

Fragments of a Theory  
of Human Plausible Reasoning

Allan Collins  
Bolt Beranek and Newman Inc.

ABSTRACT

The paper outlines a computational theory of human plausible reasoning constructed from analysis of people's answers to everyday questions. Like logic, the theory is expressed in a content-independent formalism. Unlike logic, the theory specifies how different information in memory affects the certainty of the conclusions drawn. The theory consists of a dimensionalized space of different inference types and their certainty conditions, including a variety of meta-inference types where the inference depends on the person's knowledge about his own knowledge. The protocols from people's answers to questions are analyzed in terms of the different inference types. The paper also discusses how memory is structured in multiple ways to support the different inference types, and how the information found in memory determines which inference types are triggered.

INTRODUCTION

The goal of this paper is to briefly describe a theory of human plausible reasoning I am currently developing (Collins, 1978). The theory is a procedural theory and hence one which can be implemented in a computer, as parts of it have been in the SCHOLAR and MAP-SCHOLAR systems (Carbonell & Collins, 1973; Collins & Warnock, 1974; Collins, Warnock, Aiello & Miller, 1975). The theory is expressed in the production-rule formalism of Newell (1973). Unlike logic, the theory specifies how different configurations of information affect the certainty of the conclusions drawn. These certainty conditions are in fact the major contribution of the theory.

Methodology of Constructing the Theory

To construct a theory of human plausible reasoning, I collected about 60 answers to everyday questions from 4 different subjects. The questions ranged from whether there are black princess phones to when the respondent first drank beer.

The analysis of the protocols attempts to account for the reasoning and the conclusions drawn in the protocols in terms of 1) a taxonomy of plausible inference types, 2) a taxonomy of default assumptions, and 3) what the subject must have known a priori. As will be evident, this is an inferential analysis. I am trying to construct a deep structure theory from the surface structure traces of the reasoning process.

The protocols have the following characteristics.

- 1) There are usually several different inference types used to answer any question.
- 2) The same inference types recur in many different answers.
- 3) People weigh all the evidence they find that bears on a question.
- 4) People are more or less certain depending on the certainty of the information, the certainty of the inferences, and on whether different inferences lead to the same or opposite conclusions.

I can illustrate some of these characteristics of the protocols as well as several of the inference types in the theory with a protocol taken from a tutorial session on South American geography (Carbonell & Collins, 1973):

- (T) There is some jungle in here (points to Venezuela) but this breaks into a savanna around the Orinoco (points to the Llanos in Venezuela and Colombia).
- (S) Oh right, that is where they grow the coffee up there?
- (T) I don't think that the savanna is used for growing coffee. The trouble is the savanna has a rainy season and you can't count on rain in general. But I don't know. This area around Sao Paulo (in Brazil) is coffee region, and it is sort of getting into the savanna region there.

In the protocol the tutor went through the following reasoning on the question of whether coffee is grown in the Llanos. Initially, the tutor made a hedged "no"

response for two reasons. First, the tutor did not have stored that the Llanos was used for growing coffee. Second, the tutor knew that coffee growing depends on a number of factors (e.g., rainfall, temperature, soil, and terrain), and that savannas do not have the correct value for growing coffee on at least one of those factors (i.e., reliable rainfall). However, the tutor later hedged his initial negative response, because he found some positive evidence. In particular, he thought the Brazilian savanna might overlap the coffee growing region in Brazil around Sao Paulo and that the Brazilian savanna might produce coffee. Thus by analogy the Llanos might also produce coffee. Hence, the tutor ended up saying "I don't know."

The answer exhibits a number of the important aspects of the protocols. In general, a number of inferences are used to derive an answer. Some of these are inference chains where the premise of one inference depends on the conclusion of another inference. In other cases the inferences are independent sources of evidence. When there are different sources of evidence, the subject weighs them together to determine his conclusion.

It is also apparent in this protocol how different pieces of information are found over time. What appears to happen is that the subject launches a search for relevant information (Collins & Loftus, 1975). As relevant pieces of information are found (or are found to be missing), they trigger particular inferences. The type of inference applied is determined by the relation between the information found and the question asked. For example, if the subject knew that savannas are in general good for growing coffee, that would trigger a deduction. If the subject knew of one savanna somewhere that produced coffee, that would trigger an analogy. The search for information is such that the most relevant information is found first. In the protocol, the more relevant information about the unreliable rainfall in savannas was found before the more far fetched information about the coffee growing region in Brazil and its relation to the Brazilian savanna. Thus, information seems to be found at different times by an autonomous search process, and the particular information found determines inferences that are triggered.

#### THE THEORY

The theory specifies a large number of different inference types, together with the conditions that affect the certainty of each inference type. In the theory the different types of inference are arrayed in a five dimensional space.

The dimensions of the inference space are:

#### (1) Inferences on Knowledge vs Inferences on Meta-Knowledge

There are inference patterns based on people's knowledge, such as deduction and induction, and inference patterns based on people's knowledge about their own or other's knowledge (i.e. meta-knowledge) (Brown, 1977), such as lack-of-knowledge and confusability inferences. I refer to these latter as meta-inferences. They are ubiquitous in the protocols, and yet they fall outside the scope of most theories of logic. The other four dimensions refer to the space of inferences but may also partially apply to the space of meta-inferences.

#### (2) Functional vs Set Inferences

For each type of inference, there is a functional variation and a set variation. The set variation involves mapping the property of one set (which may be a single-member set or instance) onto another set. The functional variation has an additional premise that the property to be mapped (the dependent variable) depends on other properties (the independent variables). The mapping of the property from one set to another makes use of this functional dependency. The set variation, in fact, is a degenerate form of the functional variation, which is used when people have little or no knowledge of the functional dependencies involved.

People's knowledge about functional dependencies consists of a kind of directional correlation. A judgment about whether a place can grow coffee might depend on factors that are causal precursors for coffee growing (e.g., temperature), correlated factors (e.g., other types of vegetation), or factors causally subsequent to coffee growing (e.g., export trade). For example, one might decide a place does not produce coffee, because it produces apples which seem incompatible with coffee, or because there is little export trade from the region. The directional nature of the correlation shows up in the last example. A region easily could have export trade without producing coffee, but it would be unlikely that a region would produce coffee without having export trade.

#### (3) Semantic, Spatial, vs Temporal Inferences

For each type of inference, there is a semantic, spatial, or temporal variation of the inference. Semantic inferences involve mapping properties across semantic space, spatial inferences across Euclidean space, and temporal inferences across time. These are treated as different types of inferences in the theory because the procedures for computing them are somewhat different. Semantic inferences are based on information structured in a semantic or conceptual memory (Quillian, 1968; Schank, 1972). Spatial inferences are based on information (or images) derived from a spatial structure (Collins & Warnock, 1975; Kosslyn & Schwartz, 1977). Temporal inferences are based on information derived from an event (or

episodic) structure (Tulving, 1972). Correlates of each of these types of memory structures are found in Winograd's SHRDLU (1972).

(4) Superordinate sets, similar sets, vs. subordinate sets

Inferences can involve mapping properties from superordinate sets, similar sets, or subordinate sets. The property can be mapped from one set or from many sets (either exhaustively or not). The different kinds of mappings delineated in the theory are:

- (a) Deduction (Superordinate Inferences) maps properties of the set onto subsets.
- (b) Analogy (Similarity Inferences) maps properties from one set to a similar set.
- (c) Induction maps properties of subsets of a set onto other subsets.
- (d) Generalization (proof-by-cases) maps properties of subsets of a set onto the set.
- (e) Abduction maps a subset with the same property as some set into the set.

(5) Positive vs. Negative Inferences

Each type of inference has both a positive and negative version, depending on whether the mapping involves the presence or absence of a property

Assumptions of the Theory

The theory rests on a number of assumptions about the way information is represented and processed by people. I will describe briefly what these assumptions are.

Semantic Information I assume information about different concepts is represented in a cross-referenced, semantic structure (Quillian, 1968; Schank, 1972). The nodes in the network are schemas, which are the kind of structured objects implied by the notion of frames (Minsky, 1975) or scripts (Schank & Abelson, 1977). The links between nodes represent different relations between the concepts. The correlate of this kind of semantic structure in Winograd's SHRDLU (1972) was the cross-referenced information structure constructed by MICROPLANNER.

Spatial Information. I assume spatial information about concepts, such as the size, shape, color, or location of objects and places, is represented in a spatial structure, apart from but connected to the semantic structure (Collins & Warnock, 1974). The correlate of such a spatial representation in Winograd's SHRDLU (1972) was the Cartesian representation of the blocks on the table top.

Event information. Similarly event information is assumed to be stored in a form that preserves its temporal, causal, and goal structure. This requires a hierarchical structure of events and subevents nested according to the goals and subgoals of the motors involved in the events (Brown, Collins,

& Harris, 1978). Such an event memory was constructed by Winograd's SHRDLU (1972) to record the movements of blocks and the goals they accomplished, in order to answer "why" and "how" questions about events in the Blocks World.

Retrieval I assume there are autonomous search processes that find relevant information with respect to any query (Collins & Loftus, 1975). The search process has access to semantic, spatial and temporal information in parallel, and whenever relevant information of any kind is found, it triggers an inference (Collins & Quillian, 1972, Kosslyn, Murphy, Bemesderfer & Feinstein, 1977.) The information found by the search processes determines what inference patterns are applied.

Matching Processes. I assume there are decision processes for determining whether any two concepts can be identified as the same. The semantic matching process could be that proposed by Collins & Loftus (1975) or by Smith, Shoben & Rips (1974). The spatial matching process compares places or objects to decide their spatial relation. Similarly, there must be a temporal matching process that determines the relation between two events.

Importance and Certainty. I assume that for each concept and relation a person has a notion of its relative importance (i.e. its criteriality), and his degree of certainty about its truth. In a computer, these could be stored as tags on the concepts and relations (Carbonell & Collins, 1973).

EXAMPLES OF INFERENCE RULES AND PROTOCOLS

Because it is impossible to present the entire theory here, I will give the formulations for three types of inference and show three protocols which illustrate these three types, as well as others. The three types are the lack-of-knowledge inference, the functional analogy, and the spatial superpart inference. They are all common inferences and serve to illustrate the different kinds of inferences in the theory.

The formal analysis of the protocols attempts to specify all the underlying inferences that the subject was using in his response. For the inferences that bear directly on the question, I have marked whether they are evidence for a negative or positive answer. Where a premise was not directly stored, but derived from another inference, I have indicated the inference from which it is derived. I have indicated the approximate degree of certainty by marking the conclusion with "Maybe", "Probably", or leaving it unmarked. Where a subject may be making a particular inference which the protocol does not clearly indicate, I have marked the inference "possible". Separating inferences in this manner is oversimplified, but has the virtue of being understandable.

Lack-of-Knowledge Inference

The lack-of-knowledge inference is the most common of all the meta-inferences. The protocol I selected to show the lack-of-knowledge inference shows the subject using a variety of meta-inferences to reach an initial conclusion which he then backs off a bit.

Q. Is the Nile longer than the Mekong River?

JB. I think so.

Q Why?

JB Because (pause) in junior high I read a book on rivers and I kept looking for the Hudson River because that was the river I knew about and it never appeared, and the Amazon was in there and the Nile was in there and all these rivers were in there, and they were big, and long, and important. The Mekong wasn't in there (pause) It could be just

Q. So therefore, it is not important.

JB. That's right It could be just an American view. At that time the Mekong wasn't so important

Underlying Inferences

1) Functional Abduction on Importance Level (Possible)

The importance of a river depends in part on how long it is  
The Nile is very important  
Probably the Nile is extremely long

2) Meta-Induction From Cases

I know the Amazon is extremely long  
I know the Nile is extremely long (from 1)  
I would know the Mekong is extremely long if it were

3) Lack-of-Knowledge Inference

I don't know the Mekong is extremely long  
I would know the Mekong is extremely long if it were (from 2)  
Probably the Mekong is not extremely long

4) Functional Abduction on Importance Level (Possible)

The importance of a river depends in part on length  
The Mekong is not very important  
Probably the Mekong is not extremely long

5) Simple Comparison (Positive Evidence)

The Mekong is not extremely long (from 3 and 4)  
The Nile is extremely long (from 1)  
The Nile is longer than the Mekong

6) Functional Attribution on Importance Level (Possible)

The importance of something depends on how remote it is  
The Nile is very important  
The Nile is less remote than the Mekong  
Maybe the Nile is more important than the Mekong because it's less remote

7) Functional Alternative on Importance Level (Negative Evidence) (Possible)

The importance of a river depends on how close it is and how long it is  
The Nile is more important than the Mekong because it's closer (from 6)  
Maybe the Nile is not longer than the Mekong

Contributing to the certainty of these inferences are several meta-inferences working on importance level. The functional abductions (1 and 4) are suggested by the subject's tying length to importance. He seems to know that importance depends in part on length, and since he assigns different degrees of importance to the Nile and the Mekong, he must be using that in part to infer that the Mekong is not as long as the Nile. There also is a meta-induction he is making: that since he knows the Amazon and the Nile are very long, he would know the Mekong is long if it were. This meta-induction is acting on one of the certainty conditions for the lack-of-knowledge inference: the more similar cases stored with the given property, the more certain the inference. Taken together, these inferences make the lack-of-knowledge inference very certain.

However at the end the subject backs off his conclusion because he finds another chain of reasoning that makes him less certain (inferences 6 and 7). The idea of "remoteness" only represents the underlying argument when interpreted in terms of conceptual distance. What the subject is really doing is evaluating how remote Southeast Asia was at the time he was in junior high (before the Vietnam War). This notion of remoteness is the outcome of matching processes. The Mekong was remote because it was far away culturally, historically, physically, etc. from America. Based on this the subject realizes that the Mekong's lack of importance may be due to this remoteness rather than its shortness in length. His reasoning then depends on his notion of what alternative factors importance depends on, and how it might mislead him in this case. So this chain of reasoning is also acting on the certainty conditions affecting the lack-of-knowledge inference, but in the opposite direction from the other meta-inferences.

The rule for a lack-of-knowledge inference is shown in the table below. It generally has the form. If it were true, I would know about it; I don't, so it must not be true. It is computed by comparing the importance level of the proposition in question against the depth of knowledge about the concepts involved (Collins et al, 1975; Gentner & Collins, 1978).



Lack-of-Knowledge Inference

- 1) If a person would know about a property for a given set if it were in a given range, and
- 2) if the person does not know about that property,
- 3) then infer that the property is not in the given range for that set.

Example

If Kissinger were 6'6" tall, I would know he is very tall. I don't, so he must not be that tall.

Conditions that increase certainty:

- 1) The more important the particular set.
- 2) The less likely the property is in the given range.
- 3) The more information stored about the given set.
- 4) The more similar properties stored about the given set.
- 5) The more important the given property.
- 6) The more information stored about the given property.
- 7) The more similar sets stored that have the given property.

The conditions affecting the certainty of a lack-of-knowledge inference can be illustrated by the example in the table:

- 1) Condition 1 refers to the importance of the given set. In the example Kissinger is quite important, so one is more likely to know whether he is 6'6" than whether Senator John Stennis is 6'6" for example.
- 2) Condition 2 refers to the likelihood that the property is in the given range. Likelihood affects the inference in two ways: low likelihood makes a negative inference more certain a priori, and low likelihood also makes a property more unusual and therefore more likely to come to a person's attention. For example, it is less likely that Kissinger is 7' 2" than 6' 6", because 7' 2" is more unusual. If Kissinger were a basketball player, on the other hand, his being 6' 6" would not be unusual at all.
- 3) Condition 3 relates to the depth-of-knowledge about the given set. The more one knows about Kissinger, the more certainly one would know that he is 6' 6", if he is.
- 4) Condition 4 relates to the number of similar properties stored about the set (i.e. the relatedness of the information known about the set). If one knows a lot about Kissinger's physical appearance, one feels more certain one would know he is extremely tall, if he is.
- 5) Condition 5 relates to the importance of the particular property. Being extremely tall isn't as important as missing a leg say, so people are more likely to know if Kissinger is missing a leg.

- 6) Condition 6 relates to the depth-of-knowledge about the particular property. For example, a person who has particular expertise about the physical stature of people is more likely to know that Kissinger is extremely tall, if he is.
- 7) Condition 7 relates to the number of similar sets known to have the given property. For example, if one knows that Ed Muskie and Tip O'Neil are unusually tall, then one ought to know that Kissinger is unusually tall, if in fact he is 6' 6".

Functional Analogy

The initial protocol on coffee growing in the Llanos illustrated two functional inferences: a functional calculation concerning rainfall, and a functional analogy between the Brazilian savanna and the Llanos. One of the more common functional inferences is the functional analogy. The protocol I selected to illustrate it contrasts the use of a simple analogy and a functional analogy.

Q. Can a goose quack?

BF. No, a goose - Well, its like a duck, but its not a duck. It can honk, but to say it can quack. No, I think its vocal cords are built differently. They have a beak and everything, but no, it can't quack.

Underlying Inferences

- 1) Simple Analogy (Positive Evidence)  
A goose is similar to a duck  
A duck quacks  
Maybe a goose quacks
- 2) Importance-Level Inequality (Possible)  
I know a goose honks  
Quacking is as important as honking  
Probably I would know about a goose quacking if it did
- 3) Lack-of-Knowledge Inference (Negative Evidence) (Possible)  
I don't know that a goose quacks  
I would know about a goose quacking if it did (from 2)  
Probably a goose doesn't quack
- 4) Negative Functional Analogy (Negative Evidence)  
The sound a bird makes depends on its vocal cords  
A goose' is different from a duck in its vocal cords  
A duck quacks  
Probably a goose doesn't quack

The simple analogy, which is based on a match of all the properties of ducks and geese, leads to the possible conclusion that a goose can quack, because a duck quacks. This inference shows up in the reference to "its like a duck" and in the uncertainty of the negative conclusion the student is drawing. It is positive evidence and only shows up to the degree it argues against the general negative conclusion.

The importance-level inequality and lack-of-knowledge inference are suggested by the sentence "It can honk, but to say it can quack." Here knowledge about honking seems to imply that a goose doesn't quack. I would argue that such an inference has to involve the lack-of-knowledge inference, since it is possible that a goose might sometimes honk and sometimes quack.

The functional analogy is apparent in the concern about vocal cords, which the subject thinks are the functional determinants of the sounds made. I think the sound is determined by the length of the neck, which is probably what the subject was thinking of. Honking may just be quacking resonated through a longer tube. But in any case, the mismatch the subject finds on the relevant factor leads to a negative conclusion which supports the lack-of-knowledge inference.

The table shows the rule for a functional analogy.

Functional Analogy

- 1) If a dependent variable depends on a number of independent variables, and
- 2) if one set matches another set on the independent variables, and
- 3) if the value of the dependent variable for one set is in a given range,
- 4) then infer that the value of the dependent variable for the other set is in the given range.

Example

The Brazilian savanna is like Llanos in its temperature, rainfall, soil, and vegetation. Thus, if the Brazilian savanna produces coffee, then the Llanos ought to also.

Conditions that increase certainty:

- 1) The more independent variables on which the two sets match, and the fewer on which they mismatch.
- 2) The greater the dependency on any independent variables on which the two sets match, and the less the dependency on any independent variables that mismatch.
- 3) The better the match on any independent variable.
- 4) The greater the dependency on those independent variables that match best.
- 5) The more certain the dependent variable is in the given range for the one set.
- 6) The more likely the value of the dependent variable is in the given range a priori.
- 7) The more certain the independent variables are in the given ranges for both sets.

I can illustrate the different certainty conditions for a functional analogy in terms of the example in the table.

- 1) Condition 1 refers to the number of factors on which the two sets match. If the two regions match only in climate and vegetation, that would be less strong evidence that they produce the same products than if they match on all four variables.

- 2) Condition 2 refers to the degree the dependent variable depends on different factors that match or mismatch. Coffee growing depends more on temperature and rainfall than on soil or vegetation. Thus a match on these first two factors makes the inference more certain than a match on the latter two factors.
- 3) Condition 3 relates to the quality of the match on any factor. The better the match with respect to temperature, rainfall, etc. the more certain the inference.
- 4) Condition 4 refers to the degree of dependency on those factors that match best. A good match with respect to the rainfall pattern leads to more certainty than a good match with respect to the vegetation.
- 5) Condition 5 relates to the certainty that the property is in the given range for the first set. The more certain one is that the Brazilian savanna produces coffee, the more certain the inference.
- 6) Condition 6 relates to the a priori likelihood that the property will be in the given range. The more likely that any region grows coffee, the more certain the inference.
- 7) Condition 7 relates to the certainty that the factors are in the given ranges for both sets. For example, the more certain that both savannas have the same temperature, etc., the more certain the inference.

Spatial Superpart Inference

The theory assumes that spatial inferences are made by constructing an image of the concepts involved, and making various computations on that image (Collins & Warnock, 1974; Kosslyn & Schwartz, 1977). An example of a spatial inference occurred in the earlier protocol about coffee growing, when the respondent concluded that a savanna might be used for growing coffee because he thought the coffee growing region around Sao Paulo might overlap the Brazilian savanna. This spatial matching process, which occurs in a variety of protocols, involves constructing a spatial image with both concepts in it, and finding their spatial relationship (e.g., degree of overlap, relative size or direction) from the constructed image.

The protocol I selected illustrates a spatial subpart inference, together with several other spatial and meta-inferences.

Q. Is Texas east of Seattle?

JB. Texas is south and east of Seattle.

Q. How did you get that?

JB. I essentially looked at a visual image of the U.S. where I remembered that Seattle was in Washington and know that its up in the left corner and I know that Texas is in the middle on the bottom. Sometimes you get fooled by things like that, like for example Las Vegas being further west than San Diego. This case I think we're O.K.

Underlying inferences

- 1) Spatial line slope inference  
Washington is in upper left corner of the U.S.  
Texas is on the middle bottom of U.S.  
Line from Washington to Texas slopes east.
- 2) Spatial subpart inference (Positive evidence)  
Line from Washington to Texas slopes east.  
Seattle is part of Washington.  
Line from Seattle to Texas slopes east
- 3) Meta Analogy (Negative evidence)  
People are often mistaken in thinking that Las Vegas is east of San Diego, because Las Vegas is inland and San Diego is on the Pacific Coast.  
Seattle, like San Diego, is on the Pacific coast.  
Texas, like Las Vegas, is inland.  
Maybe I am mistaken in thinking that Texas is east of Seattle.
- 4) Functional Modus Tollens (Positive evidence) (possible)  
The Pacific coast misconception depends on the inland place being north of the coastal place.  
Seattle is on the coast.  
Texas is inland.  
Texas is south of Seattle.  
The Pacific coast misconception does not apply to Texas and Seattle.

In the protocol the subject constructs a line from Washington to Texas for the purpose of evaluating its slope. The constructed line does slope east, so he answers yes. Implicit in this protocol is a spatial subpart inference or spatial deduction, that Seattle is part of Washington and the slope of the line found earlier applies to Seattle. This kind of subpart inference was found to show up in response time by Stevens (1976).

The subject briefly reconsidered his conclusion because he thought of the "Pacific Coast Misconception," that people mistakenly think that places inland are always east of places on the coast. By the meta-analogy in 3, he inferred that maybe Seattle-Texas was like San Diego-Las Vegas in that the inland location was west of the coastal location. But the subject ruled out the analogy by some inference such as that shown in 4. Actually, the functional modus tollens in 4 hides the spatial processing that the subject probably used to rule out the analogy in 3. Probably, he knew that the reason for the "Pacific Coast Misconception" has to do with the southeasterly slant of the Pacific coast. By knowing that, you can figure out that the misconception depends on the inland location being north of the coastal location. I have finessed the spatial reasoning process by stating that conclusion as a premise in 4.

The next table shows the rule for a spatial superpart inference (or spatial deduction).

Spatial Superpart Inference

- 1) If a property is in a given range for some set, and
- 2) if another set is a subpart of that set,
- 3) then infer that the property is in that range for the subpart.

Example

It is raining in New England and Boston is in New England. Therefore it may be raining in Boston.

Conditions that increase certainty:

- 1) The more central the subpart is to the set.
- 2) The greater the average spatial extent of the property.
- 3) The greater the distance of the nearest set with a contradictory property.
- 4) The greater the extent of the subpart within the set.
- 5) The more likely a priori that the property is in the given range for the subpart.
- 6) The more certain the property is in the given range for the set.

The certainty conditions can be illustrated in terms of the example in the table:

- 1) Condition 1 relates to the centrality of the subpart. For example, if it's raining in New England it is more likely to be raining in Massachusetts than Maine because Massachusetts is more central.
- 2) Condition 2 relates to whether the property tends to be spatially distributed or not. For example, rain tends to be distributed over smaller areas than electric service, so it is a less certain inference that it is raining in Maine than that there is electric service in Maine, given that the property applies to New England.
- 3) Condition 3 relates to the distance to the nearest concept with a contradictory property. For example, if you know it's not raining in New Brunswick, that is stronger evidence against it's raining in Maine than if it's not raining in Montreal.
- 4) Condition 4 relates to the extent of the subpart. For example, if it's raining in New England it is more likely to be raining in Rhode Island than in Boston, because Rhode Island is larger.
- 5) Condition 5 relates to the a priori likelihood of the property. For example, if it's raining in Washington State, it's more likely to be raining in Seattle than in Spokane because Seattle gets more rain on the average.
- 6) Condition 6 relates to the person's certainty that the property holds for the concept. For example, the more certain the person is that it is raining in New England, the more certain that it's raining in Boston.

CONCLUSION

The theory I am developing is based on these and similar analyses of a large number of human protocols. Because the same inference types recur in many different answers, it is possible to abstract the systematic patterns in the inferences

themselves, and many of the different conditions that affect people's certainty in using different inference types:

#### ACKNOWLEDGEMENTS

I want to thank my colleagues who have influenced my views about inference over the years namely Marilyn Adams, Nelleke Aiello, John Seely Brown, Jaime Carbonell, Dedre Gentner, Mark Miller, Ross Quillian, Albert Stevens, and Eleanor Warnock. I particularly would like to thank Marilyn Adams for encouraging me to fit the inference types into a dimensionalized space, and John Seely Brown for bullying me into stating the rules and protocol analyses in a form understandable to readers.

This research was supported in part by the Advanced Research Projects Agency of the Department of Defense under Contract No. MDA 903-77-C-0025, and in part by a fellowship from the John Simon Guggenheim Memorial Foundation.

#### REFERENCES

- Brown, A. L. Knowing when, where & how to remember. In R. Glaser (Ed.), Advances in instructional psychology. Hillsdale, NJ: Lawrence Erlbaum Associates, 1977, in press.
- Brown, J.S., Collins, A., & Harris, G. Artificial intelligence and learning strategies. To appear in H.F. O'Neil (Ed.), Learning strategies. New York: Academic Press, 1978, in press.
- Carbonell, J.R. & Collins, A. Natural Semantics in Artificial Intelligence. Proceedings of Third International Joint Conference on Artificial Intelligence, 1973, pp. 344-351. (Reprinted in the American Journal of Computational Linguistics, 1974, 1, Mfc. 3).
- Collins, A. & Warnock, E.H. Semantic networks. BBN Report No. 3833, Bolt Beranek and Newman Inc., Cambridge, Mass., 1974.
- Collins, A. M. & Loftus, E. F. A spreading activation theory of semantic processing. Psychological Review, 1975, 82, 407-428.
- Collins, A., Warnock, E.H., Aiello, N. & Miller, M.L. Reasoning from Incomplete Knowledge, in D. Bobrow & A. Collins (eds.). Representation & understanding. New York: Academic Press, 1975.
- Collins, A.M., & Quillian, M.R. Experiments on semantic memory and language comprehension. In L.W. Gregg (Ed.), Cognition in learning and memory. New York: Wiley, 1972.
- Collins, A.M., Adams, M.J. & Pew, R.W. The Effectiveness of an interactive map display in tutoring geography. Journal of Educational Psychology, 1978, 70, 1-7.
- Gentner, D., & Collins, A. Knowing about knowing: Effects of meta-knowledge on inference. Submitted to Cognitive Psychology.
- Kosslyn, S.M., & Schwartz, S.P. A simulation of visual imagery. Cognitive Science, 1977, 1, 265-295.
- Kosslyn, S.M., Murphy, G.L., Bemserfer, M.E., & Feinstein, K.J. Category and continuum in mental comparisons. Journal of Experimental Psychology: General, 1977, 106, 341-375.
- Minsky, M. A framework for representing knowledge. In P. H. Winston (Ed.), The psychology of computer vision. New York: McGraw-Hill, 1975.
- Quillian, M. R. Semantic memory. In M. Minsky (Ed.), Semantic information processing. Cambridge, Mass.: MIT Press, 1968.
- Schank, R. Conceptual Dependency: A Theory of Natural Language Understanding, Cognitive Psychology, 1972, 3, 552-631.
- Schank, R. & Abelson, R. Scripts, plans, goals, and understanding. Hillsdale, N.J.: Lawrence Erlbaum Associates, 1977.
- Smith, E.E., Shoben, E.J., & Rips, L.J. Comparison processes in semantic memory. Psychological Review, 1974, 81, 214-241.
- Stevens, A.L. The role of inference and internal structure in the representation of spatial information. Doctoral dissertation. University of California at San Diego, 1976.
- Tulving, E. Episodic & semantic memory. In E. Tulving & W. Donaldson (Eds.), Organization & memory. New York: Academic Press, 1972.
- Winograd, T. Understanding natural language. New York: Academic Press, 1972.