DOCUMENT RESUME

ED 244 253 CS 007 653

AUTHOR Bruce, Bertram
TITLE Robot Plans and Human Plans: Implications for Models of Communication. Technical Report No. 314.
INSTITUTION Bolt, Beranek and Newman, Inc., Cambridge, Mass.; Illinois Univ., Urbana. Center for the Study of Reading.
SPONS AGENCY National Inst. of Education (ED), Washington, DC.
PUB DATE Jun 84
CONTRACT 400-81-0030
NOTE 42p.
PUB TYPE Reports - Research/Technical (143) -- Information Analyses (070)

EDRS PRICE MF01/PC02 Plus Postage.
DESCRIPTORS *Artificial Intelligence; Behavior Patterns; *Communication (Thought Transfer); *Discourse Analysis; *Interaction; *Learning Theories; *Models; Planning; Reading Research; Robotics; Social Behavior

ABSTRACT

People in interaction with others organize their perceptions of a social situation in terms of plans even when the others' plans are poorly formulated. They use their models of others' plans in formulating their own. Much of what occurs in discourse centers on a continual communication about and reformulation of one's own plans and one's own models of others' plans. Formal methods for describing and analyzing plans are now available and make it possible to be more explicit in formulating hypotheses about social interaction. Unfortunately, the classical formulations for planning are derived from a robot world model that fails to generalize sufficiently to account for typical human planning situations. By pushing the classical model, it is possible to create a model of planning that embodies concepts such as "mutual belief," "social episode," and "goal conflict." In applying this model to a simple dialogue, it becomes clear that elements of the model are necessary for modeling such dialogues, but not sufficient. More complex dialogues and texts will undoubtedly require further elaboration of the model. (Author/FL)

CENTER FOR THE STUDY OF READING

Technical Report No. 314

ROBOT PLANS AND HUMAN PLANS: IMPLICATIONS FOR MODELS OF COMMUNICATION

Bertram Bruce

Bolt Beranek and Newman Inc.

June 1984

University of Illinois
at Urbana-Champaign
51 Gerty Drive
Champaign, Illinois 61820

Bolt Beranek and Newman Inc.
10 Moulton Street
Cambridge, Massachusetts 02238

The research reported herein was supported in part by the National Institute of Education under Contract No. NIE-400-81-0030. I would like to thank Andee Rubin, Behrooz Tavakoli, and Harry Blanchard for comments on this paper; also, Kathleen Starr for her artwork and Charyl Przekwas for preparing the manuscript.

EDITORIAL BOARD

Abstract

This paper examines the role of plan recognition and plan generation in communication. It argues that people in interaction with others organize their perceptions of a social situation in terms of plans even when the others' plans are poorly formulated. They use their models of others' plans in formulating their own. Much of what occurs in discourse centers on a continual communication about and reformulation of one's own plans and one's models of others' plans. Formal methods for describing and analyzing plans are now available. Such methods allow us to be more explicit in our hypotheses about social interaction. Unfortunately, the classical formalisms for planning are derived from a robot world model that fails to generalize sufficiently to account for typical human planning situations. By pushing the classical model we come to a model of planning that embodies concepts such as mutual belief, social episode, and goal conflict. Using the mutual belief model we look at a simple dialogue. It is clear from the analysis that elements of the model are necessary for modeling such dialogues, but not sufficient. More complex dialogues and texts will undoubtedly require further elaboration of the model.

Robot Plans and Human Plans:
Implications for Models of Communication

The Peanuts comic strip has a character named Lucy who is always trying to take advantage of one named Charlie Brown. Because Charlie is very trusting, her tricks often succeed. One of her favorite tricks occurs in the fall when football season is getting underway: Lucy holds a football and says, "Come on Charlie Brown! I'll hold the football for you. You come and kick it." Charlie runs as fast as he can, swings his leg back and tries to kick the ball. At the last second Lucy yanks the ball away, Charlie's feet fly up in the air and he lands on his back. His pride is wounded as well as his bottom.

One year, Charlie started thinking,

> "She says she's going to hold the ball so I can run up and kick it, but I know she's going to pull it away. She always pulls it away. She thinks that I don't know what she's planning to do, but, in fact, I do. I also know that she thinks that I don't know. I'm going to trick her. Instead of running up and kicking the ball I'll just run up and stop. Then she can't pull the ball away and make me fall on my back."

But then he goes further,

> "She can figure out that I know what her plan is. She's probably not going to yank the ball away after all. I'll

run up, stop, and be embarrassed, this time because there'll be no trick."

Following this line of reasoning he finally concludes that the best strategy is just to run as hard as he can and kick the ball. He runs right up and swings his leg back. Lucy jerks the ball away; he tries to kick it and lands flat on his back; and, of course, sets himself up for next year's football season.

We live in a world of intentionality, a world in which we assign meaning to objects, relations, and events. One of the primary kinds of meaning that we assign is planful behavior. That is, we look at what people do and say, "What are their goals? What are the actions that they are going to carry out to achieve those goals? What are their plans?" We seek interpretations of actions in terms of plans even in cases where behavior may not be all that planful. Referring to an action as aimless highlights the fact that it is very unusual for people to act without some goal in mind.

This paper discusses the role of plans in understanding discourse. It looks briefly at an example of what might be called a "robot model," or perhaps the "standard AI model" for plans, a model used widely in Artificial Intelligence, as well as in a wide variety of other disciplines. It then presents some problems that arise when this standard model is used to account for general human action. A contrasting model, the "mutual-belief model," is then presented by means of an example from a corpus (Hall, Linn, & Nagy, 1984) of children's conversations.

In a simple episode, we see something of the richness of human planfulness. Finally, some open questions are presented.

The Standard AI Model for Planning

Plans have been used in a wide variety of ways in the study of discourse: to look at interactions of characters in stories (Bruce & Newman, 1978; Wilensky, 1981), to look at the interactions of the plans and purposes of an author with the plans and purposes of a reader (Brewer & Lichtenstein, 1982; Bruce, 1981), to look at the various plans of participants in conversation. Most of these analyses have relied on some form of information processing model that provides an explicit representation of actions, states of the world, goals and the process of planning. There is a long tradition to this work, going back to Miller, Galanter, and Pribram (1960). The example to follow is a bit of a caricature of this view of planning and problem solving, but it suggests how people have been using terms like "operators" and "states" and the kinds of problems that can be solved.

Robot Worlds

First of all, one needs to have some way to represent a state of the world. Typically what's done is to define a world state by a list of propositions. Figure 1 shows a simple world in which there are five points, or places where something can be; there are boxes A, B and C; and there is a robot, who can move

the boxes around. To represent a state of this world we might use a set of propositions like the following:

Initial State

[On (Box-A, Floor)]
[On (Box-B, Floor)]
[On (Box-C, Box-B)]
[At (Box-A, P1)]
[At (Box-B, P2)]
[At (Robot, P3)]

This says that Box-A is on the Floor at point P1, the Robot is at point P3, and so on.

Insert Figure 1 about here.

Planning becomes relevant when we have operators which can change states of the world. An example of an operation in the micro-world might be for the robot to pick up a box. Operators are typically defined in terms of enabling conditions, and outcomes, or effects:

Operator:	pick-up (X)
Enabling conditions:	(εP) [At (Robot, P) & At (X, P)]
	(εZ) Hold (Z)
	(εZ) On (Z, X)
Effects:	Hold (X)

The pick-up operator can be used only in certain states, for example, the robot and the thing to be picked up have to be in the same location. Also, the robot cannot be holding anything and there can be nothing on the thing to be picked up. If the pick-up operator is applied, the effect is that a proposition is

added to our description of the world: The robot now holds the box.

Now, let's set up a goal state (see Figure 2), say, to have Box C on Box A and the robot at P1. This looks like the initial state with the exception of a couple of propositions:

Goal State

[On (Box-A, Floor)]
[On (Box-B, Floor)]
[On (Box-C, Box-A)]
[At (Box-A, P1)]
[At (Box-B, P2)]
[At (Robot, P1)]

Insert Figure 2 about here.

A planner can now operate to generate a sequence of operators, such as: "the Robot moves to P2, picks up Box-C, moves to P1, puts down Box-C." We call this sequence of operators a "plan."

The Plan

Move (P2)
Pick-up (Box-C)
Move (P1)
Put-down (Box-C)

More elaborate examples than this can easily be constructed, such as planning programs to write computer software or to solve complex assemly tasks. In various fields, people have used similar models to examine problem solving, composition, the reader interaction with characters in stories, conversations and

so on. It's fair to say that most planning research has relied on models of this general form.

Limitations of the Standard AI Model

Unfortunately, there are limitations to the standard AI model that appear when we try to account for human planning behavior. (Most of these are acknowledged by those who have used such a model in their research.)

The first problem is that robot planning is basically planning for an individual. There is one robot moving around in its micro-world. In contrast, human plans are essentially social, that is, people operate in a social environment in which they compete or cooperate with others. Even if there is another robot in the robot's world, the plan constructed is an individual plan; the other robot is an aspect of a passive environment. For human plans, many apparently individual actions--a teacher at home preparing a lesson for the next day--are still inherently social (the teacher must imagine social contexts for his/her actions).

Secondly, the robot operates in a world of facts. There is a set of propositions each of which are simply true or false. There may be some uncertainty about the propositions, and they may change but they're essentially treated as solid facts. In contrast, in human planning situations virtually everything has a belief status. People act on the basis of what they believe to be true, and what they believe about others' beliefs. In conversations, for example, the operators are defined in terms of

changes to others' beliefs or questions about what they believe. Moreover, these beliefs are recursive: It's not just your beliefs, and your beliefs about someone else's beliefs, but your beliefs about what they believe about you. For example, Charlie Brown is trying to figure out his beliefs about Lucy's beliefs about his beliefs.

Third, the standard AI model presupposes a fixed set of operators, state variables, and goals. Neither the operators nor the state variables change in the process of planning. There is no way to add new locations or new boxes. Also, the goal is fixed. The robot works to achieve the goal and either succeeds or fails. In the human planning situation, there is an on-going, and even a retrospective establishment of meaning. For example, Gearhart and Newman (1980) report nursery school interactions, in which a teacher talks to kids in the nursery school about the pictures they have drawn. She says, "Here, let me see your picture." Until that moment the child might have been drawing with the crayon on the paper, on the desk or on the wall. By her statement the teacher seeks to establish that the object in question is a picture and that the activity that the child was doing was drawing a picture. When she points to an orange glob on the paper and says, "Oh, that must be the sun," that orange circle suddenly takes on a new status. This establishment of meaning through interaction is a typical feature of human interactions; we define the rules of the game as we go along.

Fourth, the standard AI model tends to focus on the design and execution of plans. Given an initial state, a goal state and operators, the task is to construct the most efficient plan to get from the initial state to the goal state, or, given a plan, to execute it. In contrast, in human planning situations, explicit planning is unusual, and there are few times when one just carries out a plan. More often, people talk about plans; they communicate plans to one another; and they use plans to achieve other goals. Plans become part of what one comunicates: A plan is announced because of the effect that information is likely to have on someone else. These ways of dealing with plans require more sophisticated representations than those provided by the standard models.

Fifth, the standard AI model tends to separate the generation process from the recognition process. Producing a plan is one problem: How can we take an initial state and a sequence of operators and get to a goal state? Recognizing a plan is a separate problem: How can we infer a plan from a sequence of actions? For human plan models, these processes need to be considered together. Generating a plan is done while recognizing what plans others are carrying out. One's own plan may have to change in light of the recognition of what others are planning. Similarly, if others recognize your plan in a certain way, their plans may alter while you are in the process of recognizing them. Recognition may depend upon simulating the

generation of a plan, understanding why somebody might have produced the kind of plan they did.

Finally, in the standard AI model plans are generally reversible. In the box example, we can switch the goal state and initial state and make the same point. Exceptions occur only when the domain introduces a directionality, say, in breaking an egg. In contrast, human plans are basically historical. Actions are non-reversible because they interact with and are even defined in terms of a complex history of previous actions.

Mutual Belief Model

A number of people have been trying to find ways to represent intentions and beliefs and to model social plans as well as individual plans. Some of this work has been done with computer modelling (Cohen, 1981; Perrault & Allen, 1980; Sidner & Israel, 1981); some through analysis of conversations, stories, and skits. This section presents one such attempt based on the centrality of mutual belief (cf. Clark & Marshall, 1981; Cohen & Perrault, 1979; Schiffer, 1972). We mention here the major elements of a notation system that is being used in these analyses. For further details, see Bruce (1980) or Bruce and Newman (1978).

The notation system assumes an underlying facility for representation of actions and states. Every proposition is represented as embedded within a **belief**. A set of beliefs is indicated by a **belief space**. Figure 3 shows equivalent

representations of a proposition within the beliefs of some character or participant. We also need to represent intentions

---

Insert Figure 3 about here.

---

of at least two kinds. One is the intention to achieve something, to bring about a new state of affairs (Figure 4). The second is the intention to maintain some state (Figure 5).

Insert Figures 4 and 5 about here.

An important concept built out of individual beliefs is that of mutual belief. Mutual belief means not only that one person believes something that the other person believes, but also that each believes that the other person believes it and that the other believes that the first believes it, and so on indefinitely (Schiffer, 1972). Using the notation presented thus far we can define mutual belief (MB) of a proposition, X, between two actors, F and R, as follows:

MB (F, R, X) <==>
FB X
RB X
FB RB X
RB FB X
FB RB FB X
RB FB RB X
.
.
.

Closely related to the mutual belief concept is that of the social episode (Figure 6). People can create a frame for their

interactions within which they establish social facts to be true for the duration of that episode. Within the episode, participants have well defined roles. There are many conventional social episodes: For example, in a supermarket, there is a checkout clerk and a customer, with corresponding expectations about what each should do and know.

Insert Figure 6 about here.

The more interesting case is when people create new social episodes. There is then a set of initiating actions to establish the episode and a set of actions to close it. A social episode is represented as a mutual belief state with the participants' roles specified. Each side holds one participant's beliefs, intentions and actions. The entire episode resides within the larger belief space of one of the participants.

There are also rules for operating on these representations. For example, we have found it useful to allow the representation of a belief about one's own belief. But in general we collapse this representation by invoking a rule of the form:

SB SB X ===> SB X

Figure 7 shows how using this rule and the definition of mutual belief one can derive various equivalences within the notation system, e.g., that although "P believes it is mutually believed that X" is not equivalent to "Q believes it is mutually

believed that X," the conjunction of these two is equivalent to mutual belief. Such manipulations are useful since one of the most interesting things that has emerged in the examples that we have looked at is a shifting in and out of mutual belief--successive establishments of mutual belief, breakdowns, and then re-establishments.

Insert Figure 7 about here.

Todd and the Toys

In this section we look at a formal representation of a dialogue to see how mutual belief is established and manipulated in order to achieve personal goals. This analysis is not some magical device to see into the participants' heads, but rather a way to represent one observer's interpretation (mine) which might be compared then to someone else's. Thus the analysis is necessarily limited to a point of view.

The example, taken from a corpus of natural conversations (Hall, Linn & Nagy, 1984), is part of a conversation between a mother and her four and a half year old son. Let's call the boy, "Todd," and the mother, "Susan." Todd is in the living room playing with his father's tools and his mother is in an adjoining room:

1. Susan: Go into your room and play, Todd, and . . .
2. [She comes into the room,
3. and helps Todd put his father's tools away.

4. Todd holds a tool.]

5. S: What have you got now?

6. That's not a toy!

7. Todd: We're putting them back.

8. S: You're putting it back? Very . . .

9. T: Can I have these batteries?

10. S: No,

11. yeah, the batteries . . .

12. and how many times did Daddy say not to go in his kit?

Our analysis will allow us to address certain puzzling features of the dialogue. For example, notice that in lines 5, 6, and 8, Susan refers in the singular, whereas Todd, in line 7 uses the plural to refer to the same object(s). Note also that in line 7 Todd says that the two of them ("we") are putting away the objects, whereas, Susan in line 8 says that it is Todd's ("you") action which is in question. Are these examples of mis-communication? Do they just represent "noise" that one should expect in normal conversation? Do Todd and Susan see the object(s) in question as respectively singular and plural? Do they have different views about who is carrying out the actions in question?

The analysis to follow will make the case that, far from being noise, such apparent mis-matches are crucial to understanding what the dialogue is all about. In particular, the

difference arises because of simultaneous, but conflicting attempts to manipulate mutual belief.

Establishing Mutual Belief

Consider the dialogue at the point at which Susan says, "Go into your room and play, Todd." The essential thing Susan believes at this point is that they are entering into a social episode, working together for a common goal, which involves on Todd's part, getting the tools in the box. Figure 8 shows her belief that Todd wants to achieve the state, "tools in box." Meanwhile, Todd believes that Susan believes that it's a mutual belief that that's exactly what they are doing. But he is still playing with the tools. He also realizes that there is a conflict (represented by the dotted arrow) between his intention to play and the intention she wishes him to have--to get the tools into the box (Figure 9).

Insert Figures 8 and 9 about here.

Now let's consider Susan's view of the world in more detail. She believes Todd has the intention to play in his room, which means he must go in the room and play with his toys there. But she also wants him to get the tools in the box. Achieving that intention will take away the toys Todd is playing with, which, Susan understands, would be a conflict for Todd if he were, in fact, intending to play. But that conflict would be resolved by

the fact that once he gets into his room, there will be other toys to play with.

Todd is holding a tool in a manner that suggests play. He's clearly not putting it away. When Susan realizes that Todd doesn't have the intention of putting away the tools her belief about their mutual belief changes. The social episode starts to dissolve, and she must re-establish it as a social fact. Figure 10 says that Susan wants to re-achieve the mutual belief space that corresponds to this social episode. Within that intention is re-establishing Todd's intention to get the tools in the box.

Insert Figure 10 about here.

Recognizing the conflict between the intention she wants him to have and the one she believes he has leads her to act. She says (lines 5 and 6), "What have you got now? That's not a toy!" This is an attempt to change his intention by a retrospective establishment of meaning--defining what counts as a toy. She's also redefining the action, saying that what they are doing is not just putting tools away, it is putting non-toys away. These assertions directly address Todd's plan to play with the tools. Susan is saying, "If you look at what you are trying to do, you'll see that it's contradictory to the actual state of affairs." Note that Susan is acting on the basis of her interpretation of Todd's goals, forming hypotheses in the same way that we, as observers, are doing.

Todd (line 7) replies, "We're putting them back." He affirms that he wants to get the tools into the box; but to achieve that state he has to have the tools in his hand. Having a tool in his hand is a necessary subgoal to achieve the goal that Susan wanted all along. He may still have the intention of playing, but he clearly wants to achieve the state in which his mother believes that it's mutually believed that he is putting the tools away (see Figure 11). More precisely, his use of "we" in line 7 suggests that he either believes, or more likely, wants his mother to believe that they are engaged in a true social action in which they are both putting tools away. Regardless of whether he really wants to be putting them away, he wants to make his mother think that that's what's happening.

Insert figure 11 about here.

Susan continues (line 8) with a skeptical, "You're putting it back? Very . . ." She wants to keep things moving along so she confirms his reassurance despite her doubts about his true intentions. Using "you" to refer to Todd and "it" to refer to the tool she focuses on the specific act of putting a dangerous tool away. In effect, she asserts that if his statement and his physical actions support the assumption that the mutual belief state has been re-established, she will not question him further. She asserts that it is mutually agreed they are putting the tools

away (as in Figure 8), even though she may believe at some level that Todd still has the beliefs shown in Figure 11. Later (line 12) she reinforces this settlement and turns it into a teaching activity regarding what things count as toys.

Apparent Miscommunication

With the foregoing analysis in mind we can return now to consider the "them" versus "it" problem mentioned above. It may help in doing this to consider two hypothetical dialogues:

(A) Susan: Let's put these tools away.

Todd: OK. Where should we put them?

Susan: Put the big ones on the table and the little ones in the box.

(B) Todd: Let's play with the hammer.

Susan: OK. What should I do with it?

Todd: Use it to pound the table.

Dialogue (A) manifests the Put-Away view of the world. Note that under this view all the references are plural, i.e., they are to "tools" as a collection of undifferentiated objects, all of which are to be put away. Dialogue (B) manifests the Play view. In contrast to (A), the references in (B) are to a singular object ("the hammer"). Playing implies assigning singular, significant status to the objects used in the play. How Todd thinks of the tool may be more important than its nominal function, in fact, its usual function is only one idea for how it might be used in play.

Our analysis of the real dialogue suggests that it is Todd who has the Play view and Susan who wants to Put-Away. But notice how each refers to the tool(s). It is Todd who says, "We're putting them back,"; Susan who says, "That's not a toy" and "You're putting it back?". Thus they seem to switch roles.

There are two levels of reference to consider here. First, whether the, let's say, hammer, is one of many tools or a singular, interesting toy is a matter of individual beliefs and intentions. Thus, reference to it/them is conditioned by what participants believe. In the social world we must hypothesize about a person's beliefs even to infer what objects are physically available for reference.

Second, Susan and Todd are in a true dialogue in which they are each attempting to understand and alter each other's beliefs and intentions. What appears as a talking past each other is in fact a reflection of their coinciding attempts to take account of the other's beliefs. Thus, Susan, who wants the two to be in a Put-Away social episode, uses the singular reference, not because she sees the hammer as a plaything, but because she sees that Todd does, and wishes by her focus on it in his terms to more effectively alter his intentions. Similarly, Todd, who is undoubtedly in Play mode, probably has specific plans for the hammer as a singular, significant object. Nevertheless, he refers to it as part of a "them" because he sees that Susan is viewing it that way, and moreover, he sees that she intends for him to view it likewise. Independent of his decision whether to

go along with her desires, he decides it is useful strategically to appear to cooperate with her intention about what he should intend. His action appears in his form of reference.

The strategic use of reference to the "tools"--Susan's saying "it" because she believes Todd is thinking "they" when she wishes him to think otherwise (vice versa, in Todd's case)--is paralleled in the reference to the actor(s) involved. Thus, Todd says "we" because he believes Susan believes he is not engaged in the desired plan (i.e., that she believes that he doesn't intend to put away the tools) and he wants to have her change this belief about his intentions. Susan says "you" for converse reasons.

From this we see a resolution of the questions presented previously. Effective communication is not a function of exact reference, even if such a thing were possible. Instead, apparent mis-reference may actually reflect a high degree of coordination and understanding between participants who are operating in a dynamic social setting in which references are actions to communicate about and change one another's beliefs, intentions, and plans.

Open Questions

The approach presented here has been applied in studies of natural conversations, stories, and Sesame Street skits. These studies raise a number of important questions about the role of the author (in the case of presented social interactions), the

effects of modality, discourse conventions, representation of process, the size of planning units, complexity and background knowledge. One interesting phenomenon that emerges in the studies done thus far concerning the role of the author (or lack thereof). We saw this in our analysis of "Hansel and Gretel" (Bruce & Newman, 1978): Hansel goes along a path dropping pebbles. But he stops each time he drops one so that his mother and father become suspicious and say, "What are you stopping for now?" Hansel replies, "Oh, I'm looking at the cat that's on the roof." They say "that's not a cat, that's just the sunlight reflecting on the roof."

This interaction is unnecessary from the point of view of just marking a trail. One reason it happens is that the authors are trying to highlight the importance of Hansel's actions for the reader, and also to indicate who believes what. They show (1) that Hansel is deliberately dropping pebbles, (2) that he knows what's going on, and (3) that his parents don't. When we're analyzing "Hansel and Gretel" we might temporarily ignore the author and just look at Hansel and Gretel in their interaction with their parents. But such an analysis is necessarily tentative. In general, one must integrate the character level analysis with analysis of the interaction of plans and beliefs between the author and the reader.

We need to investigate the extent to which models of planning can be used in rhetorical theory (see Brewer & Lichtenstein, 1982), in particular, the relationship between

author and reader in terms of author's goals and reader's goals, author's beliefs about reader's beliefs, reader's beliefs about what the author believes about the reader, and so on. When there are stories within stories each of the levels adds two more participants, each of whom has beliefs about each other, and also about characters in their stories (Bruce, 1981).

Closely related to questions about author/reader relationship are questions about discourse structure and convention. How much does the carrying out of a plan by a character constrain the discourse structure that can be produced? How do we separate the conventional aspects of a story (e.g., setting, character stereotypes, and so on) from the underlying plans?

There have been some efforts at computer simulation of social planning (Cohen & Perrault, 1979; Perrault & Allen, 1980). This work addresses many of the limitations of the robot model work (see Sidner, Bates, Bobrow, Goodman, Haas, Ingria, Israel, McAllester, Moser, Schmolze, & Vilain, 1983). Designing systems that can cope with belief-based models is a formidable task.

Another question concerns the size of planning units: To what level should we break down actions? For example, should the action of putting away tools be taken as a unit? It is a subpart of the action of going into the bedroom to play and has subparts such as picking up tools and putting them in a box. The appropriate level of representation appears to be a consequence

of the interaction itself. This raises important questions about the formal representation of plans.

Another interesting question regards complexity: Producing plans and recognizing plans are both complex tasks, but what is the main contributor to complexity? Newman (1982) has analyzed several Sesame Street skits (starring Bert and Ernie) showing various degrees of complexity in the skits themselves and in viewers' interpretations. For example, one skit has six distinct interpretations. From this work it appears that the degree of embedding of intentions and beliefs is not the primary source of complexity differences. What appears more salient is the complexity of beliefs that need to be maintained outside of the mutual belief space.

A final question is how background knowledge comes into play. In the Sesame Street skits it helps to know that Ernie typically tricks Bert, despite Bert's greater knowledge about some domains. Similarly, in "Hansel and Gretel," beliefs about cutting wood, building fires, and so on, are crucial both for characters in the story and the reader.

Summary

This paper has examined the role of plan recognition and plan generation in communication. One idea that emerges from this examination is that people in interaction with others appear to organize their perceptions of a social situation in terms of plans. They do this even when the others' plans are poorly formulated. They use their models of others' plans in

formulating their own plans. Much of what occurs in discourse centers on a continual communication about and reformulation of one's own plans and one's models of others' plans.

A second idea that emerges is that formal methods for describing and analyzing plans are now available. Such methods allow us to be more explicit in our hypotheses about social interaction. Unfortunately, the classical formalisms for planning are derived from a robot world model that fails to generalize sufficiently to account for typical human planning situations.

By pushing the classical model we come to the third main idea of this chapter, namely, that a model of planning with concepts such as mutual belief at its core may take us further in our studies of discourse and other forms of social interaction.

In particular, such a model addresses a number of characteristics central to human planning: Plans are essentially social, involving cooperation and/or conflict between individuals. They operate in a space of social facts, not just physical states. As such, they are built upon participants' beliefs and their beliefs about each others' beliefs. The classical model's state variables, operators and goals are then beliefs of participants that evolve throughout the interaction. This is one reason why the human planning world emphasizes talk about plans, not just design and execution of them. Moreover, generation and recognition of plans become necessarily

intermingled since each individual's plan is only part of a larger social plan. Finally, the human plan has a historical character to it. Each action produces irreversible changes in others' beliefs, including, of course, the belief that the given action was carried out.

Using a mutual belief model that attempts to take account of these characteristics of human plans we looked at a simple dialogue. It is clear from the analysis that elements of the model, e.g., mutual belief, social episode, goal conflict, and so on, are necessary for modeling such dialogues, but not sufficient. More complex dialogues and texts will undoubtedly require further elaboration of the model.

References

Brewer, W. F., & Lichtenstein, E. H. (1982). Stories are to entertain: A structural-effect theory of stories. Journal of Pragmatics, 6, 473-486.

Bruce, B. C. (1980). Analysis of interacting plans as a guide to the understanding of story structure. Poetics, 9, 295-311.

Bruce, B. C. (1981). A social interaction model of reading. Discourse Processes, 4, 273-311.

Bruce, B. C., & Newman, D. (1978). Interacting plans. Cognitive Science, 2, 195-233.

Clark, H. H., & Marshall, C. R. (1981). Definite reference and mutual knowledge. In A. K. Joshi, I. Sag, & B. Webber (Eds.), Elements of discourse understanding. Cambridge: Cambridge University Press.

Cohen, P. R. (1981). The need for identification as a planned action. Proceedings of the International Joint Conference on Artificial Intelligence. Vancouver, B.C.

Cohen, P. R., & Perrault, C. R. (1979). Elements of a plan-based theory of speech acts. Cognitive Science, 3, 177-212.

Gearhart, M., & Newman, D. (1980). Learning to draw a picture: The social context of an individual activity. Discourse Processes, 3, 169-184.

Hall, W. S., Linn, R., & Nagy, W. (1984). Spoken words: Effects of situation and social group on oral word usage and frequency. Hillsdale, NJ: Erlbaum.

Miller, G. A., Galanter, E., & Pribram, K. H. (1960). Plans and the structure of behavior. New York: Holt.

Newman, D. (1981). Children's understanding of strategic interaction. Unpublished doctoral dissertation, City University of New York.

Perrault, C. R., & Allen, J. F. (1980). A plan-based analysis of indirect speech acts. American Journal of Computational Linguistics, 6, 167-182.

Schiffer, S. R. (1972). Meaning. London: Oxford University Press.

Sidner, C., Bates, M., Bobrow, R., Goodman, B., Haas, A., Ingria, R., Israel, D., McAllester, D., Moser, M., Schmolze, J., & Villain, M. (1983). Research in knowledge representation for natural language understanding. Cambridge, MA: Bolt Beranek and Newman, Inc., Report No. 5421, August.

Sidner, C. L., & Israel, D. J. (1981). Recognizing intended meaning and speaker's plans. Proceedings of the International Joint Conference on Artificial Intelligence, Vancouver, B.C., August.

Wilensky, R. (1981). Meta-planning. Representing and using knowledge about planning in problem solving and natural language understanding. Cognitive Science, 5, 197-234.

Figure Captions

Figure 1. Initial state in the robot world.

Figure 2. Goal state in the robot world.

Figure 3. Representation of beliefs: Person P believes proposition X.

Figure 4. Representation of intentions: P intends to achieve X.

Figure 5. Representation of intentions: P intends to maintain X.

Figure 6. A social episode.

Figure 7. Equivalences among belief representations.

Figure 8. Susan's initial model: Cooperative action.

Figure 9. Todd's initial model: Goal conflict.

Figure 10. Susan's recognition of the goal conflict.

Figure 11. Todd's goal of conflict resolution.

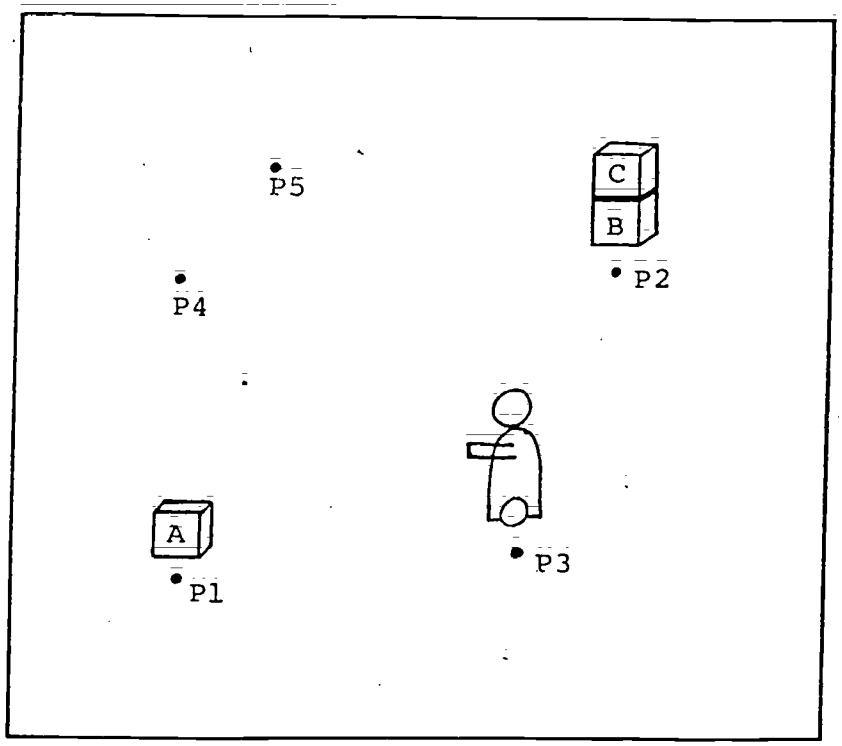
P5
C
B
P2
P4
A
P1
P3

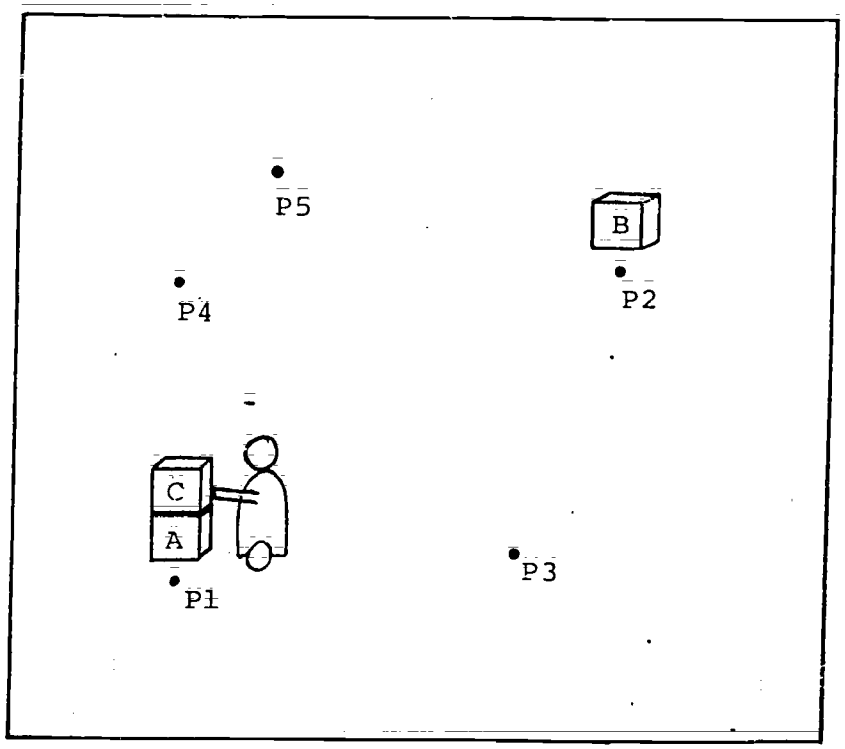
P5
B
P4
P2
C
A
P3
P1

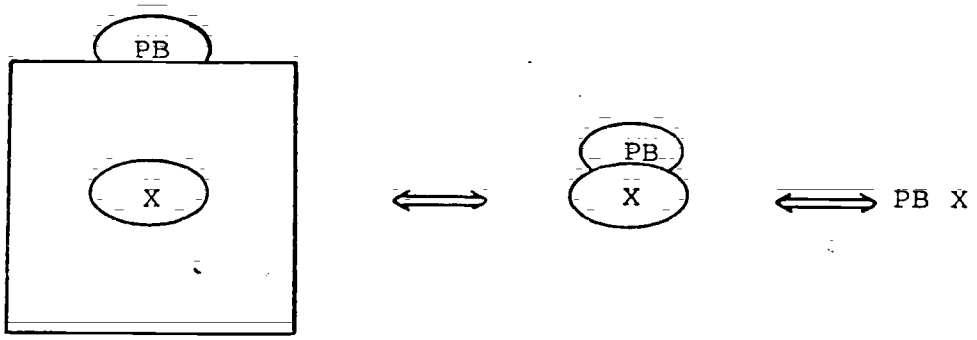
PB
X
PB
X
PB X

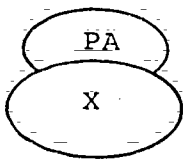
PA
X

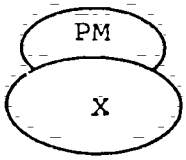
PM
X

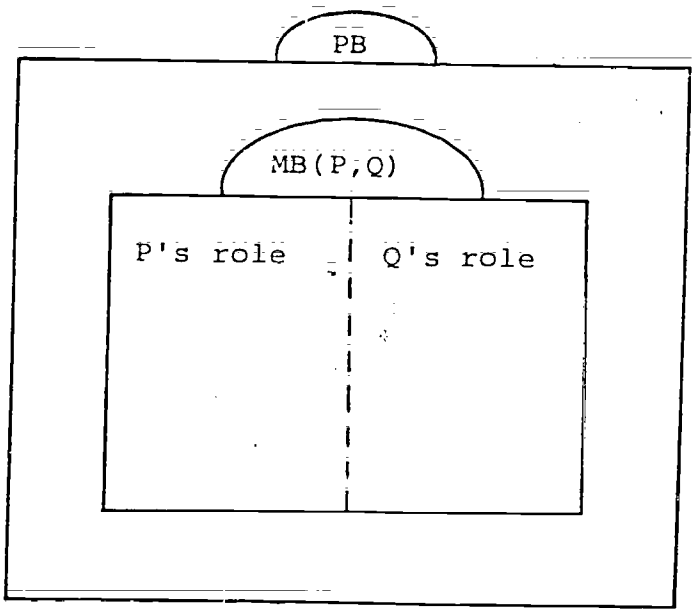
PB
MB(P,Q)
P's role
Q's role

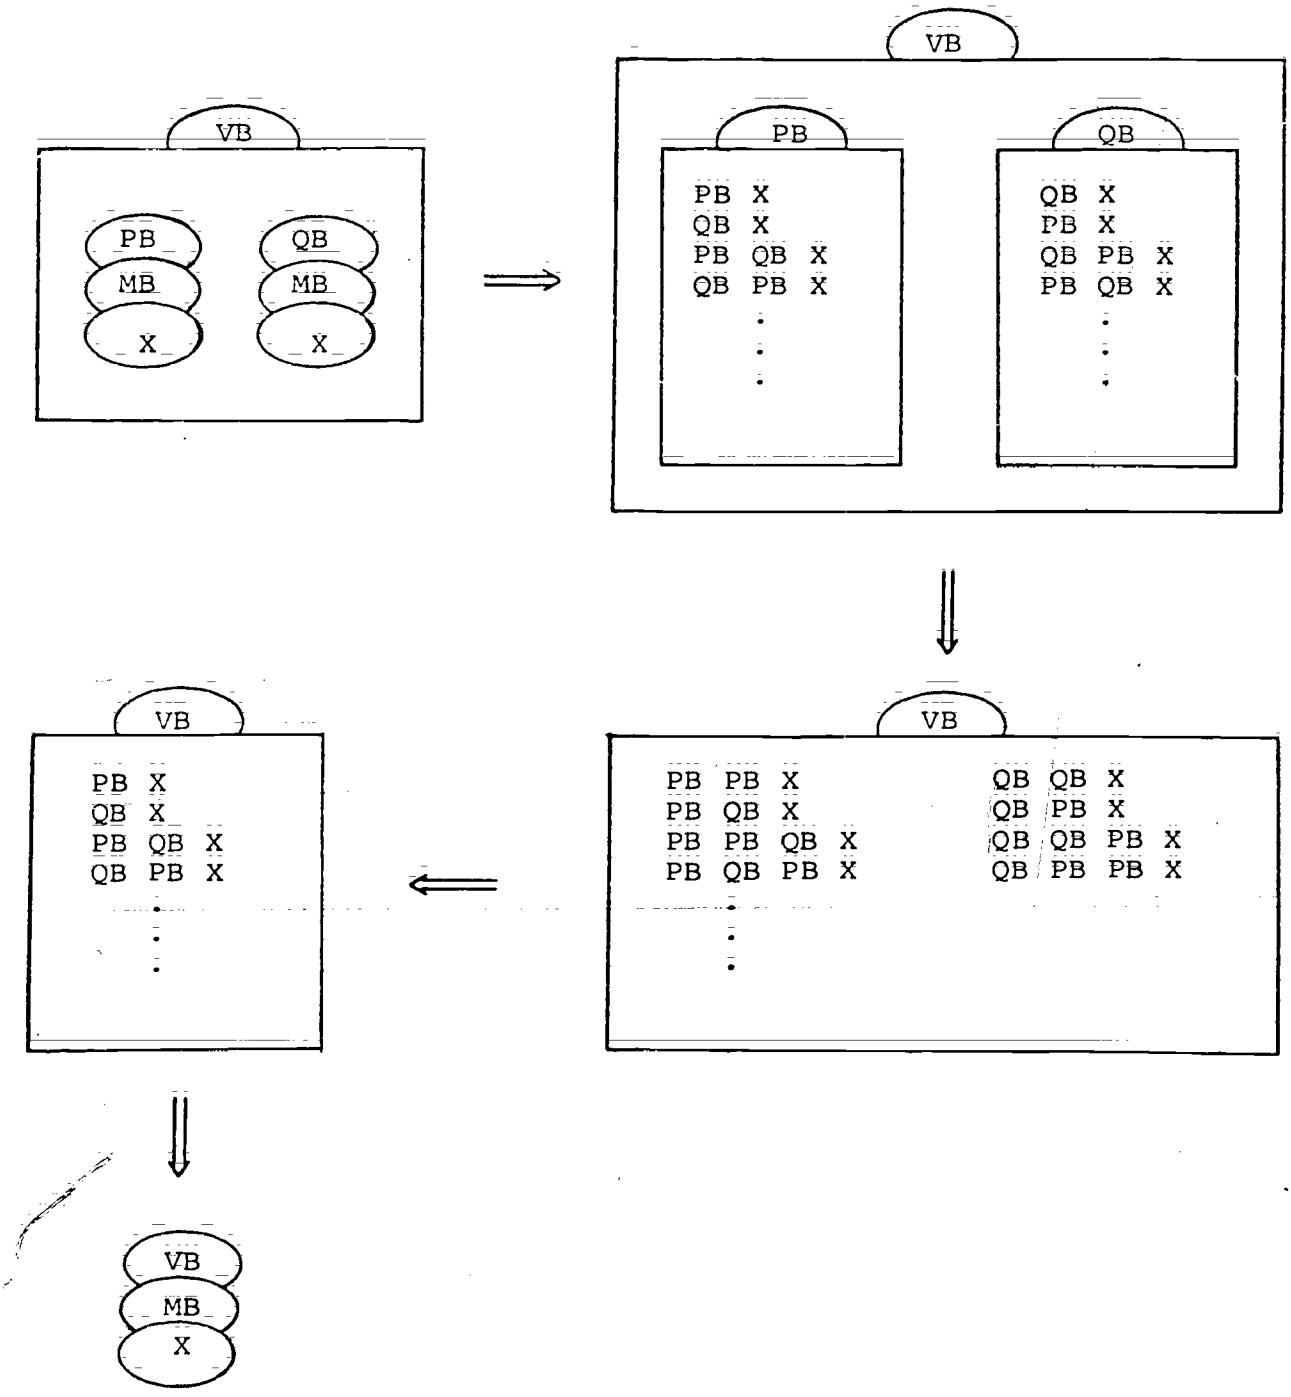
VB
PB
MB
X
QB
MB
X
VB
PB
PB X
QB X
PB QB X
QB PB X
QB
QB X
PB X
QB PB X
PB QB X
VB
PB PB X
PB QB X
PB PB QB X
PB QB PB X
QB QB X
QB PB X
QB QB PB X
QB PB PB X
VB
PB X
QB X
PB QB X
QB PB X
VB
MB
X

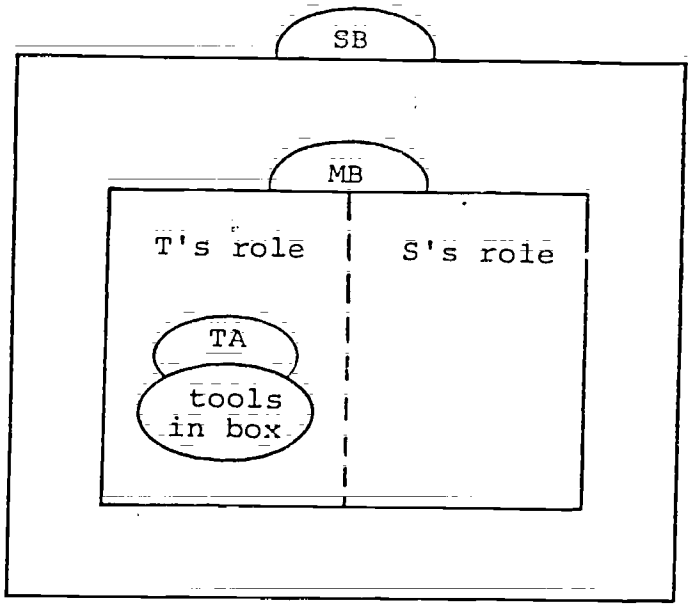
SB
MB
T's role
S's role
TA
tools
in box

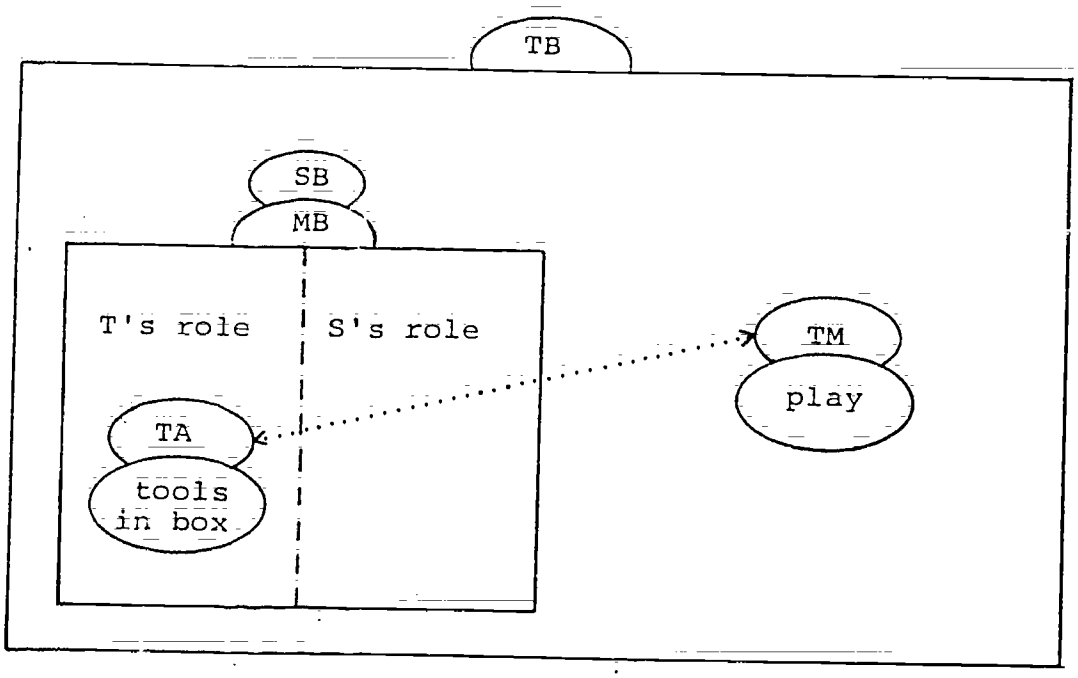
TB
SB
MB
T's role
S's role
TM
play
TA
tools
in box

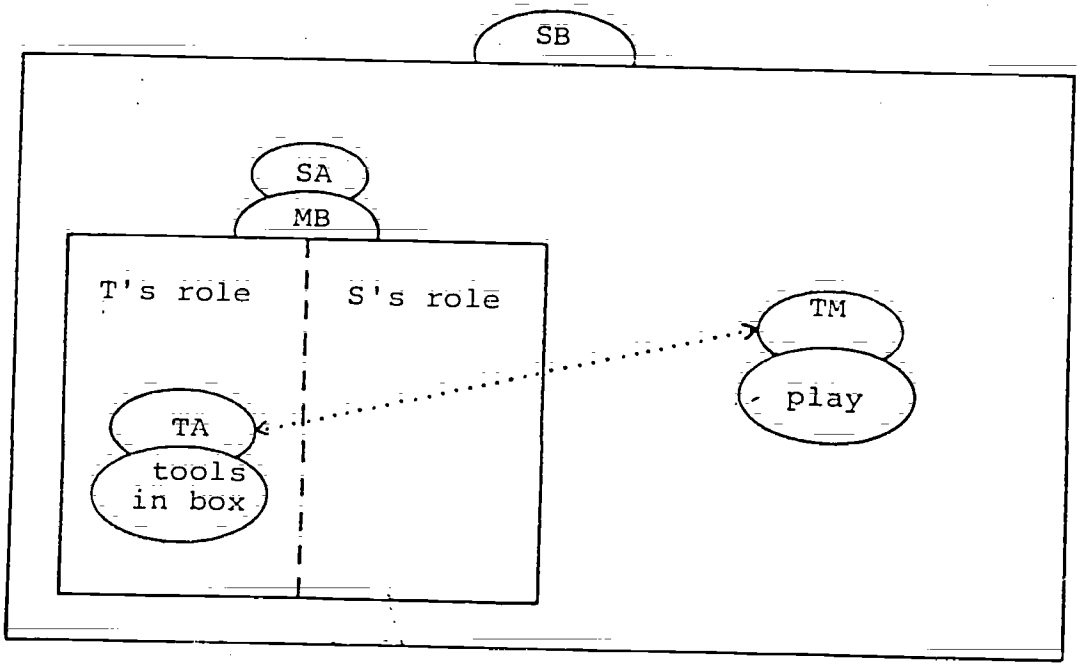
SB
SA
MB
T's role
S's role
TM
play
TA
tools
in box

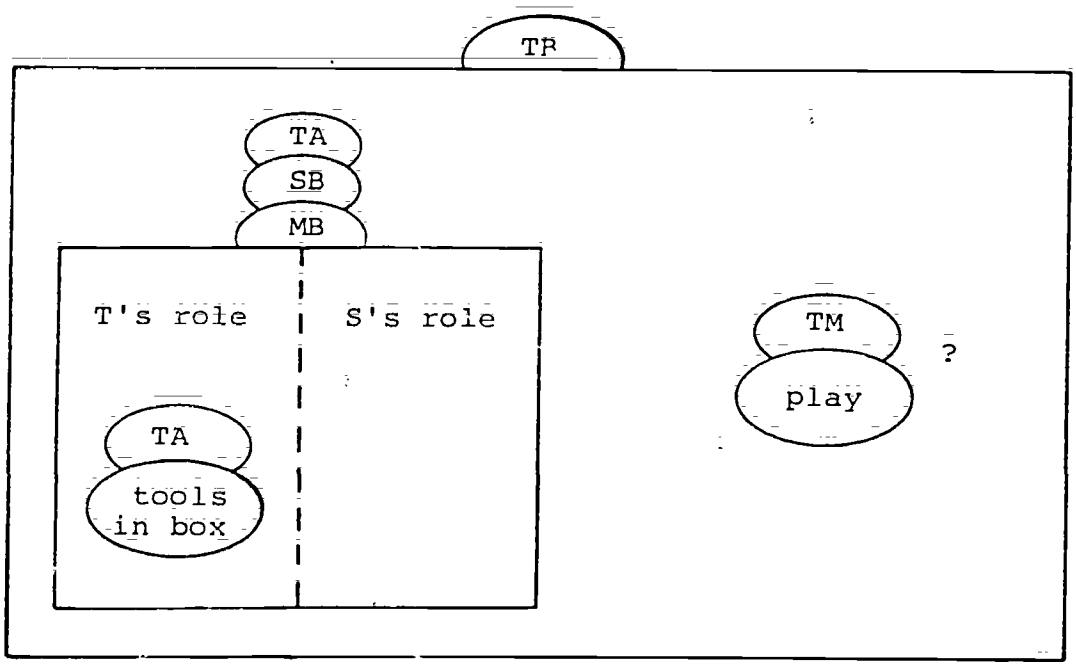
TR
TA
SB
MB
T's role
S's role
TA
tools in box
TM
play
?