in this work: the "map" analogy already mentioned, and the "picture" analogy. As Piaget [Piaget and Inhelder, 1967] and Clark [1973] point out, the horizontal plane is a "given" of our real-world experience, from earliest infancy. It is our penchant as humans for locating objects in this basic horizontal plane (the ground level, if you will), without regard for vertical considerations, which gives rise to what I call the map analogy.

It would seem as if experience does not single out for us any other plane with the special status of the basic horizontal plane, but based on language use I have come to the conclusion that there does exist a vertical plane with special status. I am referring to the vertical plane of the visual field. Granted, our visual experience is not two-dimensional and there is no actual vertical plane "out there". Nevertheless, when in the presence of that which is being described or when describing a scene from a definite point of view we talk as if what we are seeing is a vertical

plane (like a picture). We can refer to an object to the left or right of another or above or below it when in fact it is the visual image of the object which is left, right, above, or below the visual image of the reference object. Consider the phrase "the moon over Miami," where it is probably not intended that the moon is physically directly over the city. or recall the "star just above the moon" example from earlier in this text. In describing apartments we may refer to "a girl in a window". Clearly, it is the visual image of the girl which is two-dimensionally interior to the boundaries provided by the image of the window.

Having shown that a two-dimensional notion of interior is not without importance, let us discuss the uses of a three-dimensional interior. Since the three-dimensional objects we consider will be assumed to be fairly regular, a word about less regular figures is in order: convoluted three-dimensional figures appear to be taken care of, for the most part, by one of three techniques--1) approximation by a more regular solid, 2) focussing attention on only a portion of the object (hence simplifying boundaries) or 3) having recourse to the map analogy, which is two-dimensional. (It would appear that people do not like to have to keep in mind complicated three-dimensional boundaries, either.)

By a "regular" solid is meant an object at least regular enough that the following terms are meaningful in delineating parts of the solid: "top", "bottom", "front" (if no canonical front exists, the the front implied by the presence of an observer), "back", "center" and combinations like "center left side", "top right", "left front", and so on.

So far then, the notion of three-dimensional interior has required only the capability of handling fairly regular objects (and a great many of the containers with which man has surrounded himself are very regular, from soup-bowls and tissue-boxes to rooms and stadiums. Others, like buses and volumes delineated by the outlines of the branches of trees, seem to be well handled by approximations to regular "solids"). There is another type of three-dimensional interior which requires somewhat different handling--the notion of interior with respect to an environment (e.g., "in the rain", "in the grass", "in the ocean", "in this weather", "on a hot day like this", "put the peas in rapidly boiling water").*

* Under normal circumstances, this last example implies interior to a volume (provided by an implied container), as well as to the specified environment.

There are some very interesting questions about what constitutes location and what constitutes environment. A reference to the human brain can be in terms of locations, with respect to regions and landmarks, for which terms like "front", "bottom", "left", "center", and so forth are meaningful; but the brain may also be referred to as a biological environment, in which the surroundings of cells have certain characteristics and in which connotation the size and shape of the brain are beside the point. Similarly, the Mohave Desert can be both location and environment. Even the French Quarter of New Orleans, or for that matter a McDonald's Restaurant, can emphasize environment rather than location. These latter examples are actually a somewhat broader interpretation of the word environment than we wish to use for our present purposes, having more of a sense of "context" or "scenario". For the present, this model merely points out and handles somewhat separately the use of spatial terms for "spaces" whose boundaries are considered not to be of interest.

Two related distinctions would well be mentioned here: consider the difference between "there is soup in the bowl" and "there is a crack in the bowl." In many cases enclosures--rooms, bowls, boxes, desk drawers, fences--are treated as delimiters only, defining boundaries of space. Their substance is of little interest and can safely be ignored. (This enclosure type of interior correspond to the presumed "normal", or unmarked, three-dimensional interior

of the model.) The crack in the bowl, however, if it is thought of not as a discontinuity of a surface but as pertaining to the substance of the bowl itself, is an instance of embedding. Another example is that of a window in a wall, where the window is interior to the boundaries of the wall in a two-dimensional interpretation (as is a poster on the wall) but where the window is also embedded in the substance of the wall.*

In general, everything interior to an environment is embedded in the environment. However, if environment is reserved only for examples of things for which boundaries are of little or no interest, not every embedding involves an environment, since there are many instances of embedding in well-bounded objects.

So then we have two-dimensional interior, "normal" three-dimensional interior (enclosure in a volume), embedding (where the substance of the immediate surroundings is involved in the interior relationship), well-bounded or assumed-well-bounded objects, and environments (objects

* As another example, a bird in a tree is interior in the normal three-dimensional sense to the volume delimited by the outlines of the branches of the tree. A nail in the same tree is embedded (or at least a portion of it is) in the substance of the tree.

whose boundaries are inconsequential)).

For some objects the concept of a particular kind of planar surface is important--a plane with a "free" side. A body of calm water provides such a free surface, where the side in the vertically upward direction is the free side. A wall of a room has a free-surface, with the free side being the side toward the center of the room. In actuality, any free surface (not just planar ones) can be important, because our perceptions of an object are generally of the free-surface of the object (what we see and touch). Moreover, we seem to surround ourselves with objects whose free-surfaces are of more than ordinary importance--perhaps because we tend to put things on them. At any rate, in many instances, treatment of a physical object with a relevant free-surface can be simplified by considering the object as if it were only a surface for as long as the situation description allows.

Contiguity, in this model, has a fairly exact interpretation: if two objects are contiguous, then portions of their surfaces are touching--that is, in whatever coordinate system is in force, portions of the surfaces of each have identical coordinates. The relation is identified by the name CONTIG; it does not necessarily have to hold between physical objects only (a shadow can be CONTIG with a wall and a design can be CONTIG with a plate); when both objects related by CONTIG are physical, it is

often the case that the surfaces that are contiguous are part of the characteristic free-surfaces of the objects (the outer surfaces of the objects that we would visualize if we formed a mental image of them), but this is not always the case--if one object is embedded in another, then the characteristic free-surface of the embedded object is CONTIG with a part of the other object that is not normally part of the latter object's characteristic free-surface.

For a number of prepositions I make use of a concept which I call the horizontal cross-section of an object. This is potentially an ambiguous term, since most three-dimensional objects have different cross-sections at different heights between top and bottom. What I intend would perhaps be better described as the vertical projection of the object--its shadow, if you will, from a source of parallel light directly above it. Moreover, the horizontal cross-section is almost universally used in conjunction with the object's location, so that rather than being just a description of an object's dimensions, it tends to be used to obtain location restrictions.

And finally, I have borrowed one concept from current vision research being carried on at the University of Illinois by George Hadden and Dave Waltz. My own term for the concept is "limited axis of symmetry"; like the "prairie-fire" technique, it is concerned with finding "skeletons" of objects. The "skeleton" of a circle is the

point at its center. The "prairie-fire" skeleton of a rectangle is as shown in Figure 1 (a) and the Hadden-Waltz result for the same rectangle is shown in Figure 1 (b).




Figure 1:(a) Prairie-fire technique (b) Hadden-Waltz technique

What is interesting about the technique for the present work is that it also yields "symmetry axes" for objects which are not symmetric in the mathematical sense (Figure 2); it is this more generalized sense of axes of symmetry which seems to correspond to one aspect of our treatment of space relations in language.




Figure 2: Skeletons of non-symmetric objects using the Hadden-Waltz technique in two modes

Chapter two

THE PREPOSITIONS

Prepositions are one of the few closed classes of English words: new nouns, verbs, adjectives and adverbs enter the language regularly and meanings of nouns, verbs, and so forth can vary with the passage of time. Prepositions, by contrast, remain fairly stable. There are about fifty of them, and all but a few have strong spatial connotations (for is one of the exceptions and even it can be used spatially: Judy raced for home base), although they may be used in contexts where a spatial interpretation is at best a weak metaphor and perhaps not applicable at all. Nevertheless, spatial interpretations are so strongly associated with the prepositions that not so many years ago the prepositions were often described to grammar school children first learning the parts of speech by the following heuristic (or a variation): a preposition describes anywhere a squirrel can be in relation to a tree--for example, a squirrel can be in a tree, on a tree, under a tree, can run around a tree, through a tree, up a tree, to a tree, away from a tree....

The prepositions are first used by English-speaking children relatively late in the language-learning process--not surprisingly since prepositions generally require not only an object but a semantic subject also;

prepositions generally express a relationship between two physical objects or an event and an object. If the preposition is included in the utterance, the minimum length of utterance is 3. (For example, the young child stepping toys on a toy castle's staircase says "Go up stairs" and "Weeble down stairs"--including preposition, semantic subject, and object.) Adjective-noun, noun-noun, noun-verb, and verb-adverb combinations can all be expressed in utterances of length 2. Slobin [1973] points out that a locative intention can be present in a two-word utterance ("pot stove" can be glossed "the pot is on the stove") and he credits children with having the conceptual apparatus for simple locatives quite early. However, it is evidently extremely expensive computationally for children of this age to translate an idea into a surface expression, with the result that the length of surface expression is severely limited [Feldman, Goldin-Meadow, and Gleitman, 1977] and as Feldman, et al., point out, very young children are chiefly interested in expressing events or mands, anyway.

When children do express locatives there seems to be a pattern of development discernable across cultures and languages [Slobin, 1973]. As Slobin points out, some languages--Hungarian, for example--express locative and directional information in a regular, simple system of suffixes. Bilingual children fluent in such a language and another language in which locatives are more difficult to express not surprisingly express locatives in the simpler

system long before locatives show up in the less simple languages. The point Slobin stresses is that the children obviously have the concepts long before they use them in the latter language. Since the bilingual children develop locatives in the less simple language at roughly the same ages as monolingual children of that language community. Slobin argues that the monolingual children probably have the locative concepts well in advance of their ability to express them, also.

Slobin cites international. cross-linguistic evidence for the following order of development of locatives: 1) simple topological notions (in and on), 2) notions of dimensional or Euclidean space (such as might be derived directly from our previous discussion of coordinate systems--"in front of", "up", "down", and so forth) and 3) more complex spatial notions such as along and through.

In and on appear to be nearly universally the first locatives. The next most likely concepts to appear are directional contrasts to in and on--into, onto, to or toward, from, and through being the most common, and then other locatives in no particular order.

In English, with no convenient inflection, locatives are expressed relatively late in the language-learning process. In and on are as usual the first to appear, followed by the gamut of prepositions in no clear order, as if the concepts are already there and, the prepositional

trick having been learned, the new linguistic capability is quickly extended to the other concepts.*

Since it seems plausible that the general order of appearance of spoken locatives parallels the (perhaps earlier) development of locative concepts, an attempt is made in this thesis to explicate the prepositions in roughly the order that they or their counterparts appear in young children cross-linguistically. In this way, cognitive concepts basic, for instance to in and on are available for incorporation into later locatives. The major divergence from the order roughly sketched above is that Slobin's suggested 3-stage development is not followed to the letter, since it cannot be said that the concepts falling out from the coordinate system follow the development of in and on. However, if the non-prepositional use of the surface-words which serve as English prepositions is taken into account, ["Look up!" or, while patting an abrupt discontinuity in the floor, "Down"], up and down may rank in importance with in and on. (The evidence in E. Clark [1972] would tend to support this conclusion also.) And the other prepositions which take most of their meaning from the coordinate

* There is some evidence (Brown, Cazden and Bellugi [1973], and E. Clark [1972]) that the accuracy of usage of prepositions by young English-speaking children parallels the order of development of locatives in expressions by children from languages with inflectional locatives. That is, in and on are used with a high degree of accuracy, followed by common directional prepositions, followed--with still less accuracy--by the rest of the prepositions.

system--in front of, behind, left, right--are in no way prerequisite to the development of in and on in this model.

Semantic subject and object

English syntax requires what is called the object of a preposition. If a real or implied object is not in the sentence, the word is being used as something other than a preposition--perhaps a particle or an adverb. So a preposition naturally focusses attention on its syntactic object, which is usually the source of a relationship. In the spatial uses of prepositions, at least, two things are required by the relation: the object serves as the basis of the relationship and the preposition serves to state the nature of the relationship of the other thing--which for convenience I call the subject, or the semantic subject--to the object. In the phrase "the dog on the steps." the dog is the subject and the steps the object of the on relation.

Virtually every spatial preposition has a prototypical sense which presupposes that the subject and object are physical objects (PHYSOBs). The objects of the relations, however, frequently have particularly salient characteristics which are emphasized, with corresponding de-emphasis of the rest of their characteristics (not uncommonly the fact that an object of a preposition is a PHYSOB and hence has weight and requires support is totally irrelevant). The object of the preposition serves as a sort of local origin with respect to which the subject is located

or a path (if motion is involved) is oriented--e.g., "the bird flew to the tree." In many instances, particularly with in and on, the dimensions of the object of the relation put bounds on the location of the subject.

Precise specification is rare, probably because ordinarily precise location is irrelevant. For example, in a search-procedure description (telling someone how to find an object) a series of major landmarks are used: "the keys are in my purse" suffices if the searcher knows the purse is habitually kept in one place. If more detail is warranted, a suitable description might be "in the zippered pocket of my purse, on the bottom shelf of the bookcase next to the bed" (by inference in the bedroom). This is actually not a very precise location of the keys unless the bedroom is well known; but it is a good specification of search procedure, even to a stranger to the room. The bed might be anywhere on the floor of the room, the bookcase adjacent to any side of the bed, the purse anywhere along the length of the shelf, and the key in any part of the pocket. But at any point in the search, landmarks are provided sufficient to drastically limit the perceptual task. Further description might well take up as much time as it would save in the search. In a search procedure, then, even though the real-life precise location of a physical object is the goal, what is supplied by the verbal description is a sufficient number of landmarks to yield the precise location when the search is carried out.

In most non-search-procedural cases, the precise locations of objects are not even of interest. Consider "I have this friend who just loves flowers, but he lives in an apartment in the city. So he always kept a pot or two of geraniums on the edge of his balcony, where they could get a little sun...." The location of the city, the placement of the apartment within the city, the placement of the balcony relative to the apartment, the size of the balcony, and the placement of the pots on the balcony are left to the imagination of the hearer, so long as the hearer puts the pots within a minimum distance of an open side of the imagined balcony. The speaker supplies only what is required for an understanding of the narrative (often to invoke context as much as anything) and whatever the hearer supplies as default is presumed acceptable.

The subject of a prepositional relation need not be a physical object, of course. Whatever is perceptible may be given location--e.g., "a creepy sensation on my arm." Moreover, any event may take a location--for example, "I think Jeffersonian Democracy is still workable in the twentieth century" might lead one to believe that events such as thinking do not take locatives, but consider "Judy thought she would die of embarrassment at Mark's party" or "I think my most sanguine thoughts in the herb garden." (Actually, the "at Mark's party" example is not quite fair, since a party is an event itself, with an implied

location--assumed to be already known or defaulting to Mark's home, in this case.)

When an event takes a locative, participants in the event generally take the locative also. Judy in the example above was at the party as much as the event of the thinking was.

Fillmore, in describing the locative case, broke it into four types--place, source, destination, and path [1971]. The latter three of course involve motion. "Jenni walked on the sidewalk," for instance, restricts a path to within the boundaries of a (perhaps default) sidewalk. (This may not seem very helpful unless one supposes knowledge, say, that interiors of boundaries of streets and sidewalks are disjoint, in which case, the answer to "Did Jenni walk in the street?" is "no".)

In English syntax, of course, a noun can refer to an event. "A walk in the park" is still an event, as is a picnic, a fight, a dance, or as noted above, a party; since the object of a preposition must be a noun or at least noun-like (pronouns, gerunds, and so forth), nomination provides the apparatus for events to be taken as objects of prepositional relations. They seem not to be readily taken as objects of spatial relations however. To and at are the most common spatial prepositions to take events as objects, and at commonly can be glossed as during or at some time during the event. When at and other prepositions are used

spatially, it is as if the event is being treated as a visually perceptible object. e.g., a sports event or the scene of a fire. Such events usually have duration; it is tempting to replace the event with what corresponds to "the scene of the event" except that the phrase "the scene of" refers to the location even after the event is over. To illustrate, "Judy and Paul raced to the fire" refers to the location of the fire during the time that the fire was still burning. "The scene of the fire." on the other hand, refers to the same location long after the fire has become history.

To repeat, while it is not common, some events are taken as objects of locative relations; but any event can be a subject of a locative relation. In fact, some events can have more than one associated locative. Consider the following:

a punch in the face

Jim punched Bill in the face.

Jim punched Bill in the coliseum.

Jim punched Bill in the face in the coliseum.

Liz dented the car on the fender.

Liz dented the car on the showroom floor (where either the car is identified as the one on the showroom floor or the location of the denting was the showroom floor).

Amy ripped her dress in the hem at the fair.

First of all, compare that last example with "Amy lost her penknife in the sideshow at the fair." The latter is a specification of one location: the sideshow at the fair. The former is not--"the hem at the fair" doesn't make much sense under normal circumstances. The hem is the place on the dress where the rip was, and the fair was where the event of ripping took place. The dent and the punch are located on the fender and in the face, respectively, but the denting and punching took place on the showroom floor and in the coliseum. Consider this: when Amy lost her penknife, she was in the sideshow (participants in events share locations with the events), but in no way is Amy in the hem, or Liz on the fender, or Jim or Bill in the face.

I call these "contact locatives" because they first came to my attention with verbs of contact. They might just as well be called "part-whole" locatives, because they seem to specify parts of objects--the hem of a dress, fender of a car, face of an individual.

When a verb takes a contact locative, its nominalization sounds natural only with the contact locative, not with the regular locative: "a slap in the gluteus maximus" but not "a slap in the coliseum"; "a rip in the sleeve" but not "a rip in the sideshow". Nominalizations of verbs that do not take contact locatives accept regular locatives easily: "a fight in the coliseum", "meeting Jim under a bridge", "a walk on the sidewalk".

Prepositions and motion

As just mentioned, Fillmore [1971] distinguishes four locatives: place, which has been the focus of discussion thus far, and source, path, and destination, which of course apply to motion. While locatives of place, if needed for the understanding of an utterance, are virtually always supplied in conjunction with locative prepositions, locatives of motion need not be: consider the two sentences Larry hit a ball; John caught it. Hit may or may not entail motion of the object of hitting, but caught, with a physical object, ordinarily does.* The locative source of the motion was the point of impact between Larry or an implement held by Larry and the ball (or more generally, the source was approximately Larry's location) and the destination was John's location. The default path in this case is through the air.

* Caught as part of a fishing scenario is somewhat more complex--it probably was much like any other catching of an animate object, including specification of a "destination" at the end of a perhaps elaborate path; but since the actions involved in catching a fish with a fishing pole were so similar to those involved in "catching" an old boot or other inanimate object (which presumably is not in motion at the time of the catching) the word and scenario apparently have transferred to an otherwise inappropriate object for the verb. Notice that there is some humor attached to the notion of catching a boot in a fishing scenario, just as there is in a situation where a toddler runs and picks up a baseball left lying by older children after a game and gleefully announces "I caught it!"

The point is, no prepositions were required in the original description. Ordinarily locatives of place are supplied by prepositional phrases or adverbs (here and there, for example). Verbs like catch and throw specify motional locative information about their syntactic objects, and others, like leave ("he left the airport") furnish motional locative information about their subject without prepositions--not really surprising, since with verbs of motion, locative information is naturally more basic to comprehension; in some sense the syntax of the language seems to reflect this importance by including the locative information in the focal elements of the sentence--what Fillmore calls "projection" [in press].

There are verbs of motion which require an associated locative, like put (consider "she put the thingamabob ..."), which requires specification of destination, but it is more normal for a motion verb, like a stative verb, to have any or all of the three motional locatives optional. For instance,

The horses leaped out of the starting gate. (source)

Mark leaped into Daisey's arms. (destination)

I leaped out of the fryingpan into the fire. (source and destination)

The dog leaped over the hedge. (path)

The deer leaped over the fence into the underbrush (path and destination)

Notice also that "the deer leaped the fence" is acceptable.

We have already seen verbs that elevate source or destination cases to focal position (syntactic subject or object); here is an example of a verb which elevates path to focal position.

In general we will restrict our attention to verbs of motion that are as neutral as possible, involving a minimum of world knowledge about default paths or default sources and destinations. Preferably the verbs under discussion will accept all three motional cases. "Get" and "go" are fairly neutral, at least in some of their senses ("John got out of bed", "once we get to Texas, the rest of the trip is on interstate", "rabbits got into our garden, under the fence" "we went from school to the fireworks display", "James went through the kitchen").

As with place locatives, path specifications are not exact--usually the hearer is provided with whatever landmarks or path restrictions are considered important. In "James dashed through the kitchen" we do not know what his intended destination is, nor where he has come from. We know that part of his path was bounded by the interior of the kitchen. A path can have multiple restrictions: "the children raced up the hill, over the footbridge, through the yard, around the barn and back to the apple-tree" is an example of a path made up of multiple paths, presumably in temporal order, followed by a destination. A composite path can also consist of multiple destinations; "we went from

Washington, D.C., to Philadelphia, to Atlanta, to Tampa, back up to Vicksburg, and finally home. (Notice that each destination by implication becomes the next source.) In either case, a global path is specified by a series of restrictions--in the first instance, the restrictions are all that are provided; in the second, the path is restricted by the assumption that in the absence of other information a path between source and destination may be thought of as essentially a straight line (a common simplifying assumption).

Of course, multiple path restrictions can be simultaneous, as well--for example "we went under the bridge through a culvert", where what is intended is not two path segments but two restriction descriptions for one path.

As Sondheimer [1977] points out, it is not sufficient simply to treat a path as essentially linear. Under normal circumstances an object is traversing the path and the dimensions of the object may enter significantly into interpretation of the description. If we hear, for instance that Alice went through the door, we assume that her path in the horizontal plane is restricted at some point to an interval bounded in that plane by the two sides of the doorway, but we also know that at some time the whole of Alice was interior to the doorway; consider what happens when there is a disparity between Alice's size and that of the doorway, as when the door is the door at the entrance to

Wonderland. If, knowing the door, we heard the simple sentence "Alice went through the door", we would have cause to raise some questions.

Verbs of projection

There is a class of verbs--look, glance, yell, shout, aim, and point, among others--which are not actually what we would normally consider verbs of motion (none of them are mentioned among the verbs of motion in [Miller, 1972]), but which take prepositions of motion, like into, toward, and at in its directional sense, and which behave in many respects like the verb throw except that no direct object of the verb is usually offered.* In fact, these verbs are unlike a good many of the verbs of motion, since many of the latter do not take at--to go at a location, or run at a place or object, or move at, walk at,...all sound at least strained, if not outright unacceptable. Not all verbs of motion that take direct objects work, either--you can throw a thingamajig at a whatchamacallit, but you can't put a thingamajig at a whatchamacallit, and if you jump a thingamajig at a whatchamacallit, the at-phrase is being used as a place-locative, restricting the location where the event of

* although one can look daggers at someone, yell obscenities at them, or shout encouragement to them.

jumping took place, which is not the sense of at that these verbs share.

I call these verbs, verbs of projection (in spite of the fact that the verb project is not one of them) because they seem to act as if something is being projected--in the case of verbs like look, glance, stare, and so forth, the something appears to be visual focus; for yell, shout, scream, and so on, the something is the sound of one's voice; for aim and point, it appears to be the focus of attention, even though for these latter there may be an explicit direct object that is not the subject of the preposition--if you throw a gun into a room, the gun moves--if you aim a gun into a room, the path is the same, but the gun does not move along the path. The same holds true for pointing a finger or a stick at something.

While most verbs of motion tend toward the map analogy, the verbs of projection tend toward the picture analogy. The paths of ordinary verbs of motion are often placed with respect to the horizontal plane (usually by virtue of the physical objects involved). For verbs of projection the path is usually through space to a location in the visual field. Hence, for particular prepositions, say at or to, restrictions that would ordinarily be made with respect to the horizontal cross-section for most verbs of motion are made with respect to a vertical cross-section with verbs like point and look.

Geographic objects

We certainly don't perceive cities, roads, and rivers in real life as points and lines. In the presence of a road I see that it has considerable width, and a river has much more. On the other hand, we don't get to toy around much with these objects, and maps have been a part of civilization for a long time, and our experience with the majority of rivers and cities and highways (excepting those near which we may live) is with points and lines on paper. And the older we get, the more practice we have treating them as such. Besides, on the scale of countries and states, the dimensions of cities and all but the lengths of roads and rivers have shrunk to negligible proportions anyway.

Most of the time this "abstraction", if such it is, does not affect interpretation of a description except to simplify matters. The map analogy has to do with placing objects in the basic horizontal plane and anything having to do with the vertical frequently is totally ignored. Moreover, it is convenient in talking about several cities, say, to treat them as point locations without worrying about keeping track of their extents, for the time being. We talk about the distance between two cities, for instance, without worrying about which points in the two "real world" cities

we are measuring between.*

This sort of "dropping out" of one or more dimensions of objects can take place on a level somewhat less vast than the geographic level of countries and states. If one is going to a building several miles away, the one overwhelming dimension of the path--length--seems to make its width and the dimensions of the destination shrink away by comparison--at least until a later part of the discussion focusses on them.

Most of the time, the presence or absence of one dimension more or less seems not to make a lot of difference to the language of the description; sometimes, however, prepositions seem to be tuned to the difference. For example, in a relatively static situation "on the road" can

* I have always supposed map makers and road-sign builders somehow know where the "center" of a city is and use that, so that when I see a sign that says "Memphis 200 miles" I assume I am 200 miles from the center of Memphis and maybe 195 miles from a nearer edge of the city. I sometimes have wondered if it would be theoretically possible to be on the interstate within the city limits of San Diego and see a "San Diego 5 miles" sign go by. On the other hand, a civil engineering friend who has worked for the highway department now tells me that, in her state at least, the road signs indicate distance to the nearest point of the city (however that is determined) and that she thinks for road maps the distance can be between the two nearest points of the respective cities by a particular road, or between centers of cities, and if it really matters you need to know your mapmaker--all of which goes to show how we can get along on a simplifying assumption quite nicely, even if it isn't integrated into the "real world" exactly right.

mean literally on a road surface--"Mother nearly panicked when she saw little Jonathan standing on the road"--the dimensions involved are such that the road has width. In "Urbana is on I-74", none of the dimensions of Urbana are such that Urbana will fit on the road-surface of I-74. However, on a geographic scale, the length of I-74 is great enough that the highway becomes one-dimensional and Urbana shrinks to a point in comparison. There's no trouble at all with representing a point on a line. This has importance for the building of a rigorous model, however, since care must be taken whenever there is a large change in scale. If the model is told that city x is on highway y, and subsequently discussion involves points in city x, so that the city is no longer a point location but has extent, then it is obviously a mistake to infer that every part of the city is "on" the highway in question.*

Before going on to other subjects, it is at least worth mentioning that there seems to be a tendency for paths in general, not just long ones, to be taken as lines in a plane rather than traces of three-dimensional objects in motion. This may be more of a comment on how we use language than on how our minds work--if you look back at the argument in

* In fact, in the example cited, no part of Urbana is in fact within the boundaries of the actual highway I-74, since it passes the city outside the city limits.

favor of three-dimensional traces for paths of moving objects, you'll see that it appeals to extra-linguistic considerations to make its point. (There remains an open question in my mind to what extent a theoretical model based on language descriptions should be more complex than is required by the language used in the description. The answer lies somewhere between "not at all" to "that's where all the interesting stuff is taking place!") For our present purposes, I simply point out that we clearly have to have the simplifying apparatus to handle some cases (large contrasts in sizes) and apparently we use the same linguistic terms for other cases--which may be why linguists have been able to discuss motion in simple terms.

IN


The preposition in has a number of uses in English. many of them not particularly spatial (consider the "in English" phrase of this sentence, or the phrases "I am in a hurry", "in praise of wine", "in retrospect", or "John is in business for himself", or the instruction "Mix the ingredients in the usual manner"). Even limiting the word to clearly spatial prepositional uses, one finds a variety of flavors to the inferences directly derivable from the word and characteristics of its semantic subject and object. Consider the differences in the girl in the painting, the girl in the mirror, the girl in the window, and the girl in the door; the relations supposed in soup in a bowl as opposed to a crack in a bowl; the contrast between what is meant by a tadpole in a puddle and a young child standing in a puddle. We speak of sitting in the shade (where shade is certainly not a container of the ordinary sort), and of a slap in the face, without undue concern about the manner in which the slap is in the face. (Is penetration implied?...What about a pie in the face?)

For all the varieties of relationships signalled by the word in in conjunction with characteristics of the things which are the subject and object of the in relation, which is the most central meaning, the "simplest", the one most likely to be ascribed to young children? What, if anything.

is imported by the word in when used to relate a neutral subject and object, as in "the you-know-what is in the whatchamacallit"? In the absence of context, the semantic subject and object have a tendency to be taken as physical objects. (They needn't be, of course; under suitable circumstances "the you-know-what" could refer to the Watergate scandal and "the whatchamacallit" could be the lead-in spot on the 5:30 news broadcast.) There is some anticipation, moreover, that the object of the in relation will be an enclosure--delimiting a three-dimensional space which, while not necessarily "empty" is at least not solid. A room is such an enclosure, as is a closet, a birdcage, a piggybank, a milk carton, or the volume delineated by the outer leaves or branch tips of a tree. (The tree, especially, indicates how non-empty the three-dimensional interior (INTERIOR3) of the enclosure may be; nevertheless, it is relatively easy for the INTERIOR3 of a tree taken as enclosure to be the location of physical objects--e.g., a bird in a tree, or a cat, or a child, a bear cub, a tree-house in a tree.)

The enclosures mentioned thus far are bounded in all directions, and their boundaries provide the bounds on the location of the subject of the in relation. There is a large group of enclosure-type physical objects which are used as objects of in relations which are not bounded in all directions. Typically the direction without bound is the upward direction; examples include open-topped boxes,

bowls, cups and glasses, pots, wagons, dresser drawers (which become completely enclosed when "shut"), baskets, swimming pools, wells, cupped hands, grocery bags, and so forth. Such objects do, of course, have three-dimensional interiors, one boundary of which is not material but which may be perceptually salient. If all of an object is in the INTERIOR3 of such a container, it is "in" the container; if most or even some of an object is located within the boundaries of such a container, it is "in" the container. In can even be used to describe a physical object that is completely beyond the non-physical boundary of the container in a vertically upward direction, provided it is on (in the prototypical sense) an object or objects in the interior of the container: for example, a doll resting on a heap of



toys in a toybox may still be described as "in the toybox" even though no part of it is interior to the box. The same doll in the same location, hanging from a string, is clearly not in the toybox. Of course, the doll on the heap of toys example is not a prototypical case of in. It appears to be

something of an extreme case where the strong relaxation on boundary constraints seems to be mitigated by near conformity to the constraints on in relations with open-topped containers and by a habitual relation.

Notice that a doll hanging at the end of a string, which is entirely in the INTERIOR3 of the toybox is unquestionably in the toybox, even though it does not touch the box and in no way is supported by the box or its contents. The only relation between the doll and the box is the location of the doll relative to the three-dimensional space defined by the box. If the doll on the string is raised so part of it is not located in the INTERIOR3 of the box, then the judgment of whether it is rightfully in the box becomes clouded. This is in clear contrast to the many cases where support for the subject comes from within the object, under which circumstances the upper boundary is readily relaxed to allow part of the subject to project beyond the upper boundary (consider groceries in a bag, flowers in a vase, fruit in a bowl, and so forth).

People use in with reference to spaces considerably more loosely bounded than open-topped containers; a ditch is not only open-topped but open-ended--nevertheless gutters, troughs, and such-like are readily taken as objects of in relations. Street gutters frequently have but one material side--the other is rather ephemeral, but we can still speak of things being "in the gutter". Likewise the

space "enclosed" by a chair--evidently delimited by the seat and back (certainly not the legs) and arms, if any, is at best vaguely bounded. This may represent something of a limit, however, since it is almost as natural to speak of sitting on a chair as in it (sofas seem to predispose toward an on relation), and the street gutter example may be a conventional holdover from days when streets were built slightly differently. It's difficult to find other, naturally occurring three-dimensional physical objects of in relations that are quite so loosely bounded.

The in/on phenomenon with chairs, where everyone knows the conventional spatial relationship between people and chairs, and where on is certainly appropriate--deserves a bit more consideration, and we shall come back to it shortly.

What we have seen so far is that in can signal as little as confinement of location to the INTERIOR3 of the object of the relation, with no other relations involved between subject and object. However, such a situation requires a certain amount of explanation and context setting, as in the doll on a string example. In the absence of such explanation, some kind of support and CONTIG relation is assumed. In the absence of elaboration, the supporting object may be assumed to be the object of the in, and the subject and object may be thought of as CONTIG, but both relations are somewhat loose, and easily transfer from

the object of the in to any physical object interior to it. (Moreover, as has been noted, when such support is present, the strict interior relationship is no longer required.)

Now, with support, CONTIG, and the horizontal free-surface associated with chairs, on is a natural candidate for a description of the relation of objects in normal spatial relations to chairs. Stereotypical chairs have backs and it may be that the chair-back is sufficient to alter perception of the horizontal free-surface and associated free-direction and imply occupation of volume instead. but there is at least one other possible explanation worth mentioning. As is argued elsewhere in this paper, in is sometimes used to indicate restriction of location in other than a three-dimensional sense (usually boundaries in two dimensions are applicable). Since the spatial relation of a person to a chair (in canonical position, so to speak) is well known, in may be used simply to indicate location bounded relative to the chair. "in the well-known manner". This explanation might be used as well for "in bed". On a hot summer night, people "in bed" may be in the bed only in the sense that their location is two-dimensionally interior to the boundary of the bed in the horizontal plane. People in bed are certainly on the bed proper (in the strongest prototypical sense), although they may be under sheets or blankets, propped up by pillows, or sitting up and reading. What seems to be implied is what they are wearing (conventional sleepwear), as much as

anything, and the implication that they are (or ought to be, or have been) staying in the location bounded by the bed (consider someone still in bed at 10 a.m., or a sick child in bed). The in is used, as with the chair example, to indicate a well-known relationship, with some bounding of locative, as well.*

Embedding

There is a second kind of three-dimensional in relation. Consider a carpenter constructing a box out of boards, nails, and glue. And consider a salesman at a hardware store, counting nails into a box. In either case the nails are "in" a box. But the latter example is an illustration of the prototypical three-dimensional in, and the former is a case of what I call embedding. The inferences made in each case are somewhat different, though related. Perhaps a direct comparison would be instructive.

* There are, of course, other plausible explanations--for the bed example, for instance, the in convention may have been inherited from European cultures in which beds were typically cabinets in walls.

It seems pretty clear that the preposition at has the interpretation "in typical relation to" the semantic object [Johnson-Laird, personal communication with Waltz], at least for locative uses not involving motion, so that a similar interpretation in another of the most common prepositions seems at least plausible.

1) Obviously, both in relations feature a three-dimensional interior, but with prototypical in the normal inference, in the absence of other information, is that the subject of the relation is located at the bottom of the object of the relation--on the floor of a room, for instance, or the bottom of a well or ditch--gravity being what it is. For a case of embedding, no such inference is justified--in the absence of specific information the embedded object may be located anywhere in the INTERIOR3 of the semantic object with equal probability.

2) Regular three-dimensional in may indicate a direct support relation between the subject and object, but admits of chains of support relations (a doll on toys in the toybox) on the one hand, or no support relation on the other hand (a doll suspended from a string). When one physical object is embedded in another, only a direct support relation is possible.

3) The boundaries of the object of an embedding in are the boundaries of the INTERIOR3 used to restrict the location of the subject of the relation, as one would expect. Containers used as semantic objects in prototypical in relations are abstracted to provide just such boundaries--a bowl, for instance, may be treated as an abstract surface in three-space marking a volume; the thickness of its material may be ignored. On the other hand, if a container is the object of an embedding in, the

nature of the boundaries of interest change: e.g., "Hey look! There's a bubble in this glass," which has two possible disjoint regions where the bubble may be located. In the embedding case, the boundaries of the material of the glass are the description of interest and a completely different INTERIOR3 results.

4) The differences between prototypical in and embedding in mentioned thus far are of perhaps limited importance; one applies to containers as a class. the others pertain only to physical objects embedded in physical objects and may not apply when semantic subject or object are not physical objects. But there is one characteristic of embedding which appears to apply to all subjects and (particularly) objects, having to do with the type of contiguity of the subject and object. For most physical objects and for the majority of spatial relations, contiguity (CONTIG) has to do with a portion of the surfaces of the objects touching (their coordinates coincide, in other words). The surfaces are the characteristic outer free-surfaces of the physical objects (generally part of the perceptual definition of the objects). The kind of contiguity implicit in embedding differs in two respects; first of all, the entire surface of the subject of the embedding in (or whatever part is embedded, in a partial embedding) is contiguous (shares coordinates with) the semantic object of the relation; secondly, the "surface" of the semantic object of the relation which is thus contiguous

to the subject is interior to the characteristic outer free-surface of the object.

In another section of this paper, it was pointed out that some objects of three-dimensional interior relations have no characteristic shapes--no boundaries--and these are designated environments. It happens, happily, that virtually every use of in with an environment as semantic object is an embedding, using the last-described kind of contiguity; in fact, the lack of a characteristic free-surface can be used to reason backwards that this kind of contiguity is required, and hence an embedding sense of in is being used.

Two-dimensional interior

Whether a regular three-dimensional in or an embedding in, all the uses of in discussed so far have used INTERIOR3 for any object having boundaries. There are objects with perceptual boundaries for which INTERIOR3 won't do. Consider the phrases: a girl in a field, seedlings in the garden, children in the yard, toys in the driveway, a couple in the park. While there are exceptions (a girl in a corn field in August may be three-dimensionally interior to the field), by-and-large fields are not three-dimensional. Nor do we seem to think of gardens, yards, driveways or parks as delineating volumes. It seems far more likely that the locations of the subjects of these in relations are bounded in the horizontal plane by essentially two-dimensional

boundaries (the horizontal plane being marked in the case of all the semantic objects mentioned above).

For normal subjects (three-dimensional physical objects), vertical placement is the default (for example, ground level if outside), but as with three-dimensional in, it is acceptable for the subject to be supported by something else in the semantic object--the children, for instance, might be climbing in trees in the yard, but they are still in the yard. There seems to be an implied appropriate height with each of these objects, beyond which the term in is not applicable, as witness bees in the garden, or birds flying in the park, so we never seem to get entirely away from three dimensions. However, this "appropriate distance" is something that crops up with virtually every space relation not implying direct contact, and it is the special nature of the subjects in these instances which counteracts the direct contact inference and adds the appropriate distance consideration. For relatively less mobile subjects, lack of contact (the doll-on-the-string type of situation) destroys the in relation, even within the "appropriate distances" of an essentially two-dimensional region. It is hard to come up with plausible examples on the field/park level, since most naturally occurring instances of a physical object's not directly contacting such surfaces involve supporting objects which do contact the surface, and we've already accented this variation. But imagine children playing a game of

marbles, with a circle drawn on the ground. A marble may be in the circle, but if it is lifted off the ground, there is something strained about continuing to say it is in the circle.

It should be pointed out that with essentially planar semantic objects, an essentially two-dimensional mental image (corresponding to a picture) is possible. With the girl in the field example, the mental image of the girl, or at least the image of her feet, is two-dimensionally interior to the visual boundaries of the field.*

Then, of course, there are actual pictures, which are undeniably two-dimensional, and we do speak of things as being in them. A certain amount of care must be taken, since a picture (painting or photograph) can cause a double representation; on the one hand, the picture and all images in it are two-dimensional, even though the images may be of three-dimensional objects; a girl in a horizontal field still may be vertically oriented; the image of a girl in a horizontal picture is in the horizontal plane. On the other

* In fact, very young children strongly avoid drawing objects in regions with any part of the drawn object obscuring the drawn boundary of the region; evidence presented by Piaget and Inhelder [1967] suggests that only by the ages of 7 to 11 can children be expected to include among representations of objects in regions some drawings of objects with "feet" in the region and "bodies" obscuring the boundary.

hand, regardless of the orientation of the photograph and the two-dimensional images in it, the mental model engendered can be just as three-dimensional as a model constructed from "real life". (It would appear that information presented by a camera, or a drawing in a children's book, say, is not notably more constraining than visual perception, especially a few seconds after the perception. Generally we use a large number of clues [size and obscuring relations, among others] to determine three-dimensional space relations in essentially two-dimensional visual data, rather than taking full advantage of the stereo-optical capabilities of our eyes.)


An image in a mirror has the same double representation--it is two-dimensionally in a two-dimensional mirror-surface; the physical objects seen are not three-dimensionally present in the real world, but they can take all the normal three-dimensional space relations among themselves. Our vision of them can even be stereoscopic.

Now consider a girl in a window. First of all, let us assume she is not in the window in a literal three-dimensional sense. But from some perspective her visual image is two-dimensionally interior to the visual boundary provided by the opaque edges of the window. The two-dimensional nature of perception is pervasive enough

that even though we know the girl is behind the window from the observer, it is perfectly acceptable to speak of her as in it.



Finally, consider a girl in a doorway. We may still mean the same kind of relation as in the immediately preceding window example--the visual image of the girl is two-dimensionally interior to the visual boundaries of the door, and the girl may actually be spatially behind the doorway from the observer. Note that in this case, as in the window case, the support for the physical girl is not the bottom of the window or the door and there is no inference that the girl and the window- or door-frame are touching. On the other hand, a literal, three-dimensional representation is also possible, in which case,



support and contiguity are implied. (People do curl themselves up to sit in windows, too, so a three-dimensional meaning is reasonable there, as well, with attendant support and contiguity.) So we have come to rest back with prototypical in.

Geography and contact locatives

Besides fundamentally two-dimensional surfaces, there are other semantic objects of in relations which have a two-dimensional flavor. Obvious candidates are geographical entities; in the United States, in Massachusetts, in the Midwest, in the Canadian Rockies, in southern France, in Chicago, in the town of Westhampton--all suppose some sort of two-dimensional border and a restriction of location (on a somewhat more global scale than any of our previous examples) to the INTERIOR2 of the "surface" thus delineated. Questions of support and contiguity between subject and object of the in relation are generally left unanswered.

Likewise, the "contact locatives" mentioned in the introduction to all the prepositions appear to address themselves purely to locative information, logically enough since the things involved rarely require support. In the example, "Mark slapped Jim in the face", it is the point of contact implied by slap that is in the face, not the physical objects Mark or Jim, or even the event of slapping. It is interesting that the semantic object of a contact locative in is normally treated two-dimensionally, though, as for example Jim's face in the example above. Other examples include "James dented the car in the fender", "Daddy tickled Jenni in the ribs", "the bully hit him in the back". Elsewhere in this paper it is argued that precisely this tendency of contact locatives to signify

two-dimensional restriction of location and little else accounts for the ease with which in and on can be interchanged in many such uses.

Some last thoughts

Of course, what has preceded hardly pretends to cover all the spatial uses of the preposition in; however, with slight variations which might correspond to world knowledge about this particular semantic subject or that semantic object, it is encouraging to see how many of the common uses of in are encompassed. Before moving on, there are some classes that merit attention, however brief.

Consider events in general. Events themselves require no support (though their participants may) and accept locations optionally. Whether an event as subject of an in relation is located two-dimensionally or three-dimensionally may be dependent on the object of the relation, the participant(s) in the event, or the nature of the event itself. For example, flight ordinarily signals three-dimensions. "The lion walked about in the cage" restricts the path of the walking to the bottom surface of the cage, but "the insect walked about in the cage" has no such restriction--the event of walking implies a surface, but insects aren't fussy about the orientation of their support, so any surface of the cage will do. "My Daddy works in the World Trade Building," thanks to the nature of tall, many-storied buildings, has no simple space

restriction relative to the object of the in relation. Incidentally, this last example is no more nor less complicated than the direct locative statement "My Daddy is in the World Trade Building."

Some events specify relevant portions of objects with respect to an in relation. After a summer shower, the answer to "Where's Jenni?" is apt to be "Out in a puddle somewhere." No one expects Jenni to be totally immersed in the three-dimensional volume of a puddle; likely this is a two-dimensional in--she is located within the boundaries of the horizontal plane of a puddle. If we say now that she is standing in a puddle, a relevant portion of her--the feet and perhaps part of the legs--are imagined to be in the INTERIOR³ of the puddle (this, based on world knowledge of the usual depths of puddles--if she were standing in a swimming pool, we would imagine quite a bit more of her to be interior to the volume of the pool*). If she is sitting in the puddle, still other parts of her are interior to the volume of the puddle.

* This is essentially what Waltz pointed out [Waltz, in press] in his example of a dog biting a man--world knowledge affects one's estimate of the location of the resulting injury based on the size of the dog, size of the man, and position (sitting, standing, bending) of the man.

Cracks and other discontinuities are interesting subjects and objects of in relations. They are visually perceptible objects (VPOBs) and may be two- or three-dimensional. The simplest representation of a crack in the wall is of a VPOB located in the INTERIOR² of the wall treated as a surface. (We sometimes speak of "the crack on the wall.") Consider what happens when a roach disappears into the crack on the wall, though. The wall may still be treated as a surface, if we wish, but the crack may definitely have taken on a new dimension. "Mommy, there's a bug in that crack" may mean the bug is visually perceptible, waving its antennae at the speaker, and two-dimensionally interior à la the girl-in-the-window example, or it may mean the bug is three-dimensionally interior to a volume delineated by the crack as container, presumably in contact with and supported by one of the surfaces of the crack, and incidentally not necessarily visible to the speaker at all.

Holes, rips, tears, and so on share these characteristics. They may be two-dimensional VPOBs in two-dimensional bounded regions; they may be three-dimensional VPOBs in essentially two-dimensional objects (sort of the inverse of a three-dimensional tree in an essentially two-dimensional field) or they can be three-dimensional and embedded in three-dimensional objects.

We have a lot of world knowledge about cracks and holes and such. One interesting item of information about them is that they often serve as pathways--not only for animated objects, like bugs, but in some circumstances for perception (we often look through holes) or for fluids (a crack in a cereal bowl, or a hole in a bucket).

Which leads us to consider fluids, since they tend to be "in" a lot of things, mostly containers. Liquids need support, of course, but they differ from regular physical objects in containers in that they do not take support from other objects in the same container--in other words there is no chain-of-support interpretation waiting in the wings for liquids as there is for physical objects in general. Since liquids have no characteristic boundary other than the top horizontal plane, regular CONTIG does not apply, either--the relation is much more the kind of contiguity implied by embedding, with the exception of that horizontal plane, again. This is as much true for objects in the fluids as for fluids in containers--e.g., "noodles and carrots in the soup", "a fly in the soup", "soup in the bowl", "the soup's in me!"

Finally, there are some troublesome containers--envelopes and sacks, to name two. Envelopes are essentially two-dimensional but they contain things in a three-dimensional sense. They don't have bottoms except in a canonical sense, their contents are typically light enough

that friction overrides gravity. and on a macroscopic level their contents seem to be embedded as much as contiguous--in short, they're a mess to model. The main problem with sacks is that they have very definite boundaries but very context-dependent shapes. (Consider modelling what happens when you partially fill a plastic bag with water from a tap (supported only at the top of the bag) and then set the bag on a table--either with or without a closure at the top of the bag.) Some bags are so flexible that they provide support if and only if nothing else does....but then we only promised to try to capture most of the spatial uses of in. Probably we can claim to have done that--in these latter cases the problem lies with representing the objects and not so much understanding the nature of in with respect to the objects.

ON

I remember being intrigued, as a high-school junior in chemistry class, with the peculiar glob of white stuff on the ceiling, looking like a miniature upside-down mountain. Granted, such a situation is hardly a prototypical case of the on relation, yet there is nothing strained in the use of the preposition on in the description above. Moreover, young children have no apparent difficulty understanding sentences involving lights on the ceiling, flies on the ceiling, paint and shadows on walls, and designs on plates. So, while we shall define a prototypical on, we also wish on to be sufficiently well understood to handle the description of the interesting stuff on the chemistry room ceiling as well as the following phrases, among others:

the nose on his face

the light on the ceiling

the circle of light on the ceiling

a boat on the Delaware River

a city on the Delaware River

on the road to Morocco (on my way to Mandelay)

the poster on the wall

a shadow on the wall

the cheerful daisy design on her plate

a slap on the back

Murder on the Orient Express (on a bus, on a train, on a plane, on a boat but not on a car or a taxi)

As with in, there is a prototypical definition of on--one which springs to the minds of most people when they are asked to define the word. As with in, the prototype relation is between two physical objects: for convenience, in the phrase "the doll on the table" the doll will be called the subject of the relationship and the table will be called the object of the relationship. Although the object is three-dimensional, typically its main feature of interest is a horizontal free-surface, with which the subject of the relation is in contact on the free side, and by which the subject is supported. So then, in the prototypical case, there are two physical objects, the subject and object of the relation; the subject and object are related by CONTIG, and moreover the location of the contiguity is restricted relative to the object of the relation; although the object is a physical object and hence is three-dimensional and has weight and requires support and all the other cognitive baggage that physical objects carry with them, for the purposes of the prototypical on, the object could be approximated as a horizontal free-surface--on between two physical objects ordinarily designates a relevant plane, the horizontal.

Finally, in the prototypical case, the subject of the on relation is higher (farther in the positive vertical direction) than the object, in the sense that the center of visual mass of the subject has a larger z-coordinate than that associated with the horizontal surface of the object.

In the doll on the table example, doll and table are physical objects, the horizontal restriction singles out the tabletop, as opposed to the table legs, the CONTIG and SUPPORT relations between the doll and the table are noted, the doll is located relative to the table, which as object of the preposition becomes the local origin (given the relative scales of dolls and tables, the entire table obviously does not serve as a point origin; some point of the table is chosen as the local origin) and the z-coordinate of the center of visual mass of the doll is noted as being an appropriate increment above the z-coordinate associated with the top of the table.

Let us now consider an on relation between physical objects in which some of the prototypical elements are relaxed or missing. The most basic "relaxation" of all--defining horizontal in terms other than the plane exactly perpendicular to the direction defined by gravity--could hardly be considered a relaxation at all, since real-world visual systems which take into account angles provided by a necessarily inexact determination of the true horizon (particularly indoors), the less-than-perfectly-straight lines which serve as "straight lines" in the real world, and so forth, have a built-in tendency to "make allowances". If a horizontal surface were expected, a surface that was almost horizontal or almost planar would do. It seems plausible that humans do a lot of the same sort of approximating.

There are on relations between physical objects where it is unlikely that the supporting object is accepted as approximately planar: e.g. a seal balancing on a ball, a boy sitting on a pony; there are a number of things which can be inferred from such situations (such as conformability of the subject of the on relation to the shape of the object of the relation) but it is not altogether clear that such inferences properly belong to the on relationship rather than a fairly sophisticated knowledge of the world. Positing as much of the remainder of the prototype relation as possible does seem to capture much of the intention of the descriptions--a support relation, CONTIG, the usual relationship on the vertical scale, and in cases where part of the supporting object has a characteristic horizontal outline associated with a free-surface (as the backs of animals) giving first preference to that portion of the object in determining relative location. In other words, all of the prototype definition of on which can be applied, is.

On can be used in relationships between three-dimensional physical objects which are not CONTIG, under a restricted set of conditions. In the doll on the table example, for instance, even if the table is draped with a tablecloth, it is highly unlikely that the doll's location would be specified as "on the tablecloth"--when the subject of the relation is three-dimensional, such a description seems to be reserved for use as a contrastive:

if the tablecloth were rolled back for a game it would be natural to speak of the objects on the tablecloth as opposed to those on the table. Similarly, a doll which happens to be on a placemat would normally be referred to as "on the table" unless the question of whether it were specifically on a placemat or not were important. A book on a blotter on a desk may still be said to be on a desk, as may a book on a sheaf of papers or even on a pile of clutter on the desk.

On the other hand, keys on a brick on a concrete block would not ordinarily be described as being "on a block". It would appear that there are two factors (actually one major factor and a contributing factor) which determine whether the middle object is transparent to an on relation between the other two: if the depth of the intervening physical object is small with respect to the dimensions of the other two, so that the intermediate object may be treated as essentially a two-dimensional surface, then it may be treated as transparent to the on relationship. A pot on a rigid ceramic tile on a counter may still be said to be on the counter. It helps if the intervening object is non-rigid and not capable of providing support, like a tablecloth. (It is just possible that non-supportiveness alone suffices to make the object transparent to the on relation, but it is difficult to find sufficiently deep objects that don't provide some support. The pile of clutter on the desk is a possible candidate, but I suspect there is a better explanation for that particular case.)

At this point let me emphasize that I am talking about the prototypical three-dimensional on only, between physical objects. There are other uses of on which have not yet been mentioned. A number of readers may wonder why I make an issue of objects being transparent to prototypical on, since on has been known to be transitive or at least has been treated as such in AI literature for years. It is my contention that on is not always transitive. Obviously a man on a building on the ground is not on the ground. A toy on a chair on the floor is not on the floor. A toy on a rug on the floor, however, may be said to be on the floor. Now, I do not wish to take examples based on a floor or the ground as strong support for the point I wish to make, since there may be other factors at work in the non-transitivity of on when floors and the ground are involved; namely, they are the default supports in indoor and outdoor scenarios, respectively, and to mention the floor or ground as the object of an on relation ordinarily explicitly denies any other support. Nevertheless, it is at least interesting that rugs and other physical objects which conform to the caveats above may be transparent to an on relation with respect to the floor, where one might have expected transparency to be much more difficult to achieve.

Since examples with floors and the ground may be complicated by principles unrelated to the nature of the prototypical three-dimensional on only, it is fortunate that there are other persuasive examples. A child (3.11) to whom

a fan on a box on a table was pointed out (literally--verbalization was minimized: "Do you see that fan?"...pointed finger), when asked if the fan were on the table immediately responded "no". Granted, an adult might have hesitated or even have responded "yes". It is my contention that the basis of an adult's granting a sense of on which is applicable to the fan and table is a non-prototypical sense of on which is learned and practised somewhat later and which has more to do with search procedures as a separate case. (See discussion below.) Clearly, the child's unhesitating "no" indicates that there is a sense in which the fan is not on the table, and I claim that the on that fails to apply is the prototypical three-dimensional sense of on.

As it happens, it is fairly easy to create physical demonstrations of three or more physical objects having non-negligible sizes in all three dimensions in which adults show uneasiness in granting transitivity to the on relation. For some reason (I will propose a possible explanation shortly) verbal descriptions of the same demonstration are lacking in power. A bowl on a loudspeaker on a shelf in my apartment is a case in point. Confronted with the actual objects, adults admit that there is a very low probability that the bowl would be described as on the shelf, even though the search-procedure interpretation of on is applicable in this case.

What I am suggesting, then, is that the prototypical on for three-dimensional objects is not transitive. Just as the fan on the box is judged by the child not to be on the table, so there is a sense in which an object on a tablecloth on a table may be said not to be on the table. You may argue that anyone who insists the object is not on the table is being excessively legalistic, but if they persist in the end you're going to lose the argument--the object is not literally on the table.

While the prototypical three-dimensional on is not transitive, there are certain objects that are readily transparent to it, as mentioned above. The class of objects which are readily transparent to this meaning of on are precisely those which are of marginal three-dimensionality, so their tendency to be invisible to the relation seems fairly reasonable: after all, if a circle is stained on the surface of the table, without question a doll on the circle is on the table in the three-dimensional prototypical sense--the circle is not a physical object and, being of two dimensions, does not affect this meaning of on at all. The closer a physical object comes to being a depth-less surface, the greater its tendency to drop out of the relation; hence the loss of the argument mentioned above is sometimes grudgingly allowed--ordinarily the object is felt to be on the table (not the tablecloth) in the fundamental sense.

Let us take another example: suppose we stack the first ten volumes of the Encyclopedia Britannica in order on a desk, with volume one at the bottom and volume ten at the top. Now, there is a non-controversial sense in which volume ten is on volume nine. And there is a sense in which most adults, at least, would say that volume ten is on the desk. But there is no natural sense in which volume ten can be said to be on volume three. If there is a prototypical three-dimensional on relation between volumes ten and nine, and there is some kind of on relation between the physical objects volume ten and the desk, and our fundamental on is not transitive (it doesn't apply to the intervening volume three) then what is the nature of the on relationship between volume ten and the desk? Consider the answer to the following question: "Richie-e-e! Where's volume ten of the encyclopedia?" The relevant answer is "On my desk." If Richie answers "On volume nine" he is attempting either to irritate the questioner or to be humorous.

Contrast this situation with the situation where questioner and answerer have full knowledge of the desk and the stack of books. Ask "What is volume ten on?" and the simplistic answer is "on volume nine" (once someone has answered this way it is difficult to bias an answer toward the other possible response "on the desk"--the prompt "Yes, but what else is it on?" can get either a blank or a mischievous look and the answer "it's not on anything else"). My point is that the response "on volume nine" to

the latter question is an honest response--it does not violate any of Grice's rules of conversational implicature. In the former case it violates a rule.

To my knowledge, every instance of an on relationship between physical objects that has been presented to me as evidence for a transitive on has been put to me in a context where the question "where is the object (subject of the on relationship)?" is relevant. I call this the search procedure context. It is quite prevalent; grownups do a lot of it. The outstanding feature of the context is the intent to provide location, often on a relatively global scale (Richie's use of the phrase "my desk" implies the building and room as well as the relevant piece of furniture--localizing that far seems to suffice) and almost always using the map analogy. A set of keys on a brick on a desk are literally on the brick, but for someone looking for the keys they are "on the desk in the livingroom". "On the brick on the desk in the livingroom" is also an acceptable description as search procedure, of course, since once the landmark "desk" has been specified, the further landmark "brick" may be helpful. "On the brick in the livingroom" is natural only under specialized circumstances.* Natural

* Context can naturalize almost any sentence--witness the famous "I think with a fork" example. The incongruity disappears when the preceding sentence is assumed to be something like "How are you supposed to eat spaghetti?"

landmarks in an indoor map scenario include rooms, named spaces (the dining ell, the entrance way), closets, pieces of furniture, shelves (horizontal surfaces which tend to collect items, mantels and wide window sills included) and finally, in more detailed scenarios, one-of-a-kind items or uniquely-specified objects.

The naturalness of using a shelf as a landmark in a search procedure (localization procedure, if you prefer--I do not intend that a search must literally be taking place) in part accounts for the lack of power of the bowl on the speaker on the shelf example previously. In a written sentence, the landmark character of shelf is the overriding feature; in the actual situation that presents itself, the speaker is a reasonable candidate for a piece of furniture itself--a plausible landmark--and the prototypical three-dimensional on relationship begins to assert its applicability. "On the speaker on the shelf" is a reasonable description, as well as "on the speaker"; "on the shelf", in the presence of the evidence, is generally agreed to be an unlikely voluntary description.

As has undoubtedly been noticed, both types of on provide locational information--that the search-procedure on stresses location in a fairly formalized sense is not meant to imply that prototypical three-dimensional on is deficient in locative information. In a great many cases it is not clear whether the on in a description is meant to be

prototypical or search-procedural. When possible it is taken as both. If further description indicates that the subject and object of the on relation are not CONTIG and that the intervening object cannot be generalized as a depth-less surface then the stronger definition is relaxed and the search-procedural one retained.

Departures from the prototype

The prototypical on is defined to single out horizontal free-surfaces whose free sides are in the vertically upward direction--in a use of on with neutral subject and object. "the whatchamacallit is on the thingamabob." such a horizontal surface is a preferred characteristic of the object of the relation. It is, however, readily altered by known characteristics of the subject or object of the relation, particularly the object.

Suppose the object is a ceiling, for instance. Ceilings in most instances are treated as nothing but free-surfaces; their three-dimensional nature is usually ignored and they become essentially two-dimensional free-surfaces with a marked horizontal plane and marked direction in the vertically downward direction. On accepts this redefinition of salient space relation--a bug or light on the ceiling is vertically downward from the ceiling. (By contrast, above does not accept such redefinition, and over is forced into a specialized "covering" sort of meaning, as in "there must have been fifty mosquitoes bumblng and

crawling over the ceiling"). The support relation still exists, as does CONTIG, and of course the standard locative applies, with the object of the relation supplying the local origin and boundaries, and normally with the simplification that the third dimension of the object of the on does not enter into the locative description.

Walls are treated in an analogous manner, except that the free-surface is vertical and the free-direction is the horizontal perpendicular to the wall. For objects on indoor walls, the salient space relation is accordingly revised. A painting carries with it a preference for an orientation as a vertical plane, so a painting on a wall description has both painting and wall vertical planes, CONTIG (hence with orientation for maximal shared surfaces). a support relation, and the center of visual mass of the painting at a (small) horizontal displacement from the plane of the wall (assuming a painting as a physical object and not a directly applied design). A description of a shelf on a wall works equally well: the wall is vertical, the shelf is characteristically horizontal, they are CONTIG (hence the orientation of the shelf in the horizontal plane is such as to maximize contact with the wall--along the length of the shelf, in other words). with the wall supporting the shelf, and the center of the shelf displaced horizontally from the plane of the wall.

In outdoor scenarios, the alternatives are somewhat more complicated since outdoor walls may be referred to in their three-dimensional character--the subject of the on relation may have characteristics which determine whether the on is taken in its prototypical sense or with a marked vertical plane. Posters and signs usually mark the vertical plane, as do windows ("the top window on the wall of that building" is not on top of the wall). Insects do not require horizontal support, so the vertical nature of the wall predominates. Most normal physical objects do require a horizontal support, so a bottle on a wall predisposes a prototypical interpretation.

Occasionally, an essentially one-dimensional horizontal line is used as the object of an on relation. Birds on wires and circus performers on tightropes come close to being prototypical cases of on, but most others could be described as cases involving the concept of hanging; subjects of the relations often are non-rigid--laundry on a clothesline, for example, or a shirt on a coat hangar, or curtains on curtain rods--but need not be; virtually anything can be hung on a hook--clothes, hats, cups, pots, keys, tools, coiled ropes and cables and hoses, large containers with handles (gasoline cans, for instance), even ships' cargoes suitably tied. For this latter use of on, the visual (and physical) center of mass of the subject of the relation is below the object of the relation, and only a small part (typically) of the subject is contiguous with and

vertically upward from the object providing support. This "part" is not always a true part of the subject; it can be part of it by convention, like key-rings and curtain hooks, where people typically don't take the trouble to distinguish between the keys, say, and the key-ring, or it may be a (relatively) small fastener itself, like a loop of twine.

On without support

Up until now the variations of on have relaxed restrictions on the horizontal, on CONTIG, and on the planar nature of the object of the relation, but they have all had the support relation in common. There are, however, several uses of the preposition on that do not entail support.

The most obvious examples are those in which the subject of the relation is visually perceptible (VPOB) but not a physical object, e.g., shadows, patterns, a circle of light. The need for support is a characteristic of physical objects, so it is not so much that the support relation is missing as that it is not applicable. The emphases on horizontal planes and vertical directions are gone also. What is left is an extremely strong CONTIG (coincidence of coordinates) and of course the locative information bounding the location of the subject relative to the object.*

* Thin layers of liquids (blood on an arm, oil on an engine block, ketchup on a shirt) share some characteristics with this category--conformity with whatever surface they are in contact with, and effectively no depth--and some characteristics with the set of physical objects which includes insects--needing support but not being fussy about the orientation of the surface that provides it.

In some cases the subject of the on relation is an event--a slap on the back or a walk on the seashore, for example. As has been pointed out in previous sections, some types of events can take more than one locative; some events, for instance, can take normal locatives and what I have called "contact locatives". Nor is it always clear whether the subject of a relation is meant to be an object or an event (as in "Liz dented the car on the showroom floor"). For our present purposes we shall assume that the subject is unquestionably an event and that the type of locative is not in question.

Support is not required, as evidenced by the slap on the back example, although it can be inferred for whatever object is taking the walk. Contiguity is not a strong feature of the relation, either, as indicated by the sentence "Mike lost his wallet on the bridge." Participants in an event may inherit a CONTIG or support relation, but the event itself rarely does. What the event does take is the locative-bounding information (recall that any event can take a locative). The point of contact of the slap is within the boundaries of the back; the path of the (durational) walk is within the boundaries of the seashore; the location at which the wallet was lost is somewhere within the boundaries of the bridge. The free-surface interpretation is invoked, which in turn means the object of the on relation effectively becomes a bounded surface, and the location of the event is restricted to this bounded

surface.

Finally, there is another use of on which is entirely locative in nature--no support or CONTIG are implied, nor is the free-surface or upward-directional implication in force. Up until now the lack of a support relation was signalled by the characteristics of the subject of the on relation. For what follows it is primarily the object of the on which is the key to the meaning, along with general context.

In English terminology a point is usually in a region; it may be in a plane or on a plane (a point is on a map) and it is virtually always on a line. Paths are often abstracted to lines*, as are boundaries, both real and abstract. Juarez is on the Mexican border. It is also on the Rio Grande (although obviously not in the sense that a rowboat is on the same river). For anybody contemplating an east-west trip, Urbana is on I-74. A box of merchandise can be "on its way to Chicago," even if it is being flown by plane--no free-surface, no VPOB for a CONTIG relation, and no support are supplied by the on...just the metaphor of a

* Hence the otherwise strange "On our trip to Texas we saw the snow geese migrating." On cannot be glossed as during or we would have "On the lecture half the audience fell asleep." (By contrast at often can be glossed as during and works well in this sentence.) A trip is a path; paths are readily abstracted to lines, and it is difficult in English to speak of points or places on lines in any other terms than on the line. The gloss, then, is "[at a point or points (a place or places)] on our trip to Texas..."

linear path and restriction of location to (a point on) the path.

While the uses of on characterized thus far include what are probably the most common spatial uses of the preposition, this catalog is by no means exhaustive. Consider, for instance "on the south side of Chicago" (or the far side of that tree or the windward side of the factory--they all take the word on in conjunction with "side"); this particular use is directional as much as anything. On can have a proper-part-of meaning, too: "the nose on his face", "five toes on each foot", "one knee on each leg". [My first inclination was to treat these as more normal locative uses of on--after all, faces are bounded surfaces, noses and smiles are quite naturally on them, and while the knee-leg, toes-foot examples seemed to be stretching things a bit, I wasn't willing to make a separate case for them. But then, there was my daughter, lustily singing "The wheels on the bus go 'round and 'round...."] If one's neighbor cheerfully volunteers the information that the whatever-it-is on his bike is broken, one equally cheerfully assumes that the whatever-it-is is a proper part of the bike.

In/on: interchangeability

Having mentioned the wheels-on-the-bus song, we might as well consider one of the other verses of the same song: "the people on the bus go up and down..." Now it is true that buses have a horizontal free-surface on top, but in the United States, at least, anybody who knows about buses knows that's not where the people are. Likewise, people on a plane are not on the top surface of the plane, and people taking a trip on a boat are not necessarily on the top surface of the boat. Ordinarily the people are in the boat, the plane, and the bus. Yet English permits the use of the word on. On the other hand, "people on a car" has the prototypical meaning. "On a taxi" has the same flavor as "on a car," so the secret is not that the object is public transportation. People can stand up inside buses, planes, large boats and subways, but not in cars, taxis, and canoes, which just happens to parallel nicely the naturalness of using on with these objects, but if you ask for a judgment in using in or on with a ricksha or a howdah as object, informants have no clear intuition for whether on is appropriate or not and generally do not object to it, even though standing in a howdah or ricksha is not more likely than in a car. Probably this is because most English-speaking people in the United States do not have a convention in speaking about rickshas and howdahs. Given that in and on are conventionally interchangeable with some vehicles, in the case of a vehicle where the convention is

unknown but potentially applicable the interchange is allowed.

What is meant by the on in such cases? The answer seems to be pretty much a straight restriction of location. The free-surface interpretation obviously does not apply; there is no more support relation than in the interchangeable in relation; and normal CONTIG between the bus itself and its occupants is not necessary (a baby in the arms of its mother may still be on the bus).

We have already seen nearly purely locative uses of in and on (recall the search-procedural on). Because both prepositions can be used to denote locative information only, in many cases, in the absence of strong conventions, they can be used interchangeably. Consider the following contact-locative positions: Marty was hit ___ the face, ___ the nose, ___ the mouth, ___ the chin, ___ the stomach, ___ the shoulder, ___ the back, ___ the arm, ___ the leg. About the only objects that seem strange with on are face, mouth, and stomach. In could be used with all of them. But if the verb changes to slapped, the preference switches to on, with the exception of "in the face," and with either choice sounding strange with stomach (except when the individuals involved are lying down, it's difficult to imagine anyone's getting slapped on the stomach). In is still acceptable in all the cases though. Both in and on are performing the same job--locating the point of contact with reference to

named regions.

The list stretches endlessly--a scratch on my leg, a scratch in my glasses; the window in the wall, the fifth window on the wall; a boy lost on the mountain, a boy lost in the mountains. After making judgments for a while informants generally object that they feel they are being somewhat arbitrary in choosing one preposition over the other.

In sum, given that in some circumstances the prepositions in and on are indistinguishable by usage, it is my contention that their primary purpose is the same in those circumstances--namely a relatively pure locative.

TO AND FROM

While these two prepositions can and often do occur separately, there seems to be a natural tendency for them to be used in pairs, as in "How the paper got from my desk to the kitchen remains a mystery." Prototypically the pairwise use of from and to indicate source and destination (consider "the thingamajig went from x to y" or "from the x to the y"), with a probable path restriction to an approximately straight line between source and destination whenever the path is not specified.

Although typically a verb of motion is used in conjunction with to and from, they can be used to indicate a path (hence motion) without the help of other parts of the sentence: "the kids argued with each other from Detroit to Chicago" describes a path and hence implies motion, even though no other element of the sentence requires any such interpretation. Notice that when this occurs, a relatively strong sense of duration accompanies the prepositional phrases. Prepositions of time usually refer to specific events; in this case the reference event appears to be the implied traversal of the path, rather than the specifically mentioned arguing. (How much later the arrival in Chicago was after the departure from Detroit depends of course on factors not mentioned in the sentence--it could have been as little as a couple of hours if they were flying, or perhaps days under other circumstances.) The sentence is relatively

precise about locations; time considerations in this case are indirect.

So strongly connected are to and from that mention of to between two place designations ("We took the Rome to Paris train", "the wind was blowing east to west") has the same effect as a from-to pair. And a sequence of to's, as in "from the mousehole to the shoe, to my chair, to the t.v., to the windowsill," is effectively a short form for "from the mousehole to the shoe, from the shoe to my chair, from my chair to the t.v., from the t.v. to the windowsill" (provided that one object is doing the moving--if, for instance, we are talking about four or more mice, the original description may be of four distinct paths with a common source but different destinations).

Let us look for a moment at the locations thus specified. If a person goes to a chair, she may be standing in close proximity to it; if she goes to someone's office, she probably is interior to the office; if she goes to the office door, then again a closeness is implied, but not necessarily an interior relation (although if the door is open, there is no reason to suppose that she is not in the doorway in the usual sense of that phrase). If a mouse, scurrying around investigating the floor, is described as going to a chair as in the passage above, the relation between the mouse and the chair might well be "under", or within a short distance of being directly under the chair.

or for that matter on the chair.

What then do we know about the destinations of the subjects of all these to relations?

First of all, let us assume that to has nothing to tell us about relative placement in the vertical direction--if we need that sort of information and the objects in canonical position or default relation don't give it to us, something besides or in addition to to is going to be required.* In essence, to is an appeal to the map metaphor--the placement of objects relative to the basic horizontal plane.


Most of the semantic objects in the examples above were three-dimensional, and what we have left in the absence of the vertical dimension is what I call the horizontal cross-section of the object. It could be likened to the vertical projection of the object--the shadow, if you will, cast by the object when lighted from directly above. When to is used on a scale at which the objects have dimensional character (as opposed to being points), in most cases it is restricting the location of the semantic subject in the horizontal plane to within a "small" distance of wherever the horizontal cross-section of the object falls in the horizontal plane. Lest there be a misunderstanding, the

* This is a temporary supposition only--we will revise it shortly.

horizontal cross-section of an object is not a location: it is a perceptual attribute. Recall that the semantic object of a preposition serves as a local origin, so that its location can be presumed "given"; as a result the horizontal cross-section, in conjunction with the "given" location of the object and enlarged by the "small" increment mentioned, yields the actual restriction of location of the semantic subject.

The nature of the "small" distance mentioned above is not simple. It is affected by the dimensions of the semantic object, the semantic subject, and, in cases where the path is explicit, possibly by the length of the path (especially paths that are extremely long or short relative to the dimensions of the objects).

Such an interpretation effectively allows us to skirt the question of whether the mouse has gone to a position under, near, or on the chair; the position relative to the horizontal plane is restricted, but we are relatively free to envisage the mouse wherever we wish, subject to those restrictions.



In the description "the children ran to the field" the same sort of restriction holds true--the field is essentially its own cross-section, and there is at least a

potential incremental area around it that serves as the boundary for the locations of the children. For myself, the default is to place the children in the field when I visualize the field. But consider that, in my own experience, a bull pen (the kind you find on a farm) is a reasonably large field with a stout board fence around it, and the sentence "the children ran to the bull pen" leaves in my mind's eye an image of the children standing outside or perhaps on the fence, looking in, the interior of the field being excluded from the default location in this instance.

Similarly, "the children raced to the park" defaults to the inside of the park for me, in the absence of further information, but the actual restrictions imposed must be somewhat broader, since in the sentence "the children raced to the park but it was locked" the children never make it inside.

So far all the semantic objects have been open enough to allow physical objects to be interior to their horizontal cross-sections. Obviously, such is not the case for a solid semantic object--e.g., "the children raced to the telephone pole." Children and most other objects are not usually embedded in poles, so the interior of the cross-section is removed from the realm of consideration in the placing of the semantic subject in the mental model. One is tempted to say that when the interior is not possible, the boundary of

the object takes on special status, as seems to be the case with this example (the default for the children is racing to touch the pole) but consider the instruction "Go to the telephone pole and wait for me," where it is likely the speaker and hearer do not assume the pole must be touched.

For the examples discussed so far, the vertical position of the subject has been given by the characteristics of the subject and its presumed path--children and mice typically stick to the default horizontal plane--the ground or floor--though they have been known to do a little climbing around now and then. Clearly, though, our initial assumption that to has nothing to tell us about vertical placement is an oversimplification. If the semantic object of to has a strongly characteristic height above the horizontal plane (as wall clocks, posters, and ceilings do, among others) then to used in connection with these objects certainly does imply some information in the vertical direction, even to the extent of starting a search for plausible path and/or means of travel. If "Jim got to the ceiling of the gym," he didn't get there by a default path (maybe by a ladder or by climbing equipment secured to the walls). Clearly, for such objects it does not do simply to take the horizontal cross-section of the

object and the default vertical placement of the subject.* Ordinarily, of course, the characteristics of the subject (insects, birds, things that relax gravitational constraints) and verb of motion (flew, for example), or verb of projection (such as look or point) provide the strongest clues to vertical restrictions; the point of the ceiling example was that in the absence of other clues, to plus special objects can yield vertical information.

The "legal" to

Having roughly described the general flavor of to, we should at least take note that there is a more exacting version of the preposition. Suppose a family is driving westward in Louisiana, and a few feet short of the Texas border they turn around and drive east again. The answer to "did they get to Texas?" is unquestionably "no". They had to reach the border itself to have gotten "to Texas". Children, when told to go to some landmark or other seem to have a tendency to touch the landmark. And clearly in their

* It is interesting that the objects for which this seems to be true are precisely those objects whose characterizations in a computer implementation to accompany this thesis required specification of a characteristic height above the horizontal plane. A clock on a wall, for example, generally cannot be just anywhere on the wall. And the ceiling in Jonathan's room is presumed to be at the top of Jonathan's room, which is an implicit "characteristic height".

paces ("[I can] Beat you to the sand pile!") the actual boundary of the object and not some small distance from it is what counts. There is, moreover, some sense in which a mouse which comes very close to a piece of cheese but does not touch it can be argued not quite to have gone to it, in much the same sense that one could argue that a glass on a placemat on a table is not on the table.

This more legalistic sense of "to" seems to have greatest applicability in the realm of geographical descriptions and in certain well-rehearsed situations, as in races. But that it is potentially applicable to more general cases must be recognized.

Since to and from are the first motion prepositions to be discussed, it is perhaps worth repeating what was noted in the section on motion prepositions--if the implicit or explicit path is large with respect to the dimensions of the source or destination, these may be treated as point locations. In fact, if the literature is any indication, there seems to be an overwhelming temptation to treat all destinations as point locations. By this time it should be obvious that not all of them are, but clearly people find it natural to treat sources and destinations as point locations whenever possible.

AD-A085 750

ILLINOIS UNIV AT URBANA-CHAMPAIGN COORDINATED SCIENCE LAB F/0 3/7
COMPUTATIONAL INTERPRETATION: ENGLISH SPATIAL PREPOSITIONS. (U)
FEB 79 L C BOGGESS


N00014-67-A-0305-0026

ML


UNCLASSIFIED T-75

2 of 2

ALL INFORMATION CONTAINED
HEREIN IS UNCLASSIFIED



END
DATE
FILMED
8-80
DTIC



To invokes scenarios

There is at least one other major use of to without which a discussion of the preposition is clearly incomplete. Suppose a child, when asked "Where's your mother?", responds "She's gone to the store." Usually the questioner has no idea which store "the store" is *, and the child may not either. In fact, the mother may have gone to several stores. In effect, the child has said that her mother has "gone shopping". The child has invoked a scenario.

We use to plus a number of phrases to evoke specialized scenarios--it seems to be an important linguistic/cognitive function of the preposition (at serves the same function): to go "to the doctor" is one such scenario, "to the hospital" is a slightly different one. "The children have gone to school" (perhaps several different schools); a spouse is late "to work", Mommy goes "to the office", or maybe "to the laundromat". I have a helpful neighbor who often tells me she is going "to the library" and offers to take books back for me; now in many towns, going "to the library" gives locative information, as well as invoking a scenario--it happens that in this case we both know that we

* My own default is the nearest grocery store, but I wouldn't bet money on it.

both frequent the two public libraries available (I assume I can probably ignore the two dozen or so campus libraries available to us), but not usually both libraries on the same trip, so even though I know what my neighbor's activities are likely to be, I don't know what her location will be and before I can take advantage of her offer I have to ask. There is occasionally a public service announcement on television about all the things one can do "at the library" that is clearly invocation of a scenario, since no one location is implied--indeed, no one library is expected to fit their description of a library; in terms of this particular scenario, "going to the library" is an opportunity to "borrow classic films, attend a crafts class, see a puppet show, learn to belly dance..." and so forth, all--as they repeatedly chorus--"at the library". And the list of scenarios that may be invoked in a similar manner is seemingly endless (we go to the dentist, the beach, the movies, to church, to a funeral, a wedding...). If the default location for a scenario is well known, then invoking the scenario also provides locative information. If the same neighbor mentioned above talks about going to church I assume a particular church whose location I happen to know. Other neighbors, however, can talk about going to church and neither they nor I am bothered that I don't know which of many churches they intend, nor where the church is located. The location of the "destination" in these examples is incidental to the scenario-setting characteristic of the

sentences.

Some particular uses of from

Most of the foregoing discussion of to applies equally well to from. There are at least two uses of from which are not particularly characteristic of to (except when it is paired with from) and which are worth some discussion, at least: 1) the syntactic object of the preposition from can be another prepositional phrase, and 2) the restrictive information of from can be directional in nature, rather than place-locative.

Consider the following: "John came from behind the tree," "a voice came from above me," "the mechanic rolled out from underneath the car." "Underneath the car" is a place restriction, and we understand that the source of the mechanic's motion was the "place" indicated by the phrase "underneath the car." The occurrence of a prepositional phrase instead of a noun phrase scarcely affects interpretation.

The most likely interpretation of "behind the tree" is also place-locative, since usually an appropriate distance is implied with behind--that is, a gloss of the phrase might be "on the opposite side of the tree from the implied origin (point of view) and within a small distance of the tree." The "behind the tree" phrase, then, provides location restrictions, and the "from" indicates that the place so

restricted is the source of John's path.

The "voice above me" example yields to the same analysis, since above can have the same sort of appropriate distance assumption as behind. However, above not infrequently has a very relaxed appropriate distance (as does, say, beyond), so that the emphasis in interpretation shifts to restriction of direction rather than restriction of place-location. (Suppose, for example, the "voice from above me" phrase were in the context of the Old Testament; in my own understanding of such descriptions, distance drops out of the interpretation altogether.) The object of from can also be an explicit direction, as in "The storm is approaching from the south."

Before now, all the restrictions we have discussed have been place restrictions, even when they were restrictions on a path. We have delineated paths by assuming restrictions on a source, perhaps, and a destination, and perhaps a few intervening parts of the path, and assuming (in the absence of information to the contrary) that the path is essentially linear. A variation on this method is to bound a location in space and restrict the path to the location but not worry about what a description of the path would be--as in "the butterflies flitted above the flowers," where the location of the flowers bounds the path in the horizontal plane, and above is assumed to have the appropriate distance constraint recently mentioned. To repeat, all restrictions up until

now have been location restrictions; this is the first time we have had a direction restriction applied to a path. Essentially the origin is assumed known--sometimes it is supplied by to and a destination; in the storm example the origin is assumed to be whomever the storm is approaching, probably the hearers of the sentence--at any rate it is typically supplied by another element of the description. The path specified by the from phrase then is assumed to have the origin as destination and is bounded not implicitly by constraints on its parts but directly by constraints on its orientation with respect to the origin. In cylindrical or spherical coordinates this translates directly into restrictions on angles from the origin.

In sum, from with a directional argument is interpreted as describing a path that is radially toward a known origin, with restrictions on the orientation of the path relative to the origin.

Occasionally, from is used as a short form of away from. Away from actually takes a physical object or location as its argument, but its interpretation is analogous to the discussion above. In this case, the semantic object of away from is the origin; it is also the source (as opposed to destination) and again the path is radial with respect to the origin. In "the workmen scattered from the bomb site," the source is the bomb site (plus perhaps an appropriate distance), the destinations are

unspecified. the individual paths of the workmen are each assumed to be approximately radial with respect to the origin, and the verb "scattered" effectively informs that the paths are randomly spaced in the plane around the origin.

AT

At has many of the characteristics of to, although in its most characteristic uses it is used as a place locative and not in descriptions of motion. That consideration aside, however, "Mary stood at the door" has much the same locative flavor as "Mary went to the door"--the same area interior to the vertical projection of the horizontal cross-section of the doorway and within a small distance thereof is indicated. Again, judging from the literature, a lot of people treat place-locatives as if they are point locations, in many instances. I suspect this is done when the location being mentioned is by implication distant from the speaker or whatever location has been serving as focus of attention (e.g., "meanwhile, back at the ranch...")--it certainly does not apply for "the chair is at the table" or "Jenni is at the window."*

* Cooper [1968] glosses at as in or near, but it is entirely possible that part of the chair is under the table: allowing the location of the chair to be interior to the vertical projection of the horizontal cross-section of the table gets around this problem--the only other constraints are general real-world constraints that the back of the chair and the top of the table cannot occupy the same space. Notice also that "window" has a characteristic height so some part of Jenni is expected to be as high as the window, or nearly so--a baby who can't quite reach the window might still be "at the window."

What has been mentioned seems to work fairly well for deriving mental models from descriptions using at in its "literal" sense: certain location restrictions are implied, objects are placed subject to the restrictions, in canonical position by default. But the issue of deciding whether at is a suitable description of the spatial relation between objects is not so simple. Johnson-Laird has pointed out [to Waltz, personal communication] that at has the sense "in normal relation to"--so that an upside-down chair next to a table would not be described as at the table, and while we can say "Jenni is at the window" or "at the door", we don't say she is "at the drapes" or "at the wall". In these latter instances there is no norm to which we have recourse. This raises the serious possibility that, while there may be a simple interpretation available for understanding a description using at, the criteria for using at may be learned on a case by case basis. It is not enough simply for the locative restrictions to be fulfilled and the object to be in canonical position, for why, then, is it not natural to say "Jenni is at the wall"?

Moreover, at, like to, is frequently used to invoke scenarios; "at the doctor's", "at the beach", "at work", "at school". What is more, the same phrase could be used as scenario-invocation or for imparting locative information. If one harrassed spouse says to the other "Pammy left her shoes at the laundry, today!" it is entirely possible that which of several laundries is unimportant. On the other

hand, if the situation is that the parents are looking for Pammy's shoes and one suddenly blurts "she must have left them at the laundry!", clearly locative information is intended, and the particular laundry the speaker has in mind is quite important.

This scenario-setting use of at is so natural that we often don't realize we are doing it. Rumelhart and Levin [Norman and Rumelhart, 1975] use the sentence "John was located at UCSD from 1967 to 1969" as a characteristic example of the locative for the LNR system. But nobody hearing the example would expect that John's location was restricted to UCSD for two whole years. When the school is a college or university, the "at school" scenario is a very broad one indeed, allowing for daily excursions into a surrounding city or countryside, weekends even farther afield (if John's mother uttered the sentence to her neighbor, the neighbor might have volunteered something sympathetic about college boys bringing home sacks of laundry on weekends, without in any way negating the original statement), and even summers thousands of miles from campus.

In sum, the techniques of this paper have limited applicability to the place-locative uses of at, not being suited to the invocation of scenarios. Where locative information is intended, the theory cannot predict that at is inappropriate in all cases that a native speaker

would--the distinguishing features being not only perceptual properties and location restrictions but also whether the semantic subject might conventionally be found within the location restrictions of the semantic object.

There are two senses of at which are not strictly place-locative. If the semantic object of at is an event--"at the picnic", "at a lecture", "at their wedding" then the at may be invoking a scenario, or, if the location of the event is known, it may be providing locative information; not infrequently, though, if the subject is also an event, the at may also be glossed as during: "They ran out of champagne at the reception", "half the audience fell asleep at the lecture". This is not to say that the locative information is not also there; but events have strong identification with time and this is evidenced in those prepositions that take events as objects. "We were coming home from the lecture" assumes the coming-home event occurred after copresence with the lecture and "going to the lecture" assumes copresence with the event is anticipated in the future of whatever is currently being described.

At can be used with verbs of projection; you can throw something at a wall, or look at the wall, or, if you want to, even yell at it. In the motional context, at is not fussy about what objects it will take. Its sense is essentially that of toward, except that the path of toward is explicitly incomplete. The destination is clearly the

semantic object; the path is presumed straight in the absence of other constraints (a physical object's path in air may be straight with respect to the horizontal plane but not necessarily with respect to height, of course). In the case of projectiles, the target may be the vertical cross-section of the object plus a small amount (I once threw what I thought to be an empty milk carton at a friend, with every good intention of just missing him--even given that intention I can still say I threw "at" my friend), although strictly interior restrictions clearly seem more likely defaults for vertical targets.

It would appear, then, that the directional uses of at pose no particular problem for the types of modeling processes that are the particular province of this paper.

MOTIONAL IN AND ON--INTO AND ONTO

The prepositions in and on have been treated essentially as place-locatives, and in their most prototypical uses this seems to be their function. Clearly, though, they can be used as prepositions of motion, as in, for example, "put the toys in your toybox," "they got in the car and drove away," "come in the house. now," "the geese were flying in that direction," "we got on the bus," "I put it on my desk," "she tossed her books on the bed." With the exception of the "in that direction" example, the interpretations of these descriptions are exactly the same as if into or onto had been used.

Children from a wide variety of languages and cultures appear to learn to use the prepositions into and onto fairly rapidly after mastering in and on in their prototypical senses as place locatives [Slobin, 1973]. This seems natural enough, since interpretation of into and onto seems to consist of little more than adding the location restrictions of in or on to the destination description of the path. Consider the following: toys are put into a toybox or onto a table; a person runs into a room; a girl walks into a field; paint goes on (onto) a wall; a marble rolls onto a rug. As usual, the verb and characteristics of the subject determine most of what is known about the path. The prepositions mark the destination, essentially by giving locative restrictions of the subject relative to the

semantic object. These restrictions for into correspond to those for place-locative in, and for onto correspond to those for place-locative on.

Of course, for verbs of projection, it is the projective object of the verb which is the semantic subject of the preposition, so that in "he aimed the gun into the room" the locative restrictions on the destination are still the interior of the room, but the path is not that of the person with the gun nor that of the gun, but rather of what I have chosen to call the focus of attention.

There are, of course, exceptions. "The geese were flying in that direction" clearly does not provide a locative restriction for a destination. Rather, it restricts path directionally, much as toward and away from do.

Furthermore, there are at least two exceptions to handling into as motion with specified destination. The phrase "face into the wind" appears to imply direction (notice that the counterpart phrase is "face away from the wind"); this may be an idiom or it may be a function of the characteristics of the object: a ship may turn into the wind, but when a car turns into a driveway we have a normal interpretation--the location restrictions on the car at the end of the described motion are the boundaries of the driveway in the horizontal plane.

The other exceptional into occurs in phrases like "run into" or "bump into" or the like, where the semantic object is often solid, as opposed to a container.* Examples are "Bob wasn't looking and walked into a tree," "She backed into the wall," "Groping around in the dark I bumped into one piece of furniture after another." The descriptions still specify destinations of sorts, but the relevant relation between subject and object is contiguity rather than an interior relationship.

* Recall that containers are objects with open interiors, like rooms or boxes. To "run into a room" clearly invokes the normal sense of into. However, "bumped into" plus a container still implies this latest meaning of into.

TOWARD

Toward is the counterpart of away from. Its semantic object is the ostensible destination, but the path is explicitly incomplete--the destination is used to bound the orientation of the path and to put an upper bound on the extent of the path: Toward does not imply that the semantic subject reaches the destination or that the location restrictions apply to the semantic object, in contrast to to which may have such an interpretation.

For verbs of projection, toward can indicate that a path does not extend as far as the destination--an archer can aim toward a distant target in the expectation that his arrow will fall short. But toward may also contrast with at in the following manner: to aim at a target, as has been mentioned, tends to default to the interior of the vertical cross-section of the target; to aim toward the target can imply a strong relaxation of the "small" distance around the cross-section usually attendant on at. That is, toward can imply not so much that the path does not reach the vertical plane of the target so much as that the destination restrictions include a non-negligible area around the target, like taking sloppy aim or pointing vaguely.

NEAR

Denofsky [1976] presents an extended discussion of the uses of near and a set of formalisms to interpret a nearness relation between two physical objects. The following are comments on what I perceive to be some of the important points of the paper relevant to the present topic.*

Clearly the interpretation of near depends on, among other things, the dimensions of both of the objects related. Moreover, there can be an implicit scale in effect, from microscopic to astronomical, and particularly where the dimensions of the objects are extremely large or extremely small with respect to the scale in effect, the scale exerts considerable effect of the interpretation of near. Other considerations are population density of objects, and whether the two related objects are part of an orderly, recurring pattern.

Of particular interest were Denofsky's recourse to "standards". He points out, for instance, that our judgment of whether a car is near us on a freeway has one standard for normal driving and another for a traffic jam. He also suggests that in the case of a little boy on a subway platform cautioned by his mother not to go near the tracks,

* My interpretations may be rather free--my apologies to the author for any inadvertant misrepresentations.

the judgment of nearness may be made on the basis of how close other people are to the tracks. To me this suggests that where there are a number of similar objects, nearness may be based on a norm for the objects, where the norm is determined either directly from visual evidence when the evidence is present as in the preceding examples, or from inspection of a mental model or mental image, for well-known situations. As Denofsky points out our standards for normal, comfortable spacing of furniture in a mental model of a livingroom are quite different from our standards for a bathroom. Even though the sizes of the furnishings are not notably different, the norms for distances are, and these norms affect criteria for nearness.

Finally, Denofsky points out that standards of nearness are different for a man walking and a man driving. It seems to me that what is happening here is not so much that there are two different standards as that (as so often is the case with motion) the criteria have switched from distance to time. Whether a point on a path is near or not depends on how long it will take to get there, and that depends on the speed with which the path is traversed. This perhaps odd-seeming time dependency is not just a quirk that crops up with the preposition near but a persistent tendency for objects in motion, particularly when the object in motion is human. (See for example the discussion of from/to on page 82.)

The notion of "appropriate distance", which is used with respect to several prepositions in this model, has much in common with Denofsky's account of near. The same characteristics of the objects are involved; differences appear to be differences of degree--for to and at the appropriate distance would be, in Denofsky's terms, very near, or even extremely near (as opposed to quite near or merely near, all of which have distinct interpretations in his system); for above, the appropriate distance is much more relaxed--in my understanding it would not even qualify as near in Denofsky's system.

OTHER PREPOSITIONS

Many prepositions yield to analyses not much different in character from those already described; some of these have been hinted at in other portions of this paper. For example, through signifies that a path is interior to the boundaries of the semantic object, but the source and destination of the motion are not. Prototypical above indicates that the semantic subject is vertically upward from the semantic object and within an appropriate distance (somewhat larger than that associated with to and at) from the object, and moreover, either the vertical projection of the horizontal cross-section of the subject is interior to that of the object or vice versa (usually one can tell which should be understood as interior to the other by comparing characteristic sizes). The two-dimensional version of above used in picture-analogies simply takes the abstraction of these concepts in the vertical plane. And finally, directional above can be captured in many instances by using as boundaries an appropriate angular interval from the vertical, with the semantic object as origin, of course.

Over, under, beneath, below, and so on generally are variations on above or its counterparts. An obvious exception is the covering sense of over (e.g., "she put a handkerchief over her mouth"), which seems to be amenable to analysis in terms of contiguity and/or interior relationships of cross-sections of the objects involved.

Underneath and under can signal the same relationships except that the roles of the semantic subject and object are switched.

Under also has an interesting variation for objects like tables and desks (usually furniture with legs). The basic meaning for these objects appears to be under the body of the object but not under the legs (e.g., "I found the you-know-what under Jenni's chest of drawers"). Trees also are often treated as if their trunks were negligible (Saint Luke tells us Zacchaeus climbed into a sycamore tree; when we speak of sitting under a tree we generally mean under the part of the tree delineated by the branches and leaves), but otherwise we usually mean under the entire object when we say under. (My four-year-old daughter recently asked why Little Boy Blue was under a haystack--the idea obviously bothered her, even though she has seen plausible illustrations of the nursery rhyme. Clearly her default sense of under does not give a reasonable interpretation at this time.)


Some of the examples in other parts of this paper have used many objects or the location of a path over time as potentially indicating a surface. The covering sense of over, for example, usually expects a surface characteristic in the semantic subject, but we have used an example where the covering "surface" was a large number of flying mosquitoes. The same sort of extension often applies to the

semantic subject of around, which frequently describes a situation in which multiple objects are treated as one entity which is then what is understood to be around the semantic object. The location restrictions are once more subject to an appropriate distance consideration, hence dependent on both subject(s) and object. Around, by the way, does not have to mean completely encircling, depending on physical constraints of the situation and the population density of the semantic subject, if it is made up of many discrete parts. For example, there may be many people around a popular guest at a party, but if he is standing next to a wall it is not expected that the people do more than half-surround him. Similarly, if the population density of the semantic subject is low, large breaks in the surrounding "surface" are acceptable.


Cresswell [1978] cites an example of around with a different sense: "We walked around the museum," where he glosses around as "frequently deviating from a straight line." In my own idiolect, all the examples of such an interpretation for around occur with in, as, "The children ran around in the department store," for which the analysis would seem to be an unspecified complex path bounded by the interior of the semantic object. (An analogous interpretation is necessary for sentences like "The butterflies flitted above the flowers"; obviously the interpretation is not unique to around.)

Up and down as motional prepositions of course depend on features of their objects ("up a hill," "down the well") but seem to hold no real surprises. Out of is a counterpart to into, and so on. Indeed, of the remaining prepositions, the one whose basic interpretation seems not to have been reasonably begun by our analysis so far is across.

The key to the analysis of across is attention to axes of symmetry. Consider children arranged in a circle. The child across the circle from Jenni is the one who is at the diametrically opposite point of the circle. If I am at one wall of a rectangular room, then if I speak of something across the room from me, I mean something within an appropriate small distance of the boundary of the room on the opposite side of the relevant axis of symmetry.



If I am at a corner of a rectangular room, then I may mean the diametrically opposite corner of the room, or either of




the two other corners, since they are at boundaries symmetrically opposite to my position also.

In practice, of course, we often relax this rigorous definition (which perhaps is what we have in mind when we say directly across) and tend to allow anything near the boundary such that an axis of symmetry is between it and the semantic object. Consider what is meant by "across the meadow from me." If the meadow



is as shown, fairly elongated, then we may mean in symmetric position on the other side of the partial axis of symmetry, but we may just as well mean any point near the boundary of the other side of this partial symmetry axis from us. For an irregular shape as shown, this still provides reasonable location restrictions.



Unfortunately, there is neither time nor space to cover in one paper all the aspects of all the English prepositions. I do feel that the principles disclosed thus far carry considerable descriptive power and do much to reveal the characteristics underlying the uses of prepositions in general and the objects they relate.

Chapter three

THE IMPLEMENTATION

A small computer model was written and run on the DEC-10 system at the Coordinated Science Laboratory of the University of Illinois to illustrate some of the features of the theory. Programs were written in the MACLISP dialect of LISP, and included about 45 functions; these were for the most part short functions--average length for a pretty-printed function was less than twenty-five lines. Data for the implementation consisted of twenty-two "definitions" of objects and the definitions of the prepositions themselves. Input to the program consists of English sentences--either statements or questions. Statements are expected to be either "naming" statements ("Tweety is a bird" or "Volume-1 is a book") or locative statements ("A book is on a table", "In the room is a bed"). Output is either a set of coordinates for each object in the "mental model" or a response to the question. Response to a naming statement does not generate significant output.

Program data

The system is initialized by running a file called PREP.MIC which reads in the remaining files needed, including the data files. Some sample definitions of visually perceptible objects follow:

```

(JENNILYN INDIVIDUAL
  (CHARACTERISTIC-SHAPE ((CROSS-SEC 0.33 0.22)
                        (HEIGHT 1.0)))
  (WEIGHT 20.))

(HOUSE PROTOTYPE
  (CHARACTERISTIC-SHAPE ((CROSS-SEC 15. 10.)
                        (HEIGHT 7.)))
  (FEATURES (CONTAINER)))

(TABLE PROTOTYPE
  (INSTANCE-OF FURNITURE)
  (CHARACTERISTIC-SHAPE ((HEIGHT 0.75)
                        (CROSS-SEC 1.2 0.9)))
  (FREE-SURFACE (((PLANE HORIZONTAL)
                    (FREE-DIRECTION /+Z)
                    (HEIGHT 0.75)
                    (DIMENSIONS 1.2 0.9))))
  (WEIGHT 25.0))

(BOX PROTOTYPE
  (CHARACTERISTIC-SHAPE ((HEIGHT 0.3)
                        (CROSS-SEC 0.3 0.3)))
  (FEATURES (CONTAINER OPEN-TOP))
  (WEIGHT 1.0))

(FLY PROTOTYPE
  (CHARACTERISTIC-SHAPE ((HEIGHT 3.0E-3)
                        (CROSS-SEC 3.0E-3 5.0E-3)))

(WALL PROTOTYPE
  (CHARACTERISTIC-SHAPE ((HEIGHT 2.5) (CROSS-SEC 4.)))
  (FREE-SURFACE (((PLANE VERTICAL)
                    (HEIGHT 2.5)
                    (WIDTH 4.)))))

(CEILING PROTOTYPE
  (CHARACTERISTIC-SHAPE ((CROSS-SEC 3.6 4.)))
  (FREE-SURFACE (((PLANE HORIZONTAL)
                    (FREE-DIRECTION /-Z)
                    (DIMENSIONS 3.6 4.0))))
  (FEATURES ((CHARACTERISTIC-HEIGHT 2.5)))

(SHELF PROTOTYPE
  (INSTANCE-OF FURNITURE)
  (FEATURES ((CHARACTERISTIC-HEIGHT 1.0)))
  (WEIGHT 1.5)
  (CHARACTERISTIC-SHAPE ((HEIGHT 0.025)
                        (CROSS-SEC 0.2 1.2)))
  (FREE-SURFACE (((PLANE HORIZONTAL)
                    (HEIGHT 0.025)
                    (DIMENSIONS 0.2 1.2)
                    (FREE-DIRECTION /+Z)))))

```

Units of measurement are meters and kilograms.

At initialization, the components of the definitions are added to the property lists of the objects. In addition, the components of the characteristic-shapes are used to create a simple "mental picture" of the object, in the form of coordinates of a right parallelepiped. The coordinates are always given in a particular order: bottom front right, bottom front left, bottom back left, and so on. The object is assumed to be in canonical position and the program potentially has access to its canonical top, bottom, front, back, left and right sides. This permanent mental picture is kept under a "local coordinates" property, with the bottom right front taken as local origin.

The definitions specifically single out planar free surfaces on a free-surface list, since it is impossible to judge from the representation whether a planar surface is a characteristic of the object itself.

Also included in the definitions is an indication of whether the object is essentially hollow as opposed to essentially "solid" throughout. The surfaces of the latter are the boundaries of matter; the surfaces of the former enclose space. The feature CONTAINER is used to indicate an object whose interior is canonically empty. Another feature applies to CONTAINERS only and is used to indicate whether they are OPEN-TOPped or not.

Height normally refers to a dimension of the object. Height as part of the description of a plane horizontal free-surface indicates where the plane is located relative to the object's bottom. And CHARACTERISTIC-HEIGHT as part of a feature list indicates that an object normally would be found at a given height above the default ground-level (either the floor or the actual ground). Otherwise a clock placed randomly on a wall might end up very close to the floor. After toying around with the program for awhile, it became obvious that we have such default characteristic heights for a number of items--clocks and windows, of course, and to a lesser extent shelves, certainly counters, cabinets, and so forth.

To facilitate the handling of English input from the user, a parsing system called LINGOL, developed at MIT, interfaces between the user and the modeling portion of the program. Part of the initialization process is to create dynamically from the data the definitions that LINGOL requires to recognize the nouns and prepositions. A similar process occurs during the interactive portion of a run if a new name is introduced by a user.

Input sentences

Once past the initialization phase, the program enters a loop in which the user types in input, LINGOL parses the user's sentence, and the model responds.

The user may type one of three kinds of input: a "naming" sentence, a location description, or a question. Suppose the user types "Tweety is a bird." The word Tweety is unknown to the system prior to this input. On discovering that the parse of the sentence fits a naming pattern, the program creates a syntactic definition so that LINGOL will recognize the name in the future; furthermore, it creates an individual bird (up until now the system has known only a prototypical bird) whose characteristics are assumed to be identical to those of the prototype. (Under different circumstances than these, say if Tweety's size were already known, he would not inherit the size of the prototypical bird. Only those characteristics not already specified are inherited from the prototype.)

This facility for creating individuals and naming them was developed for two reasons: it allows introducing and naming people and animals at will, as long as the system knows the appropriate prototype; and it provides a simple means of distinguishing between several of an item--say three chairs or ten books. It is possible to type in ten sentences to the effect that volume-1 through volume-10 are books, and thereafter the system would know about ten individual books, and be able to take account of their locations separately.

A second type of input from the user is a location description. Suppose the user types "A book is on a shelf." An individual book and individual shelf are created (later references to "the book" will be assumed to mean this particular book) and LINGOL fills a subject register, object register, and preposition register before calling the function that actually creates the "mental model". This function chooses a point on the semantic object (in this case the shelf) as the origin of a "global" coordinate system, invokes the definition of the preposition on to find what restrictions are placed on the location of the semantic subject, and places the semantic subject in the global coordinate system. The location of the subject and object are recorded under their "global coordinates" property, and the location restrictions on the subject are recorded under the subject's location restriction property. Response to the user consists of printing out all the global coordinates of each object affected by the description (sometimes just the subject, sometimes every object already in the model) in canonical order, canonical bottom right front corner first. (Hence, if an object, say a fly, were upside down, one of its highest points--whichever corresponded to its right front foot in the upside-down position--would be printed first.)

The third type of input accepted is a question which is interpreted as asking for a judgment of whether a preposition suitably describes the relationship between a

subject and object. For example, if the input is "Is volume-10 on the table?" the LINGOL portion of the program fills the registers as usual, but since the input is recognized as a question, rather than call the modeler a routine is invoked to determine whether the requirements for on are fulfilled by the relationship of volume-10 and the table. The response is either "yes", "no", or "not directly, but...."

The prepositions and primitive relations

Prepositions are essentially short blocks of LISP code.

The definition of ON is shown below:

```
(ON (SUBJECT OBJECT)
  (PROG (SURFACE-LIST OBJECT-SURFACE)
    (COND
      ((NULL (GET SUBJECT 'WEIGHT))
        (SETQ OBJECT-SURFACE
          (FIND-AVAILABLE-SURFACE OBJECT)))
      ((GET SUBJECT 'WEIGHT)
        (SUPPORT SUBJECT OBJECT)
        (SETQ SURFACE-LIST (GET OBJECT
          'FREE-SURFACE))
        (COND
          ((SETQ OBJECT-SURFACE
            (PREFER '((PLANE HORIZONTAL)
              (FREE-DIRECTION /+Z))
            SURFACE-LIST))
            (COND
              ((AND NEWORIGIN (THREEDIM-P OBJECT))
                (LOCATE-OBJECT-OF-ORIGIN
                  OBJECT
                  (LIST 'SURFACE OBJECT-SURFACE))))
              ((COND ((AND (THREE-DIM OBJECT)
                (NOT (MEMQ 'OPEN-TOP
                  (GET OBJECT
                    'FEATURES))))
                  (SETQ OBJECT-SURFACE (TOP-OF OBJECT))))
                ((SETQ OBJECT-SURFACE (CAR SURFACE-LIST)))
                ((SETQ OBJECT-SURFACE
                  (FIND-AVAILABLE-SURFACE OBJECT))))
              (COND ((GET SUBJECT 'FREE-SURFACE)
                (CONTIG OBJECT OBJECT-SURFACE SUBJECT))
```

(T (CONTIG OBJECT
 OBJECT-SURFACE
 SUBJECT
 (BOTTOM-OF SUBJECT))))))

On is faced with two decisions: it must decide which surface of the object the subject is contiguous to, and it must decide which side of the subject is contiguous to the object. The two outer conditional clauses perform these functions.

If the subject does not behave normally with respect to gravity (shadows, visual patterns, thin films of liquids and many insects exhibit gravity-defying behavior) then any available surface of the object will do.

If the subject is under gravitational constraints, then the routine looks for one of four possibilities: in order of preference, 1) a horizontal plane in the object, 2) if the object is three-dimensional and is not an open-topped container, then the top of the object, 3) failing either of these, then any planar free-surface, and finally 4) any available surface. In any of these cases, the object requires support and by supposition the semantic object furnishes it.

Having found the surface of the object, on looks for a probable surface of the subject. The check to see if the subject has a marked free-surface is actually a back-handed way to see if the subject has a preferred orientation. If it has, the preferred orientation is presumed to be the

canonical one, and on passes CONTIG, not a surface of the subject but the entire subject, thereby instructing CONTIG to translate the subject in whatever direction necessary to bring it into contact with the object-surface indicated, but not to rotate it in any way. On the other hand, if the subject has no preferred orientation, on selects the canonical bottom of the subject.

CONTIG essentially does what has already been suggested: if given the entire subject, it merely performs a translation of the subject to the object-surface. If given a particular surface of the subject, it rotates the subject if necessary before translating it. The model does not know about such relationships as being partially on something (for which translate "part of the subject is on the object") so CONTIG makes the simple assumption that one of the two surfaces involved is two-dimensionally interior to the other.


The definition of the preposition in has to decide if it is dealing with a container, whether the container is open-topped, and whether the subject behaves normally with respect to gravitational constraints. It then calls one of the INTERIOR functions and, sometimes, CONTIG (when the subject is assumed to be in the bottom of a container, for instance).

The INTERIOR functions could have been written as one function which interrogated the characteristics of the arguments passed to it. They were kept separate to make clearer what the calling functions were expecting. At present, the system has two- and three-dimensional interior functions, which restrict the location of their subject with respect to a plane of their object or the volume delineated by the object, respectively.

Some examples

Suppose a user types: A book is on a shelf.


The LINGOL portion of the program echos (BOOK IS (ON SHELF)); this echo is not used for anything by the program and is provided mostly as a double check for the user against the misinterpretation of input by the syntactic analyzer. An individual book and individual shelf are created; the modeling portion of the program records the location restrictions for the book, chooses a location for it and gives the user the global coordinates of the book and shelf. These correspond to the following illustration:



Notice the upper surface of the shelf is chosen for the $Z=0$ plane--leaving all but the top surface of the shelf below the plane with negative coordinates. It is entirely possible that we are about to enter an extended description of many objects, all of which are on the shelf or above it, in which case treating the top surface of the shelf as our basic horizontal plane makes sense.


Suppose now that the next sentence is: The shelf is on a wall.

LINGOL echos (THE-SHELF IS (ON WALL)). (We will omit the echoed response after this example.) Since the object that has been serving as our origin has just been treated as a semantic subject and related to another object, the location of the shelf and everything related to it is accordingly revised. We use the symbol $\curvearrowright$ to indicate the origin of the global coordinate system in the illustrations.



Actually, as far as the implementation is concerned, it is an accident that the long edge of the shelf is aligned with the wall. But for the order in which the dimensions of the shelf were given in the data, the program might just as well have set the short end of the shelf flush with the wall. However, the program would not place the top surface of the shelf against the wall, even if it were bare, since that surface is marked as characteristically horizontal. Incidentally, were it not for the fact that shelves are marked as having a characteristic height (a little bit of "world knowledge") the program would have put the shelf considerably lower.

Suppose an input were: A light is on a ceiling.




A corner of the ceiling is taken as the origin. Since the ceiling has a marked free-direction vertically downward, the light ends up on the correct side of the ceiling surface. Notice that, while people would ordinarily put the light in the middle of the ceiling, the program doesn't know enough about ceilings and lights to do so.

finally, let's follow an extended example.

Input: A glass is in a box.


Resulting model:



Comments: The glass has weight, so it ends up not only in the box, but at the bottom of it.

Input: The box is on a table.


Resulting model:



Comments: There is only one individual box known to the system, so the phrase "the box" can be interpreted with no difficulty. Notice the surface of the table is taken as the basic plane for the discussion so far, rather than putting the origin at, say, a point at the bottom of the table.

Input: The table is on a floor.

Resulting model:




Comments: "A floor" sounds strange, but the system doesn't know for the present that tables are almost always on floors, so mentioning a particular table does not allow it to presuppose a particular floor that it could reference as the floor.

As is probably becoming obvious, the model does not choose locations randomly. Rather, it tends toward a particular corner. This choice was made in hopes of avoiding the "findspace" problem when several objects must be located on one surface.

Input: The floor is in a room.

Resulting model:


(3.6, 4.0, 2.5)



Comments: Again, the model doesn't know that a floor is part of a room. Naturally, a default-sized floor exactly fits a default-sized room, but the model has to know that a floor belongs at "ground level" or it would try to put the floor at a more or less arbitrary level in the room. While this particular sentence sounds unusual, it is natural to speak, say, of "the floor in Jonathan's room".

Input: The room is in a house.


Resulting model:



Comments: It is not very evident from the illustration, but the room is actually several feet off the ground in the model. Obviously, this "goof" could have been fixed (there are several natural solutions that would not have involved extending the capabilities of the model), but it was left in for two reasons--it illustrates what the model does with an in relation involving a weightless object, and secondly, it was pointed out to me that if the building had been a hotel rather than a house, a room in the same relative position in the larger building would have seemed quite reasonable.

Input: The house is in a field.

Resulting model:



Comments: A house has no weight, either, as far as the model is concerned, but a field is a two-dimensional object, and the in relation implies contiguity under those circumstances.

Questions

Suppose after all this the user types: Is the box on the table?

The response from the system is YES.

To the input: Is the box on the floor? the system responds NO.

Is the glass in the box? YES

Is the box in the room? YES

Is the glass on the table? NOT DIRECTLY, BUT ON IS STILL AN ACCEPTABLE DESCRIPTION.

The program answers these questions by directly interrogating the three-dimensional model, not by knowing that, say, if A is on B and B is in C then A is probably in C. At no time did we say that the box was in the room. But thanks to the sizes of boxes and tables and the locations of floors relative to the rest of a room, there is no question but that the box must be in the room in the most rigorous sense of the word.

It is possible to distinguish between a glass on a tall object in a box and a glass on a small object in the box. If the tall object were large enough that the glass was exterior to the box, then this sort of model could reasonably balk at calling the glass in the box--or at least hedge, as a person might. A system built on the sort of inference rules mentioned above could have trouble distinguishing between these cases.

Incidentally, that "NOT DIRECTLY, BUT..." response was not entirely anticipated. The response is a canned phrase, of course, but the piece of program that produced it was written with multiple on relations and furniture in mind. (Recall that items of furniture serve as landmarks and tend to collect on relations.) The possibility of one object in

another object on a piece of furniture was not considered at the time.

The direct query of a model is not without its pitfalls, as will be shown in Chapter four. However, when used judiciously, it does seem to hold promise as having certain advantages over competing methods.

Chapter four

WHAT CAN WE CONCLUDE FROM ALL THIS?

Perceptual attributes

If there is one point I wish to emphasize about this work, it is that the prepositions are not used in a perceptual/cognitive vacuum: it is the prepositions in conjunction with perceptual attributes of the objects related by the prepositions that enable the construction of a mental model of the spatial relations of the objects. For example, it would appear that a major function of the preposition on is to signal contiguity with a suitable surface. The semantic object of the preposition contributes to the total interpretation--it may be that the semantic object has only one eligible surface. The semantic subject may influence the interpretation--if it is a "dead weight", like a book, there is a preference for finding in the semantic object a nice, safe horizontal plane with free direction upward, or as second choice something roughly horizontal, and failing that for the topmost part of the semantic object; on the other hand, if it is a fly, which has very free gravitational restraints, any surface will do. The end result is that on, rather than having four or five spatial "definitions" ends up having one "definition" having various components, each of which is or is not responsive to a given characteristic of the semantic subject or object.

Possibly it could be argued that the two methods are equivalent--I could have written many versions of on, ON1, ON2, ON3, and so on, each having a complex conditional that had to be satisfied by the subject and object. I have to admit that I would find such a solution intuitively less satisfying at this time, partly because even though people readily agree that a light on the ceiling, a shelf on a wall, and a toy on a table involve different spatial relations, it is not clear that they feel it is the on that is different in each case. The concept of one spatial definition which is responsive to the attributes of the objects seems intuitively more satisfactory.

In any event, the perceptual characteristics of the objects would still have to be taken into account--they cannot be ignored! This may seem a very simple point, but to my knowledge every previous attempt to account for the use of the prepositions has done just that--ignored the characteristics of the objects involved. Clearly, we cannot afford to dismiss them.

Once the decision is made to take account of the perceptual characteristics of the objects, it quickly becomes clear that a small number of attributes yield a considerable amount of interpretive power.

To be frank, this study was not initiated with the expectation that the interpretation of prepositions would hinge so strongly on the objects related, either. It was only after a number of false starts that the key became clear. Once noted, it was interesting to see parallels in other facets of English and in other languages. Whorf [1956] discusses what he calls English cryptotypes, one of which is "the transitive verbs of a covering, enclosing, and surface-attaching meaning." He argues for such a category on the evidence that the prefix un-, used to denote reversal, may be used only with verbs from the category. We may say "uncover, undress, unfold," but not "unbreak, unlift, unheat."* Moreover, Whorf points out that a coined verb flimmick can take un- as a prefix if it means, say "tie a tin can to" (e.g., "he unflimmicked the dog") but not if it means "put together" (e.g., "he unflimmicked the set of radio parts").

It would appear then that native speakers unconsciously categorize English verbs according to, among other things, perceptual attributes not unlike those utilized in comprehending prepositions.

* In a footnote to Whorf's paper, it is pointed out that the adjectives "unbroken, unheated" and so forth express condition, not a reversal of action; hence they are not counter-examples to the premise.

In a variety of diverse languages--Taraskan, Bantu, and several American Indian languages, among others--various markers classify objects into categories, among which are enclosed/nonenclosed, solid/plastic/liquid, and rigid/non-rigid discriminations, and various classifications based on relative width versus height ([Basso, 1968], and [Denny, in press], cited in [Pylyshyn, 1977]). It is especially striking that the categories mentioned recur in widely diverse cultures.

At least from the perspective of this work, the perceptual attributes of the objects are not held to be semantic markers. To the extent possible, the modeling process in the computer implementation is based on a three-dimensional model of the object (in canonical position prior to being placed in context). Where keywords were used in the interpretation process, it was to differentiate planar surfaces (which are otherwise indistinguishable from other surfaces in the simple representation scheme used) and to distinguish between totally enclosed volumes (rooms, mailboxes, and cereal boxes, for example) and containers with open tops (bowls and open-topped boxes).^{*} Clearly, for

* There was also world knowledge represented in knowing for example that certain objects, like clocks or windows, occur at characteristic heights, and that some things, like flies, have relaxed gravitational constraints. There is little likelihood that this sort of information would be mistaken for semantic markers, however.

humans and models with a richer descriptive scheme, such distinctions would be available by direct appeal to the three-dimensional image in memory.

At least for the interpretation process, then, we don't need keywords or semantic markers for those objects that are "regular". The information required is not linguistic in nature or necessarily attached to the surface words used in descriptions. When presented with an object whose name was totally unfamiliar, we could nevertheless use the prepositions easily with the new word and easily be understood, provided that the perceptual attributes of the object are known. It is highly unlikely, then, that the necessary information is stored at the level of linguistic knowledge or as lexical entries (cf. Whorf's flimmick example).

Location restrictions

It seems clear that only very rarely does a discription enable us as listeners to interpret the precise location of an object. Even in search procedures, where the goal is the precise location of an object in the real world (and it does seem to be true that the most detailed descriptions of locations tend to be in search procedure situations), what is given is a series of landmarks, each of which may be very imprecisely located; however, each landmark is deemed to be sufficiently outstanding in context that further description is unnecessary and perhaps wasteful of time and energy.

In the non-search-procedural case, what happens is that the description yields bounds on locations, and the listener is free to construct a mental model placing objects at will within these bounds. In the simplest case--a semantic subject, preposition, and semantic object--the location of the semantic object is presumed "given", the listener has some idea of the shape and size of the subject and object, and combining these with the preposition yields a set of restrictions on the location of the subject. The listener, in visualizing the scene in his mind, probably goes one step further and actually places the subject at some location within the boundaries specified. The result is what I have called the "mental model".

Inferencing

One of the nice features of this "mental model" is that it holds out hope for doing inferencing and deduction by direct reference to the model, under optimum conditions, and by reference to the model plus the location restrictions under other circumstances.

As was shown in a previous section, a very simple set of procedures, by consulting the model resulting from a description of a glass in a box on a table on the floor was able to decide that the box was on the table, the glass was on the table, the glass was not on the floor, and so on, without recourse to inference rules of the following kind: if A is <prep> B and B is in C then A is <prep> C for any of

one set of prepositions but not for prepositions outside the set. I find the potential for avoiding such rules encouraging.


Two cautions are in order, however. In interpreting a description (building the model in the first place), it suffices to place objects in simplest possible relationships. If a description mentions a book on a desk, we probably visualize the book as being directly on the desk. In the presence of more detailed knowledge, ofcourse, such a default may be contra-indicated; if I were to tell someone they could find an item on my desk, it is unlikely that they would visualize the item as directly on the desk, if they have ever seen the desk. The reverse process--judging from a mental model whether a particular preposition is an appropriate description of the relation between two objects--is not so simple. In deciding whether "above" is an acceptable description, for instance, there is little question when one object is directly above the other, but clearly the word is acceptable even when the direct case is not applicable, and deciding these more marginal cases often leads to a lot of hedging, even from native speakers.

The simple implementation of this theory never had to deal with extremely cluttered surfaces (partly to avoid the findspace problem), so it was always able to assume on meant directly on in constructing the spatial model. However, in deciding whether on was an appropriate description of an

existing relation, it did need to know whether the semantic object of the proposed relation was a piece of furniture under some circumstances, since items of furniture serve as landmarks and horizontal surfaces on furniture tend to collect on relations. Knowing this extra bit of information, it could decide, for instance, in the example of the stack of encyclopedias on a desk that volume 10 was on volume 9 but not on volume 3, and that volume 10 was on the desk though not directly on the desk. If the same stack of encyclopedias is put on the floor, the implementation does not judge that volume 10 is on the floor.

It would appear then, on this rather simple evidence, that decisions as to the appropriateness of prepositions are more likely than the interpretation phase to bring non-simple strategies into play. This is not to say that the non-simple strategies are never needed for interpretation--only that there are a lot of instances where we don't need them.

The second caution is best put by describing a session with the implementation: as it happened, the particular mental model produced after "a shelf is on a wall" and "a fly is on the wall" was the equivalent of the illustration below.



Now suppose we were to ask if the fly is under the shelf. The correct answer, of course, is "I don't know", since on the basis of the description the fly might be under the shelf, but it might be elsewhere, too. (If the implementation had been set up to try putting the fly under the shelf, and, subsequently, at a place violating the location restrictions of under, in response to the question, it would have found neither violated the location restrictions placed on the fly by the original description and hence would have had reason to suspect that it couldn't answer the question one way or the other.)

Clearly, then, the simple expedient of directly consulting the constructed model is a little too simple. The more freedom of choice a model presents in actually choosing the location of an object, the more incidental any relations between various objects may be. In the end what we know are the location restrictions and it is based on

them that we need to make judgments.

Nevertheless, there still seems room to hope that we may be able to avoid sets of inference rules of the "if A <PREP1> B and B <PREP2> C then A <PREP2> C " sort.

Regularities

For all the hedges and caveats of the preceding pages, it was evident from the implementation that paying attention to a very small set of attributes of objects yields an astonishing amount of descriptive power. The attributes included a very rudimentary surface description, the concept of a free-surface with associated free-direction, the essential "emptiness" of containers, some notion of gravity, of contiguity, of the interior relation in two or three dimensions, of partial axes of symmetry, some awareness of scale, and a coordinate system with marked vertical direction. Clearly, these concepts do not handle all cases of descriptions using place locatives. It might even be said that they do not handle some of the most common cases (we will come back to this in a moment). But they do handle the most typical cases-- the regular uses of in, on, and the other prepositions--the uses we are most likely to think of as standard. In so doing, they capture much of the descriptive power of the prepositions.

Why then could it be said that they do not handle some of the most common cases? Recall, for a moment, that the most commonly occurring English verb is the verb be. Recall further that the verb be is also the most irregular verb of the language. It is commonly accepted that there is a direct connection between these two observations. Moreover, as a general rule, the most frequently occurring verbs in English are also the most irregular. Because of their daily rehearsal, we are able to absorb them, irregularities and all, over a period of time. As everyone knows, children tend to regularize verbs early in their language-learning experience and it takes many years of practice before they gain command of the irregular forms. It seems likely that over many generations, verbs which may have been less regular but were not so frequently used simply retained the regularized form that language-learners tended to give them.

The point of all this is that most-common items tend to be least regular. Something of the same sort seems to apply to uses of the prepositions with common objects. If an object is something that we see and speak about daily, then over the period of the formative years we have no trouble absorbing an irregular construction with regard to that object. Tables, for instance, have a tendency to be treated as if they were essentially the table top--"under the table" for most objects means under the table top but definitely not under the legs. Rugs are an exception, of course, as are floors, and there are undoubtedly other exceptions to

the mini-rule of treating the table as top only. One occasionally gets the disheartening feeling that the use of prepositions to relate other objects and tables is learned on a case by case basis. Unfortunately, tables (and beds and chairs, among other things) are enough involved in our everyday lives that we cannot discount such a possibility--we certainly have had plenty of opportunities to rehearse any irregular behavior of these objects.

Fortunately, even the most common objects (including tables) appear to be regular most of the time, with most of the prepositions. It is interesting that some of the irregularities fall into classes, like classes of irregular verbs (sing, sang, sung; drink, drank, drunk; sink, sank, sunk; swing, swang, swung). For example, "the people on the bus" are actually in the bus--they aren't on the bus in the same sense that "the people on the car" would be on the car. On has the same interpretation in "on the plane", "on the subway", or "on the boat"--indeed for anything that can be boarded,* but evidently not for anything else. So at least potentially there may be classes of irregular objects.

After all is said and done, though, it is still the case that the system seems to work, and work well, for the great majority of regular objects, and even for the

* my thanks to Dave Waltz for pointing out the direct relation with the verb board.

irregular ones most of the time. It seems clear that basic understanding of the use of the prepositions is ours if only we pay attention to a small set of perceptually salient characteristics of the objects related.

List of References

- Basso, K.H. The Western Apache classificatory verb system: A formal analysis. Southwestern Journal of Anthropology, 1968, 24, 252-266.
- Berlin, B. and Berlin, E. Aguaruna color categories. American Ethnologist, 1975, 2(1), 61-88.
- Berlin, B. and Kay, P. Basic color terms: Their universality and evolution. Berkeley and Los Angeles: University of California Press, 1969.
- Brown, R., Cazden, C. and Bellugi, U. The child's grammar from one to three. In C.A. Ferguson and D.I. Slobin (Eds.), Studies of child language development. New York: Holt, Rinehart & Winston, 1973.
- Clark, E.V. On the child's acquisition of antonyms in two semantic fields. Journal of Verbal Learning and Verbal Behavior, 1972, 11, 750-758.
- Clark, H.H. Space, time, semantics, and the child. In T.E. Moore (Ed.), Cognitive development and the acquisition of language, New York: Academic Press, 1973.
- Cooper, G.S. A semantic analysis of English locative expressions. (Technical Report No. 1587), Cambridge, MA: Bolt, Baranek & Newman, Inc., 1968.
- Cresswell, M.J. Prepositions and points of view. Linguistics and Philosophy, 1978, 2, 1-41.
- Denny, J.P. The "extendedness" variables in classifier semantics: Universal features and cultural variations. In M. Mathiot (Ed.), Boas, Sapir and Whorf revisited. The Hague: Mouton, in press.
- Denofsky, M.E. How near is near? a near specialist. (AI Memo No. 344), MIT Artificial Intelligence Lab, 1976.
- Feldman, H., Goldin-Meadow, S. and Gleitman, L. Beyond Herodotus: The creation of language by linguistically deprived deaf children. In A. Lock (Ed.), Action, gesture and symbol: The emergence of language, New York: Academic Press, 1977.
- Fillmore, C.J. Some problems for case grammar. In R.J. O'Brien (Ed.), Georgetown monograph series on languages and linguistics, 24, Washington D.C.: Georgetown University Press, 1971.

- Fillmore, C.J. The case for case reopened. In J.P. Kimball (Ed.), Syntax and semantics. New York: Academic Press. in press.
- Ginsberg, H. and Opper, S. Piaget's theory of intellectual development: An introduction. Englewood Cliffs, NJ: Prentice-Hall, 1969.
- Heider, E.R. Probabilities, sampling, and ethnographic method: The case of Dani color names. Man, 1972, 7(3), 448-456.
- Lenneburg, E.H. Biological foundations of language. New York: John Wiley & Sons, 1967.
- Miller, G.A. English verbs of motion: A case study in semantics and lexical memory. In A.W. Melton and E. Martin (Eds.), Coding processes in human memory. Washington, D.C.: Winston, 1972.
- Norman, D.A., Rumelhart, D.E., and the LNR Research Group. Explorations in cognition. San Francisco: W.H. Freeman & Co., 1975.
- Piaget, J. and Inhelder, B. The child's conception of space, translated by F.J. Langdon and J.L. Lunzer. New York: W.W. Norton & Co., 1967.
- Pylyshyn, Z.W. Children's internal descriptions. In Language, learning, and thought. New York: Academic Press, 1977.
- Rumelhart, D.E. and Levin, J.A. A language comprehension system. In D.A. Norman and D.E. Rumelhart, Explorations in cognition, San Francisco: W.H. Freeman & Co., 1975.
- Schank, R. (Ed.) Conceptual information processing. Amsterdam: North Holland Publishing, 1975.
- Schank, R.C. and Abelson, R.P. Scripts, plans, goals and undersatnding: An inquiry into human knowledge structures. Hillsdale, NJ: Lawrence Erlbaum Associates, 1977.
- Slobin, D.I. Cognitive prerequisites for the development of grammar. In C.A. Ferguson and D.I. Slobin (Eds.), Studies of child language development. New York: Holt, Reinhart & Winston, 1973.
- Sondheimer, N.K. Spatial reference and semantic nets. American Journal of Computational Linguistics, 1977, 71, 1-67.

- Waltz, D.L. On the interdependence of language and perception. In TINLAP-2: Theoretical issues in natural language processing-2. University of Illinois, 1978.
- Waltz, D.L. A model for low level vision. In A. Hanson and E. Riseman (Eds.), Machine vision. New York: Academic Press, in press.
- Wilks, Y. Grammar, meaning and the machine analysis of language. London and Boston: Routledge, 1972.
- Whorf, B.L. A linguistic consideration of thinking in primitive communities. In J.B. Carroll (Ed.). Language, thought and reality. Cambridge, MA: The M.I.T. Press, 1956.

Appendix

A WORD ABOUT PRIMITIVES

In spite of the usefulness of primitives in models of knowledge representation and natural language understanding (they are central to successful and apparently adaptable systems by Wilks [1972], by Norman, Rumelhart and the LNR group [1975], and several systems by Schank and his collaborators [Schank, et al., 1975], and [Schank and Abelson, 1977]), there seems to be a genuine reluctance on the part of many to the use of primitives. Although the reason for the resistance seems to be hard to express precisely, in broad terms the objections seem to boil down to "it's not natural." In other words, there's no convincing evidence that people have primitives.

Actually, there are a number of theories in psychology and linguistics which rely on the notion of innate concepts and processes. among them ideas from the Gestalt psychologists, child psychologist Piaget, and linguist Chomsky. The Gestalt psychologists held that all humans are born with built-in principles for organizing the mass of perceptual information which the senses constantly provide. Chomsky pointed out that by the age of four most children have mastered the components of the mature grammars of their native languages, having only to increase the sophistication

with which they combine the components as they grow older. Given the complexity of this learning task and the speed with which it is accomplished, Chomsky postulated that the fundamental language-learning processes must be innate--that humans are born with a predisposition toward learning language and with the fundamental concepts necessary for that learning already in place and ready to function (numerous arguments by Lenneberg and Chomsky for the necessity of postulating innate language-learning processes may be found in Lenneberg [1967]). More recently psycholinguists have argued that what is innate is not so much language-learning specifically, but a range of cognitive concepts, among which are some which facilitate language-learning. In so doing, they have approached the theories of Jean Piaget who for half a century has maintained that cognitive development in children is a progression through a series of stages: some children may progress through the stages more rapidly than others, but the order of progression is invariant [Ginsburg and Oppen, 1969].

Obviously, concepts born in humans or otherwise universal to them ought to be primitives in models of cognition. Unfortunately, with the exception of the work of Piaget, there is little experimental evidence to support the theories directly--they remain simply theory, at least insofar as the innate processes are

concerned--powerful. and. to many. convincing. but disappointingly uninformative as to what the universals are, if indeed they exist. Even the results of Piaget's work with children are open to considerable interpretation, at least in part because it is several levels of abstraction removed from what modelers of cognition are interested in (using a baby as an informant creates some unavoidable difficulties).

Clearly, humans are not born with total knowledge of the world. They spend a considerable portion of their lives learning, both by direct experience and through the transmitted experiences of others. Much of this experience must be highly individual. On a more abstract level, a person's family, culture, and linguistic community provide many of the conceptual building blocks with which he builds his model of the world. Of course, the primitives thus created could work against communication between persons with widely varying experiences, on an individual level, and members of dissimilar cultures, on a larger level.

We would like, then, some evidence that people from different cultures and with widely varying experiences do nevertheless share some common perception of the world. Fortunately, in recent years, convincing evidence has begun to be accumulated in the field of cognitive anthropology.

The experiments were originally designed to test the Whorfian hypothesis--broadly stated "language determines how you see reality." Different cultures divide the color spectrum with differing degrees of fineness--one culture having terms only for black and white, one having perhaps four color terms for black, white, the "warm" colors and the "cool" colors, still another subdividing the warm colors into red and yellow, another subdividing the cool colors into blue and green, others having, in addition, terms for purple or orange, or having names for other colors as well (in English we have terms for brown, pink, chartreuse, puce, magenta, lavender, scarlet, and so forth). Not only do cultures subdivide the spectrum with differing degrees of fineness, but the boundaries of the subdivisions differ considerably from culture to culture. A typical culture might have names for red, yellow, and "grue" (green/blue), where a sample of light yellow-green might be considered "yellow" and the darker greens would be judged as belonging to "grue". In general, a member of such a culture not only would have difficulty making a verbal distinction between green and blue, but he would also find it a bit strange that anyone would consider such a distinction important. At first it was thought that members of cultures which divided the spectrum differently actually perceived the colors differently--a case of a

culture's providing the primitives with which one builds one's concept of reality.

Berlin and Kay [1969] noted that, if a culture had a term for a color, then the "best" example of that color--the focal color for that part of the spectrum--tended to vary astonishingly little from culture to culture. Virtually all cultures have a term for "red", for example, and agreement on which of many reds is the "best" example of red is, from culture to culture, almost invariant.

Heider [1972], in working with about 200 samples of colors from all parts of the spectrum with various hues and values, found that certain colors were well remembered whether the subject (and the other members of his culture) had names for them or not--even if a subject had terms only for white and black, for example, there was a particular red (focal red) that he would have no trouble remembering, as well as a focal blue, a focal green, and a focal yellow--and that he could communicate these to his fellows significantly more easily than non-focal colors. There were found to be a total of eleven focal colors, counting white and black, and all of them are, by nature, easier to remember than non-focal colors. Furthermore, our language seems to have no direct effect on our ability to respond to them.

Moreover, the order in which color-terms are acquired by a culture appears to be the order in which young children develop the ability to differentiate the colors. Berlin and Kay [1969] and Berlin and Berlin [1975] outline seven stages: the first differentiates black and white, where "white" encompasses the light and warm (long wave-length) colors.* At the second stage the light/warm category splits to become two: one for which white is the focus, stressing lightness, and one for which red becomes the focus, including the warm chromatic colors. At stage III, either the light/warm category may further split between red and yellow, with their respective foci, or black may split from the dark/cool category, leaving a term which has its focus in black and another term for the green/blue part of the spectrum with focus at either focal green or focal blue. By stage IV, whichever split is not made in stage III has been accomplished. In the fifth stage a clear distinction is made between blue and green; in stage six, brown emerges from the light/warm category, and in stage seven, terms for purple, pink,

* Foci for the two color terms of a stage I culture tend not to be as predictable as for cultures at more advanced stages, but evidently focal red tends to be the focus for the light/warm category and black tends to dominate the dark/cool category.

and orange appear. Secondary color terms (chartreuse, lavender, crimson, and the like in English) can enter the vocabulary of a culture at a relatively early stage, but for early and late stages alike, there is wide variation in judgment of what a particular term refers to, and it is virtually impossible to derive a focal color for a given secondary term.

In any case, the order specified by the seven stages appears to be invariant. No culture, for instance, has clearly separated categories of blue and green, but does not have a color term with focus at yellow.

The response to the focal colors appears to be "hard-wired," although the basis for the response is not yet well understood. The cones of the retina respond maximally to certain ranges of frequencies of light. Likewise, the lateral geniculate nucleus has a higher frequency of response and a greater intensity of firing to certain frequencies in the red, green, yellow, and blue ranges, as well as black and white. but in neither case do the frequencies that are singled out correspond to the focal colors. In any case, the point is that in spite of culture and linguistic background, which seem to be working against any universal recognition of colors, in general, for non-colorblind individuals, the focal colors have a

special cognitive status which is not idiosyncratic.

Here then is a clear indication that there truly are universals in human conception of reality--an argument for primitives which are tied to something deeper than an individual's particular experiences.

Studies are beginning to be made in a search for other "hard-wired" perceptual mechanisms. To date there is evidence of at least one other probable universal construct--members of a fairly primitive culture which live in rounded huts and have no experience with right angles (other than, I suppose, tree trunks which tend toward right angles with respect to the ground) nevertheless appear to remember right angles more readily than other types of line junctions, and again the distinguished status appears to be due to "wired-in" sensory perception mechanisms.

The point of all this is that humans can communicate with one another because at basis they truly do experience the same reality. Granted, in some sense an Ancient Greek, or a Druid, or a Hopi experience a different universe than a European-based twentieth-century American. Culture does provide its own viewpoint, its own primitives, if you will, to help structure knowledge of the universe and interpret experience. But in a deeper sense, below the cultural level, we all do share common experiences.

DAT
ILMI