

Appears in P. R. Cohen, J. Morgan, and M. E. Pollack, eds., *Intentions in Communication*, MIT Press, 1990

PLANS AS COMPLEX MENTAL ATTITUDES

Martha E. Pollack
Artificial Intelligence Center
SRI International
and
Center for the Study of Language and Information

1 Introduction

There are plans and there are plans. There are the plans that an agent “knows”: essentially recipes for performing particular actions or for achieving particular goal states. And there are the plans that an agent adopts and that subsequently guide his action. The distinction is between knowing that a plan for assassinating the president is shooting him, and actually planning

⁰Preparation of this paper was supported by a gift from the Systems Development Foundation. The research was done as part of my doctoral thesis [19], which was supported by a gift from the Systems Development Foundation, by an IBM Graduate Fellowship, by the Defense Advanced Research Projects Agency under Contract N00039-84-K-0078 and by the Office of Naval Research under Contract N00014-855-C-0013. My sincere thanks to Barbara Grosz, and to all the others, too numerous to list here, who contributed to my thesis effort.

to assassinate the president by shooting him.

To keep matters straight, we can refer to what one knows when he knows a way to do something as a *recipe-for-action*. Schematically, if we let r denote some recipe-for-action, A denote an agent, and PLANS be a binary relation between agents and the recipes-for-action they adopt, $\text{PLANS}(A, r)$ will denote that A has a plan to do r . (I am temporarily suppressing issues of time.) A potential ambiguity remains, however. For there is A 's plan r , the particular recipe-for-action he has adopted, and there is the state of mind that A is in when $\text{PLANS}(A, r)$ is true. It is this latter distinction that Bratman is addressing when he notes that in speaking of an agent's plan we might "mean an appropriate abstract structure—some sort of partial function from circumstances to actions, perhaps. On the other hand, [we might] mean an appropriate state of mind, one naturally describable in terms of such structures" [5, p. 271]. We thus need more terminology: if $\text{PLANS}(A, r)$ is true, we can say that the recipe-for-action r is *the object* of A 's plan, which is itself what Bratman calls a "state of mind." Indeed, I will argue that the plan itself is usefully seen as a complex mental attitude, one comprising a structured collection of beliefs and intentions.

Artificial Intelligence research in plan generation has always been primarily concerned with the recipes-for-action that are the object of an agent's plans. The focus of such research has been on automating the process by which an agent can compute a recipe-for-action. Thus Nilsson describes

plan generation as “the problem of synthesizing a sequence of robot actions that will (if properly executed) achieve some stated goal, given some initial situation” [18, p. 275]. Although the assumption has always been that the computed recipe-for-action would be adopted as the object of a plan, and some robots have been designed that did actually adopt and execute the recipe-for-action, the plan generation literature has paid very little attention to plans *qua* mental attitudes.¹ The same is true of much of the plan inference literature in AI, which began later than the work on plan generation and which inherited many of its techniques.

In this paper, I claim that a model of plan inference adequate to support a theory of cooperative communication must concern itself with the structure of the complex mental attitude of having a plan, as well as with the structure of the objects of that attitude. In Section 2, I recount the traditional AI approach to performing plan inference in communication, and describe some limitations of that approach. In Section 3, I present an alternative analysis of plans, one which emphasizes their nature as action-guiding, complex mental phenomena. Section 4 describes the role of this analysis of plans in a theory of cooperative communication that is not subject to the limitations discussed in Section 2. In Section 5, I return to the traditional AI model of plans and show how it is subsumed by the model presented here.

¹Recent work by Georgeff and Lansky [12] is an exception.

The theory of cooperative communication discussed in Section 4 has been developed in some detail, and has been implemented in SPIRIT, a system that reasons about the plans underlying the queries that it is asked and that generates appropriate responses. SPIRIT is able to infer plans even when they are invalid, i.e., when their objects are not recipes-for-action that SPIRIT would itself construct. This ability distinguishes SPIRIT from earlier plan inference systems, which do not rely on a careful analysis of plans as complex mental attitudes. Details of SPIRIT and of the plan inference framework it embodies can be found in [19].

2 The Traditional Approach

If you overheard the following conversation, you would probably find it quite unremarkable:

A: “I want to talk to Kathy, so I need to find out the phone number of St. Eligius.”

S: “St. Eligius closed last month. Kathy was at Boston General, but she’s already been discharged. You can call her at home. Her number is 555-1238.”

Intuitively, it is quite clear what is occurring in this discourse. Agent *A* believes that Kathy is at St. Eligius, and plans to call her there. Agent *S* believes that *A*’s intended act of calling St. Eligius cannot be performed,

since St. Eligius is closed. Moreover, S believes that even if A could call St. Eligius, it would not contribute to his goal of talking to Kathy, because she is actually at home. Consequently S , being cooperative, provides A with information she believes will contribute to his goal: she tells him the correct phone number for reaching Kathy. She also tells him why she believes the information he requested was not appropriate to his goal.

Conversations such as this one, in which the beliefs of the inferring agent differ significantly from the beliefs of the actor whose plan she is inferring, provide a serious challenge to most existing AI systems for plan inference in communication (e.g., [1, 7, 17, 21]). To see why, it is instructive to consider briefly how these systems work. Each of them has an *operator library*, which encodes a set of recipes-for-action. These operators are a direct outgrowth of the representations first developed in the STRIPS system [10] and later expanded in the NOAH system [20]. Each operator may contain some or all of the following parts:

- *a header*, which names the action α for which the operator is a recipe;
- *a precondition list*, which describes what must be true for α to be performed;
- *an effect list*, which describes what will be true after α is performed;

- *a list of constraints*, which describes restrictions on legal instantiations of the operator ²;
- *a body*, which may be a set of subactions whose performance constitutes performance of α , or a set of subgoals whose achievement constitutes performance of α .

The operators in the library represent relatively simple recipes-for-action. More-complex recipes can be constructed out of the simpler ones. Both in plan generation and plan inference, the more-complex recipes-for-action are represented as directed acyclic graphs, whose nodes are labeled with operator headers and propositions; when a node is labeled with a proposition P , it should be interpreted as denoting any action that would achieve P . Graphs, rather than linear orderings, are necessary because the recipes are seen as having both a hierarchical and a temporal dimension. Recipe graphs are sometimes referred to as macro-operators.

All the plan inference systems include a set of rules for constructing recipe graphs out of the simpler recipes in the operator library. Each rule states the conditions under which a piece of a recipe graph can be constructed. The conditions for constructing a subgraph always refer both to

²For example, restrictions on the types of parameters, on the relations between parameters, and on the ordering of the subactions into which an action is decomposed. Constraints can be thought of as preconditions that the planning agent never attempts to achieve.

the subgraphs that have already been constructed and to operators in the operator library. A typical rule, for example, allows a subgraph that includes a node labeled α to be expanded by adding an arc from α to a new node labeled β , if there is an operator whose header is β and which includes α in its body. In all, there are only three basic conditions under which the various plan inference systems will construct a recipe subgraph. These are shown in Figure 1. Note that each condition corresponds to one of the ways that actions and propositions can be related in an operator.³

Plan inference systems begin with some action α believed to be part of the recipe-for-action that is the the object of the actor’s plan. They then attempt to construct larger recipes-for-action by repeatedly applying inference rules, until one of the computed recipes satisfies some termination condition. (Typically, the termination condition is that the recipe constructed is for an action that an agent in the domain is likely to want to perform.) The computed recipe is taken to be the object of the actor’s plan.

³Several of the plan inference systems have additional rules which apply only to nodes that encode information-seeking actions. For example, many systems will construct a link from a node encoding “finding out whether P” to one encoding “achieving P.” This construction is meant to capture the intuition that if an agent want to know whether P is true, he may want P to be true. Kautz [15] has shown how these rules, along with rules for handling nesting plan inference, can be viewed as special cases of the three relations shown in Figure 1.

Figure 1: Rules for Constructing Plan Graphs

In fact, the inference rules are often written in such a way as to acknowledge the importance of plans as mental attitudes; after all, inferring another agent’s plan means figuring out what actions he “has in mind.” Consider Allen’s model, which was one of the earliest accounts of plan inference in conversation and inspired much of the subsequent work in the field [1, 2]. Allen writes $AW(P)$, where P is some proposition, to express the fact that A has a plan to achieve P ; $AW(ACT)$, where ACT names some action, expresses the fact that A has a plan to perform ACT . A typical plan inference rule, then, is expressed as

$$SBAW(P) \rightarrow SBAW(ACT), \text{ if } P \text{ is a precondition of } ACT.^4$$

This rule corresponds to one direction of application of Condition (1) of Figure 1.⁵ It can be glossed as follows: if the system (inferring agent) believes that the actor wants some proposition P to be true, then the system

⁴This rule should not be seen as a sound rule of logical inference, but rather as a rule that suggests an inference that is likely to be true. Cf. Kautz and Allen[16].

⁵Condition (1), as well as Conditions (2) and (3), can be applied in “either direction”: they can be used to add β when α is already in the subgraph, and they can be used to add α when β is already in the subgraph. Allen’s Precondition-Action Rule, given above, corresponds to the former case for Condition (1). His system also includes the converse rule:

$$SBAW(ACT) \rightarrow SBAW(P), \text{ if } P \text{ is a precondition of } ACT.$$

may draw the inference that the actor wants to perform some action ACT of which P is a precondition. Notice that it is left unstated precisely who it is—the system or the actor—that believes that P is a precondition of ACT. If we take this to be a belief of the system, it is not clear that the system will infer the actor’s plan; but, on the other hand, if we take it to be a belief of the actor, it is unclear how the system comes to have direct access to it. In practice, there is only a single set of operators relating preconditions and actions in Allen’s system: these represent recipes-for-action that are assumed to be mutually known to the system and the actor.

In effect the “SBAW” context is transparent to the reasoning process that is performed by Allen’s system: its reasoning is all performed directly on its object, and the B and W operators are carried directly from antecedent to consequent in each inference rule. Without any repercussions, the B and W operators can be omitted, resulting in rules that are completely equivalent to those in Figure 1. In fact, in his example plan graphs, Allen often omits the B and W operators entirely, and this practice has been continued in more recent work in plan inference [7, 17]. For Allen, as well as for most of those whose work he inspired, “the properties of the W operator [are] specified only by the . . . plan inference rules.” Allen’s W operator is equivalent to the PLAN relation introduced in the introduction to this paper: it is the

relationship that holds between an agent and a recipe-for-action he adopts.⁶ His plan inference rules apply essentially directly to the object of the plan under construction. Thus Allen, and his followers, analyze the state of having a plan only in terms of the structure of its object.

In sum, the traditional approach to plan inference has been to reason directly about the object of an actor’s plan, constructing a recipe-for-action using a library of simpler recipes that are assumed to be mutually known to the actor and the inferring agent. While this approach has proved to be successful for modeling a number of conversational phenomena, it has at least three shortcomings that preclude it from handling in a principled way dialogues like the one presented at the beginning of this section—dialogues in which one of the participants has a plan that the other deems to be invalid.

⁶Actually, there seems to be some tension regarding what is really meant by the W operator. Allen states that $AW(P)$ means “A has a goal to achieve P,” which seems to imply that P is a single action or property, not a whole plan. Consistent with this, he says that “ $SBAW(X) \rightarrow SBAW(Y)$ ” should be taken to mean that “if S believes A has a goal of X, then S may infer that A has a goal of Y” [2, p. 120]. But he uses these rules to infer not just that A has a goal of Y, i.e., that his plan contains Y, but also that it contains X and Y related to one another in some particular way specified by the rule. So the rule relation preconditions and actions mentioned above should probably be written “ $SBAW(P) \rightarrow SBAW(P \rightarrow (\text{is-a-precondition-of}) ACT)$, if P is a precondition of ACT.” Writing the rule this way would clarify his model, but it would not affect the claim that the B and W operators are transparent to the inference rules.

First, such a system will infer a plan subgraph linking nodes labeled X and Y, where X and Y are either operator headers or propositions, only if X and Y are both encoded in the system’s operator library in one of the configurations shown in Figure 1. This fact, combined with the fact that the operator library includes only valid domain information, may at first seem to suggest that traditional systems cannot infer *any* invalid plans. It turns out, however, that there are certain types of invalid plans that can, in principle, be inferred by a traditional system. For example, such a system could use a rule based on Condition 3 of Figure 1 to infer a subgraph joining some actions α and β , even if the system also believes that one of the preconditions in the operator relating α and β is false.⁷ However, without understanding why such a subgraph might be part of the object of the actor’s plan, the system will not be able to determine whether or not it is reasonable to include such a subgraph in the inferred plan. In fact, existing systems have not focused on the inference of such invalid plans, and indeed have included various control heuristics that bias the inference process against finding them (see, e.g. Allen’s Heuristic (H1) [2, p. 127]). The first shortcoming of the traditional approach is an inability to state within it general heuristics for

⁷Strictly speaking, the truth or falsity of the precondition at the time of the inference should not be what is at issue, but rather its truth or falsity at the time that the agent executes the action. Since the traditional systems have no explicit representation of time, however, this distinction collapses for them.

rule application.

Second, there are other types of invalid plans that cannot be inferred at all under the traditional approach. For instance, an actor A may be relying upon a simple recipe-for-action that is not in the system's operator library. In the case mentioned in the previous paragraph, S and A mutually know a recipe-for-action linking α and β : A 's plan is invalid because the necessary preconditions for doing β by doing α do not hold. However, there are other cases in which A plans to do some action β by doing some α , where S does not at all believe that α will, under any set of preconditions, lead to β . (Indeed, S may even believe that α —or β , or both—cannot be done at all.)

One way to enable a system to reason about plans with such invalidities is to encode directly sets of erroneous beliefs that handle such cases by encoding directly sets of erroneous recipes that their users are likely to have; this is the approach taken by many computer-aided instruction (CAI) systems [6, 11, 8, 22]. Although this seems to be a useful strategy, it is necessarily incomplete. It is impossible for any person to have complete knowledge of the potential beliefs of other people, since the range of beliefs is, in principle, infinite. This means that system designers cannot anticipate *a priori* all of the potential misconceptions the users of their system may have. It also means that the intelligent agent that such systems emulate—the human being—cannot know *a priori* all of the potential misconceptions that people asking her questions might have. Sometimes she will have to

deal with a novel (to her) belief. A more general approach is to try to understand the structure of the beliefs that an agent has when she has some plan, thereby making it possible to apply various belief-attribution techniques to the problem of plan inference.

A third weakness of the traditional approach concerns the generation of appropriate responses in conversation. To explain why *S*'s response in the example dialogue is a cooperative one, it is necessary to understand what beliefs *S* is attributing to *A* by virtue of attributing to her some plan.

To handle these issues—to make room in the theory of plan inference for discrepancies between the beliefs of the actor and the inferring agent, thereby allowing for the proper treatment of “invalid” plans—it is necessary to undertake an analysis of what the context “SBAW” means: to analyze plans as mental phenomena. It is that task to which we now turn.

3 An Alternative Model of Plans

To develop an account of plans as mental phenomena, we could begin with the traditional AI models, which, as we have seen, have focused on the objects of plans: we could attempt to determine what sort of mental states would reasonably have such objects. However, it will prove to be more fruitful to begin with our commonsense conceptions of what it means to have a plan. The resulting model can be shown to subsume and make more

precise the model implicit in the traditional AI accounts.

3.1 The Belief Component of Plans

Let us begin with the plan (referred to earlier) to ask Kathy how she is feeling. I plan to do this, believing that Kathy is at St. Eligius, by finding out the phone number of St. Eligius, calling there, and then saying to Kathy “How are you doing?”.⁸ The performance of these acts is meant to “entail”—in a sense of entail yet to be further specified—the performance of my goal act.

One’s plans, however, may fail. If, unbeknownst to me, Kathy has already gone home, then my plan will not lead to my goal of asking her how she is feeling. For me to have a plan to do β , which consists in the doing of some collection of acts Π , it is not necessary that the performance of Π actually lead to the performance of β . What is necessary is that I *believe* that its performance will do so.⁹ This insight is at the core of a view of plans

⁸Throughout the rest of this paper, I will make the simplifying assumption that when one dials a hospital, one reaches directly the person with whom one wants to speak. If the reader is uncomfortable with this simplification, the action of asking for Kathy can be inserted before the action of saying to Kathy “How are you doing?”.

⁹In fact this condition may be slightly too strong: the agent need not be sure that performing his plan will entail performing his goal. In the normal state of affairs, though, he will at least think this likely and will act as if he believed it. See the further discussion in Section 3.3.

as mental phenomena; on this view plans “exist”—i.e. gain their status as plans—by virtue of the beliefs of the person whose plan they are.

So far, then, I have associated the state of having a plan to do β with a belief that executing some collection of acts Π will lead to doing β . Note that the temporal ordering of the acts is an essential part of the plan: I may well have a plan to prepare onions for a sauce by chopping them and then sauteeing them—believing that by so doing I will perform my goal—and not yet have a plan that involves sauteeing and then chopping the onions. However, the ordering need not be total: as Sacerdoti demonstrated [20], there are many plans with objects that include acts whose temporal order with respect to one another is irrelevant. For example, I may have a plan to set the table that includes the following acts: carrying the flatware, plates, and glasses to the table; setting out the flatware; setting out the plates; and setting out the glasses. I may believe that it is essential that the first act in the list be performed prior to the others, but may also believe that the other three acts can be performed in any order with respect to one another, and can even be interleaved. Of course, when I actually execute my plan, the acts I perform will be totally ordered with respect to one another. So there is a sense in which the beliefs that are part of my plan are partial.

There is also another sense in which my beliefs may be partial: they may concern acts only to an arbitrary level of abstraction. The set of acts that I believe will entail my asking Kathy how she is feeling includes the act of

finding out the phone number of St. Eligius. It includes this despite the fact that I may not yet have considered how I will do this—for example, whether I will call the information operator or look in a phone book.¹⁰

An agent A 's belief that performing the acts in Π will entail performing β is not by itself sufficient to guarantee that A has a plan, consisting of doing Π , to do β . To see why not, consider the following scenario. Suppose that I decide that while I am finding out the phone number of St. Eligius (say by looking in the phone book), I might as well at the same time find out the phone number of my bank (perhaps because I know I have to call there later to check on a wire transfer). As before, I believe that finding out and dialing the phone number of St. Eligius, and then saying certain words, will entail my asking Kathy how she is feeling. I do not believe that this will cease to be true if I also find out my bank's phone number. Thus, there is a temporally ordered collection of acts Π , which equals finding out the phone number of St. Eligius, finding out the phone number of my bank, calling St. Eligius, and saying to Kathy "How are you doing?", such that I believe that executing Π will lead to my goal of asking Kathy how she is feeling.¹¹ However, it seems incorrect to say that my plan to ask Kathy how

¹⁰Bratman [4] discusses the significance of partiality of plans in resource-bounded agents like humans and robots.

¹¹Notice that Π in this example need only be partially ordered: I may consider it irrelevant whether I first find out the phone number of St. Eligius and then the phone number of my bank, or vice versa.

she is feeling includes my act of finding out the phone number of my bank; instead, this act is part of another plan (to check on my wire transfer) that I intend to interleave with my original plan. For an act to be included in my plan, I must believe that it plays a role in that plan; for the case at hand I do not believe that finding out the phone number of the bank plays any role in my plan to ask Kathy how she is feeling.

3.2 Playing a Role in a Plan

What does it mean for an act to play a role in a plan? Consider once more my plan to ask Kathy how she is feeling. Part of this plan, I claimed, involved my calling St. Eligius. What would we say about my beliefs about these two acts: my asking Kathy how she is feeling and my calling St. Eligius? We might say that I believe that the latter will enable the former; i.e. that by calling St. Eligius I will establish a communication channel—a phone link—to Kathy, which will enable my saying something to her. Similarly, we might say that I believe that finding out the phone number of St. Eligius plays a role in my plan because I believe that doing the former will enable my calling St. Eligius, which itself plays a role in my plan. In general, then, if an agent believes that doing one act α will enable either his goal or some other act γ that plays a role in his plan, then α may play a role in his plan. In order to strengthen the “may” to “will,” we need to consider the agent’s intentions. I will discuss this presently.

Next consider the relationship between my acts of saying to Kathy “How are you doing?” and asking her how she is feeling. We would not say that I believe that the first act will enable the second. Instead we could describe my beliefs using the “*by* locution” in English: we could say that I believe that *by* saying “How are you doing?” I will be asking Kathy how she is feeling. Similarly, consider a slightly more detailed analysis of my plan to ask Kathy how she is feeling, which makes explicit the way in which I go about finding out the phone number of St. Eligius. One way in which I might plan to do this is by looking the number up in the phone book. We would not say that I believe that my looking up St. Eligius’s phone number will *enable* my finding it out; rather, we would say that I believe that by looking it up I will find it out. Or I might, instead of looking up the phone number, plan to discover it by getting my office mate to tell it to me, as a result of my asking her to tell it to me, which consists in my uttering the question “Do you know the phone number of St. Eligius?” Once again, we would not say that I believe that my uttering the question “Do you know the phone number of St. Eligius?” will *enable* my asking my office mate to tell me the phone number of St. Eligius, nor would we say that my asking her to tell me the phone number will *enable* my getting her to tell it to me. What we would say is that I believe that *by* uttering the question I will be asking my office mate to tell me the phone number; that *by* asking her I will be getting her to tell it to me; and that *by* getting her to tell it to me, I will

be finding it out. Notice that the *by*-locution does *not* completely correlate with causation: while it is true that asking to be told the phone number seems to relate causally to getting a listener to tell you the phone number, there is no such causal flavor to the relation between uttering the question and asking to be told the phone number.

The claim is that if an agent believes that by doing α he will be doing either his goal act β or some other act γ that plays a role in her plan, then α may play a role in his plan. Of course, to some extent this just begs the question, since I have left vague the conditions under which one act can be said to be done *by* doing another. Goldman's concept of *generation* [13] can be used to make these conditions more precise [19, pp. 52-67]. I will adopt the term *generation* to describe the relation between two acts that we commonly express with the *by* locution. What is important to notice here is that there does seem to be an intuitive difference between the relation between acts that we describe as enablement and the one I am now calling generation. Just as we would not describe the relations between the acts discussed in the previous paragraph as enablement, so we would not, in general, use *by* to describe the relations earlier discussed as examples of enablement. We would not, for example, say that I believe that I can call St. Eligius by finding out the phone number there.¹²

¹²This claim may be less clear for the acts of asking Kathy how she is feeling and calling St. Eligius: it seems possible to say that I plan to ask Kathy how she is feeling by calling

What does the difference between enablement and generation consist in? Most importantly, it is the case that when one action α generates another action β , then the agent need only do α and β will automatically be done also. However, when α enables β , then the agent needs to do something more than α to guarantee that β will be done. I cannot simply find out the phone number of St. Eligius and then rationally expect that I will have called St. Eligius. But if I utter the words “Do you know the phone number of St. Eligius?” in the appropriate circumstances, then I need do nothing more to have asked my office mate the phone number of St. Eligius. And having done that, if the circumstances are right (e.g. my office mate knows the phone number of St. Eligius and is willing to tell it to me), then I need do nothing more to have caused my office mate to tell me St. Eligius’s phone number.

Of course, if the circumstances are not right, then by asking my office mate for the phone number, I will not have caused her to tell it to me. For instance, she might not hear me ask her. In that case, I may repeat my question. Now it is important not to confuse things here. It is true that in this case, I *do* “do something more” than my original act of asking for the phone number in order to perform the act of causing my office mate to tell it to me. But here my first act—of asking for the phone number—does St. Eligius. The ordinary-language test is a rough one, and occasionally fails to correlate with the phenomenal distinctions I want to draw.

not generate my act of causing my office mate to tell it to me. My second act—repeating my request—might do so, if my office mate hears me this time and responds, but then I have not “done anything more” than this second act, of repeating my question, to cause my office mate to tell it to me.

Analogously, I might dial the phone number and get a busy signal. I then need to do something more to establish a communication channel to Kathy. I might dial the number again, or I might drive to St. Eligius. In either case, if I succeed in establishing a communication channel, I succeed in doing so by dialing the second time, or by going to St. Eligius: it is my act of *redialing* or my act of going to St. Eligius, and not my original act of dialing, that generates my act of establishing the communication channel.

3.3 The Intention Component of Plans

So far we have seen that for an agent A to have a plan to do β that consists of doing Π , he must have a certain set of beliefs about the acts that Π comprises. Specifically,

- (1) A must believe that executing the acts in Π , in their (possibly partial) temporal order, will entail his performance of β

and

(2) A must believe that each act α in Π plays a role in his plan, i.e. either he believes that by doing α he will do β or some other γ that plays a role in his plan (i.e. that α generates β or γ); or he believes that doing α will enable doing β or some other γ that plays a role in his plan.

Condition (1) can be thought of as a sufficiency condition: it guarantees that if A believes that α is part of what he will do to achieve his goal, then α is in the plan. Condition (2) can be thought of as a necessity condition: it guarantees that if α is part of A 's plan, he believes it will play a role in achieving his goal.

Although these beliefs are necessary, they are not sufficient to guarantee that doing Π is A 's plan to do β . It is also necessary that A have a certain set of intentions with respect to Π . In particular, for my plan to do β to consist in my doing Π , I must intend to execute each of the acts in Π . Let Π once more be finding out the phone number of St. Eligius, calling St. Eligius, and saying "How are you doing?". I may believe that executing Π would entail my asking Kathy how she is feeling, but if I intend instead to wait until next Thursday and then talk to her face to face, Π is not a plan I have to ask Kathy how she is feeling.

Further, in order for doing Π to count as my plan to β , not only must I intend to execute Π , but I must also intend it *as a way* of doing β ; that

is, I must intend to do the acts in Π in order to do β . Imagine that I am a teenager whose parents have forbidden me to stay out past midnight. Imagine further that there is some club, “The Mayday Late-Night Club,” whose membership is limited to those who show up at the movies at 1 a.m. on May the first. Now let Π be the singleton set going to the movies at 1 a.m. on May the first. I may intend to execute Π , and I may believe that doing so will entail both becoming a member of the Mayday Late-Night Club and aggravating my parents. But if I intend my execution of Π as a way to do the former—and *not the latter*—then doing Π will count as a plan I have to join the Late-Night Club, but not as a plan I have to aggravate my parents. Notice that this is true even if I intend to aggravate my parent in some other way, say by getting a Mohawk haircut.

An argument similar to that made in the belief case also applies here. Thus, not only must I intend to execute Π in order to do β , I must also intend each act α in Π to play a role in my doing β ; i.e., I must intend each α either to generate or enable β or some other γ that itself plays a role in my plan. Thus we can state three more conditions on A ’s having a plan to do β that consists in doing Π , namely:

(3) A must intend to execute each act α in Π , in the (possibly partial) temporal order.

(4) A must intend to execute Π as a way of doing β .

and

(5) A must intend each act α in Π to play a role in his plan, i.e. either he must intend by doing α to do β or some other γ that plays a role in his plan; or he must intend by doing α to enable doing β or some other γ that plays a role in his plan.

The close parallel between Conditions (2) and (5) should lead us to ask whether one subsumes the other. The discussion so far has shown that having the beliefs defined in Condition (2) does not entail having the intentions in Condition (5). But does having the intentions in Condition (5) entail the beliefs in Condition (2)?

The question of whether an agent's intention to do α entails a belief that the agent will do α has been debated in the philosophical literature, and no consensus seems to have been reached. Where Grice [14] for one, answers in the affirmative, Davidson [9] goes to great pains to provide a counterargument.¹³ A telling comment on the controversy is provided by Bratman [5] who notes that "plans normally support expectations of their

¹³Davidson's case rests upon examples such as the following: I might intend to make ten legible copies of what I am writing by pressing hard on carbon paper, without believing with any confidence that I will succeed. Grice maintains that many such examples are actually elliptical versions of conditional intentions. Contra Grice, Davidson argues that such an intention is *not* an elliptical version of a conditional intention to make the ten copies if I can, for since one cannot intend to do what is impossible, intending to do X if one can is equivalent to intending to do X *simpliciter*; nor is it an elliptical version of

successful execution. . . , [although] there may still be cases in which I plan to A but do not believe I will” (p. 286, footnote 4). Within this paper, it will prove to be sufficient to assume this “normal” state of affairs and accept the view that if one intends to X, one must believe that one will.

Given this, we can see that Condition (5) directly entails Condition (2). For if, for each α in Π , A intends α to play a role in his plan, then A also believes that α will play a role in his plan. If A intends to do β or some other γ by doing α , he also believes that he will do β or this other γ by doing α ; if he intends to enable β or some other γ by doing α , he also believes that he will enable β or this other γ by doing α .

Similarly, Condition (4) entails Condition (1). For if A intends to do Π as a way of doing β , then A must believe he will do β by doing Π , and this is exactly Condition (1).

To state the commonsense requirements on A ’s having a plan, it is thus sufficient to state the conditions on intending—Conditions (3) through (5). However, while the conditions on belief—Conditions (1) and (2)—are entailed by Conditions (4) and (5) respectively, and are thus redundant, it some more detailed conditional intention to do X if, e.g., the carbon paper is particularly good, my hand muscles are more powerful than I thought, etc. His argument against this is that “there can be no finite list of things we think might prevent us from doing what we intend, or of circumstances that might cause us to stay our hand” [14, p. 94]. This is obviously a description of the notorious “frame problem” that plagues AI.

is worth keeping them in mind, for it will turn out that when an inferring agent deems an actor's plan invalid, it is because she believes it includes invalid beliefs. In fact, it is even worthwhile to make explicit yet another belief that is requisite to having a plan. Condition (3) above asserts that A intends to execute each act α in Π ; this then entails that A believes he will do each act α in Π . And this belief, in turn, entails a belief that A can do each act α in Π . Adding this condition to the definition, we can summarize the analysis of "having a plan" as follows:

(P0) An agent A has a plan to do β , that consists in doing some set of acts Π , provided that

1. A believes that he can execute each act in Π .
2. A believes that executing the acts in Π will entail the performance of β .
3. A believes that each act in Π plays a role in his plan. (See discussion below.)
4. A intends to execute each act in Π .
5. A intends to execute Π as a way of doing β .
6. A intends each act in Π to play a role in his plan.

4 Plans in Cooperative Communication

We can now consider how the model of plans just developed can be incorporated in a theory of communication that does not suffer from the shortcomings described in Section 2—that is, one in which even invalid plans can be reasoned about properly. To begin, it is useful to develop a more formal representation of the analysis given in (P0).

4.1 A Representation for Simple Plans

I will restrict my attention to a subset of plans, which I call *simple plans*. An agent has a simple plan if and only if he believes that all the acts in that plan play a role in it by generating another act; i.e., if it includes no acts that he believes are related to one another by enablement. The representation language I will use builds upon Allen’s interval-based temporal logic [3], a typed first-order predicate calculus. In particular I will make use of the predicates HOLDS and OCCURS. HOLDS is a binary relation over propositions and time intervals: $\text{HOLDS}(P,t)$ is taken to be true if and only if proposition P holds throughout time interval t . OCCURS, as I will use it, is a ternary relation: $\text{OCCURS}(\alpha,A,t)$ is true if and only if agent A performs an act of type α during time interval t .

Note that the first argument to OCCURS is an act-type. The distinction between act-types and actions is crucial. Actions or acts—I will use the two

terms interchangeably—can be thought of as triples of act-type, agent, and time. Thus typing **DEL** . is an example of an act-type, while my typing **DEL** . right now is an example of an action. Generation is a relation over actions, not over act-types. Sometimes an act of typing **DEL** . will generate an act of deleting the current mail message (say, when the former act is performed while the agent is using a particular electronic mail system). But not every case of an agent typing **DEL** . will result in the agent deleting the current message; for example, my typing it just now did not, because I was not typing it to a computer mail system. However, when an act of A 's doing α at time t generates an act of his doing β at time t , there are certain conditions C such that any time there occurs an act of α while C holds, there will also occur a simultaneous act of β . The regularity of the generation-enabling conditions C is what enables us to reason about whether by doing some action we will do another action, and consequently what enables us to construct and reason about (simple) plans.

To encode these regularities, I will introduce the following abbreviatory device: I will say that act-type α *conditionally generates* act-type β under conditions C , and will write $CGEN(\alpha, \beta, C)$, where C denotes the generation-enabling conditions relating α and β . Thus the $CGEN$ predicate is defined as follows:

$$(C1) \quad CGEN(\alpha, \beta, C) \leftrightarrow$$

- (i) $\forall A \forall t_1 [[HOLDS(C, t_1) \wedge OCCURS(\alpha, A, t_1)] \rightarrow OCCURS(\beta, A, t_1)] \wedge$
- (ii) $\exists A \exists t_2 [OCCURS(\alpha, A, t_2) \wedge \neg OCCURS(\beta, A, t_2)] \wedge$
- (iii) $\exists t_3 \forall A [HOLDS(C, t_3) \wedge \neg OCCURS(\beta, A, t_3)]$.

It is then straightforward to define the generates relation, GEN, in terms of CGEN:¹⁴

$$(G1) \text{ GEN}(\alpha, \beta, A, t) \leftrightarrow \exists C [CGEN(\alpha, \beta, C) \wedge HOLDS(C, t)].$$

That is, agent A 's doing α at time t will generate his doing β if and only if there are some generation-enabling conditions relating α and β , and further, those conditions hold at time t .

In translating Definition (P0) into the representation language, I will also make use of several other relations, which I will treat as primitive in this paper. (See, however, [19, pp. 67-72].) The relation $EXEC(\alpha, A, t)$ will be taken to be true if and only if the act of A doing α during time interval t is *executable*. The relation $BEL(A, P, t)$ is true if and only if agent A believes proposition P throughout time interval t ; $INT(A, \alpha, t_2, t_1)$ is true if and only if throughout time t_1 , A intends to do α at time t_2 . I also make use of one function: *by* maps two act-types into a third, composite act-type— $by(\alpha, \beta)$ denotes the act of doing β by doing α .

¹⁴Throughout, all variables should be taken to be universally quantified with widest possible scope, unless otherwise noted.

These various components of the representation language then can be combined in Definition (P1), which encodes in the representation language the definition of having a simple plan:

(P1) $SIMPLE - PLAN(A, \alpha_n, [\alpha_1, \dots, \alpha_{n-1}], t_2, t_1) \leftrightarrow$

- (i) $BEL(A, EXEC(\alpha_i, A, t_2), t_1)$, for $i = 1, \dots, n \wedge$
- (ii) $BEL(A, GEN(\alpha_i, \alpha_{i+1}, A, t_2), t_1)$, for $i = 1, \dots, n-1 \wedge$
- (iii) $INT(A, \alpha_i, t_2, t_1)$, for $i = 1, \dots, n \wedge$
- (iv) $INT(A, by(\alpha_i, \alpha_{i+1}), t_2, t_1)$, for $i = 1, \dots, n-1$.

The left-hand side of (P1) denotes that the agent A has, at time t_1 , a simple plan to do α_n , consisting of doing the set of acts $\{\alpha_1, \dots, \alpha_{n-1}\}$ at t_2 . Note that all these are simultaneous acts; this is a consequence of the restriction to simple plans. The right-hand side of (P1) corresponds directly to (P0), except that, in keeping with the restriction to simple plans, specific assertions about each act generating another replace the more general statement regarding the fact that each act plays a role in the plan. Clause (i) of (P1) captures clause (1) of (P0).¹⁵ Clause (ii) of (P1) captures both clauses (2)

¹⁵In fact, it captures more: to encode Clause (i) of (P0), the parameter i in Clause (i) of (P1) need only vary between 1 and $n-1$. However, given the following relationship between EXEC and GEN

$$EXEC(\alpha, A, t) \wedge GEN(\alpha, \beta, A, t) \rightarrow EXEC(\beta, A, t)$$

the instance of Clause (i) of (P1) with $i=n$ is a consequence of the instance of Clause (i)

and (3) of (P0): when i takes the value $n-1$, clause (ii) of (P1) captures the requirement, stated in clause (2) of (P0), that A believes his acts will entail his goal; when i takes values between 1 and $n-2$, it captures the requirement of clause (3) of (P0), that A believes each of his acts plays a role in his plan. Similarly, clause (iii) of (P1) captures clause (4) of (P0), and clause (iv) of (P1) captures clauses (5) and (6) of (P0).

4.2 Invalid Plans and Cooperative Responses

Given Definition (P1) it is straightforward to state what it means for an agent to have an invalid simple plan: A has an invalid simple plan if and only if he has the set of beliefs and intentions listed in Definition (P1), where one or more of those beliefs is incorrect, and, consequently, one or more of the intentions is unrealizable. We can also see what it means for one agent to believe that another has a simple plan: I will say that S believes that A has some (simple) plan to do the action β by doing the actions $\alpha_1, \dots, \alpha_n$ if S believes that A has the configuration of beliefs and intentions represented in (P1). And these two statements can be combined, so that S will be said to believe that A has a simple plan that is invalid if S believes to be false some belief she attributes to A by virtue of attributing to her some simple plan (and consequently, S also believes to be unrealizable some intention she

with $i=n-1$ and the instance of Clause (ii) with $i=n-1$. A similar argument can be made about Clause (iii).

attributes to *A* in virtue of attributing to her that same simple plan).

The structure of Definition (P1) suggests that there are two types of plan invalidities, corresponding to the two types of beliefs that are part of the mental attitude of having a simple plan. An incorrect belief corresponding to Clause (i) of (P1) will indicate a plan that includes an intention to do an *unexecutable act*, while an incorrect belief corresponding to Clause (ii) will indicate what I will call an *ill-formed* plan: one in which the intended acts will not lead to the goal. Of course, in accounting for cooperative conversation, what is at issue is not the absolute correctness of the actor's beliefs, but, as explained above, the inferring agent's beliefs about the beliefs she attributes to the actor. So, I will say that *S* judges *A*'s plan to contain an unexecutable act if *S* believes to be false one of the beliefs she attributes to *A* in satisfaction of Clause (i) of (P1). Likewise, I will say that *S* judges *A*'s plan to be ill-formed if she believes to be false one of the beliefs she attributes to *A* in satisfaction of Clause (ii) of (P1). Of course this is not to say that any agent *S* would assent to a description of *A*'s plan as "ill-formed": to say that an agent believes a plan is ill-formed is to describe a configuration of beliefs that agent has.

Consider again the dialogue presented in Section 2:

A: "I want to talk to Kathy, so I need to find out the phone number of St. Eligius."

$$\begin{aligned}
& \text{BEL}(S, \text{BEL}(A, \text{EXEC}(\text{call}(\text{St. Eligius}), A, t_2), t_1), t_1) \\
& \text{BEL}(S, \text{BEL}(A, \text{EXEC}(\text{establish-channel}(\text{Kathy}), A, t_2), t_1), t_1) \\
& \text{BEL}(S, \text{BEL}(A, \text{GEN}(\text{call}(\text{St. Eligius}), \\
& \qquad \qquad \qquad \text{establish-channel}(\text{Kathy}), \\
& \qquad \qquad \qquad A, t_2), t_1), t_1) \\
& \text{BEL}(S, \text{INT}(A, \text{call}(\text{St. Eligius}), t_2, t_1), t_1) \\
& \text{BEL}(S, \text{INT}(A, \text{establish-channel}(\text{Kathy}), t_2, t_1), t_1) \\
& \text{BEL}(S, \text{INT}(A, \text{by}(\text{call}(\text{St. Eligius}), \\
& \qquad \qquad \qquad \text{establish-channel}(\text{Kathy}), \\
& \qquad \qquad \qquad t_2, t_1), t_1)
\end{aligned}$$

Figure 2: The Plan S Infers for A

S : “St. Eligius closed last month. Kathy was at Boston General,
but she’s already been discharged. You can call her at home.
Her number is 555-1238.”

We can account for S ’s response there by assuming that she believes that A has a certain set of beliefs and intentions satisfying (P1). A portion of those beliefs is shown in Figure 2. S believes that A believes that calling St. Eligius at time t_2 is executable, though S believes it is not executable,

and she informs *A* of this in her response. Also, *S* believes that *A* believes that the act of calling St. Eligius will generate the act of establishing a communication channel to Kathy. *S* believes that it will not—that even if *A* could call St. Eligius, that act would not have the desired effect. This belief also affects *S*'s response.

Strategies for producing cooperative responses to questions must incorporate decisions about what information to include. The view of plans developed here—in which plans are seen as complex mental attitudes, and plan inference is seen as the process of attributing such attitudes to an actor—provides the basis for determining a class of information that may need to be included in a cooperative response to a question. Specifically, to be cooperative a response may need to include information about the particular discrepancies *S* finds between her own beliefs and those she attributes to *A* as part of her belief that he has some particular plan. Note, however, that a cooperative response may not necessarily include all such information, for *S* may deem some or all of it to be irrelevant. The plan inferred to underlie a query and any invalidities it is judged to have are but two factors affecting the response generation process, the most significant others being relevance and salience.

4.3 Explanatory Plans

When S judges A 's plan to be invalid in one of the ways discussed above, she has intuitively “made sense” of the plan and understands the source of the invalidities. However, there are also case in which an inferring agent simply cannot make sense of an actor's query. As a somewhat whimsical example, imagine A saying:

A : “I want to talk to Kathy, so I need to find out how to stand on my head.”

In many contexts, a perfectly reasonable response to this query is “Huh?”. A 's query is *incoherent*: a listener S may be unable understand why A believes that finding out how to stand on his head (or standing on his head) will lead to talking with Kathy. One can, of course, construct scenarios in which A 's query makes perfect sense: Kathy might, for example, be currently hanging by her feet in gravity boots. The point here is not to imagine such circumstances in which A 's query would be coherent, but instead to realize that there are many circumstances in which it would not.

The model of plans as I have so far presented it does not distinguish between a query of this type and one in which the inferred underlying plan is ill-formed. The reason is that, given a reasonable account of semantic interpretation, it is transparent from the query just given above that A intends to talk to Kathy, intends to find out how to stand on his head, and

intends his doing the latter to play a role in his doing the former. Further, as consequences of these intentions A believes that he can talk to Kathy, believes that he can find out how to stand on his head, and believes that his doing the latter will play a role in his doing the former.¹⁶ But these beliefs and intentions are precisely what are required by Definition (P0) to have a plan; and if S could determine that the intended role of the supporting act of standing on his head was generation, then these beliefs and intentions would also be exactly what is required by Definition (P1). Consequently, after hearing the query, S can in fact infer a plan underlying A 's query, namely the obvious one: to find out how to stand on his head in order to talk to Kathy. Then since S does not herself believe that the former act will lead to the latter, on the analysis so far given, we would regard S as judging A 's plan to be ill-formed. Unfortunately, this is not the desired analysis: the model should instead capture the fact that S cannot make sense of A 's query here—that it is *incoherent*.

To capture the difference between ill-formedness and incoherence, I will claim that, when an agent S is asked a question by an actor A , S needs to attempt to ascribe to A more than just a set of beliefs and intentions satisfying (P1). Specifically, for each belief satisfying Clause (ii) of (P1), S must also ascribe to A another belief that explains the former in a certain

¹⁶Recall the discussion in Section 3.3 about assuming that an intention to α entails a belief that the agent will do α .

specifiable way. The beliefs that satisfy Clause (ii) are beliefs about the relation between two particular actions. For instance, the plan underlying the example query includes A 's belief that his action of calling St. Eligius at t_2 will generate his action of establishing a communication channel to Kathy at t_2 . This belief can be explained by a belief A has about the relation between the act-types "calling a location" and "establishing a communication channel to an agent." A may believe that acts of the former type generate acts of the latter type, provided that the agent to whom the communication channel is to be established is at the location to be called. Such a belief can be encoded using the CGEN relation introduced earlier. So, for instance, S may attribute to A a belief that we can express as follows:

$$BEL(A, CGEN(call(X), establish-channel(Y), at(X, Y)), t_1).$$

This belief, combined with a belief that Kathy will be at St. Eligius at time t_2 , explains A 's belief that, by calling St. Eligius at t_2 , he will establish a communication channel to Kathy. In contrast, S may have no basis for ascribing to A beliefs that will explain why he thinks that standing on his head will lead to talking with Kathy. Consequently she will deem the second example query to be incoherent.

Explanatory beliefs are incorporated in the PI model by the introduction of *explanatory plans*, or *eplans*. Saying that an agent S believes that another agent A has some eplan is shorthand for describing a set of beliefs possessed

by S , specifically:

(P2) $BEL(S, EPLAN(A, \alpha_n, [\alpha_1, \dots, \alpha_{n-1}], [\rho_1, \dots, \rho_{n-1}], t_2, t_1), t_1) \leftrightarrow$

(i) $BEL(S, BEL(A, EXEC(\alpha_i, A, t_2), t_1), t_1)$, for $i = 1, \dots, n \wedge$

(ii) $BEL(S, BEL(A, GEN(\alpha_i, \alpha_{i+1}, A, t_2), t_1), t_1)$, for $i = 1, \dots, n-1 \wedge$

(iii) $BEL(S, INT(A, \alpha_i, t_2, t_1), t_1)$, for $i = 1, \dots, n \wedge$

(iv) $BEL(S, INT(A, by(\alpha_i, \alpha_{i+1}), t_2, t_1), t_1)$, for $i = 1, \dots, n-1 \wedge$

(v) $BEL(S, BEL(A, \rho_i, t_1), t_1)$,

where each ρ_i is $CGEN(\alpha_i, \alpha_{i+1}, C_i) \wedge HOLDS(C_i, t_2)$.

Clauses (i) - (iv) of (P2) are identical to clauses (i) - (iv) of (P1). The key clause in the definition of eplans is Clause (v): for S to believe that A has some eplan, she must attribute to A beliefs that explain the other beliefs that are constituents in his plan. I will call the beliefs that S attributes to A in satisfaction of Clause (v) *explanatory beliefs*.

It is now possible to sketch the process of inferring the eplan that underlies a query. To begin, S will believe that A intends to do some act α , possibly because A tells her this in the query. S thus may believe that A has a trivial eplan, i.e.

$$BEL(S, EPLAN(A, \alpha, [], \rho, t_2, t_1), t_1)$$

where ρ is nil (logically true)

Let us suppose then that *S* has reason to believe—on the basis of something other than *A*’s query itself—that it is plausible that *A* believes that act-type α conditionally generates some other act-type γ , under some specific condition *C*, and that she has no reason to suppose that *A* believes that *C* will not hold at the intended performance time of his plan. Then *S* can decide that it is plausible for *A* to believe that by his act of α , he will do γ , and further, that it is plausible for him to intend to do α in order to do γ , and to intend to do γ . The plausibility of *A*’s having these intentions depends upon the plausibility of his having the aforementioned beliefs, so *S* will attribute to *A* as a bundle the plausible intentions and supporting beliefs. That is, *S* will reason from the trivial eplan to a larger one. The process of belief and intention ascription can be iterated: believing that it is plausible that *A* intends to do γ , *S* can then reason about what *A* might believe he can do by this act. *S* can also reason about what *A* might plausibly intend to do in order to do his goal act β , and can then iterate as well in the “backward” direction. Again, the reasoning is from one plausible eplan to another.

What is crucial in this picture is the way in which *S* attributes to *A* particular explanatory beliefs. As noted, when she decides that it is plausible for him to believe that by doing α he will be doing some γ , the plausibility of that belief must be established on the basis of something other than *A*’s query itself; otherwise, the very distinction that eplans were set up to

support—i.e., the distinction between incoherent and ill-formed plans—will collapse. While it is true that if A says “I want to talk to Kathy, so I need to find out how to stand on my head,” S can figure out that A believes that by standing on his head he can talk to Kathy (or at least enable his talking to Kathy), this belief can only be attributed to A on the basis of his query. The types of evidence that can be used to attribute explanatory beliefs to an agent are encoded in a set of *plan inference rules* that describe the reasoning that is permitted in going from one plausible eplan to another. If the plan inference process proceeds only through the use of plan inference rules, the distinction between ill-formed plans and incoherent queries is maintained.

In the simplest plan inference rule, S attributes to A a belief that she herself has:¹⁷

$$\begin{aligned}
 & \textbf{(PI1)} \quad \text{BEL}(S, \text{EPLAN}(A, \alpha_n, [\alpha_1, \dots, \alpha_{n-1}], [\rho_1, \dots, \rho_{n-1}], t_2, t_1), t_1) \wedge \\
 & \quad \text{BEL}(S, \text{CGEN}(\alpha_n, \gamma, C), t_1) \\
 & \quad \rightarrow \\
 & \quad \text{BEL}(S, \text{EPLAN}(A, \gamma, [\alpha_1, \dots, \alpha_n], [\rho_1, \dots, \rho_n], t_2, t_1), t_1) \\
 & \quad \text{where } \rho_n = \text{CGEN}(\alpha_n, \gamma, C) \wedge \text{HOLDS}(C, t_2).
 \end{aligned}$$

¹⁷The plan inference rules should really be stated in terms of plausible eplans, i.e., collections of beliefs and intentions that S thinks A plausibly has. When S has found some set of these that is large enough to account for A ’s query, their epistemic status can be upgraded to beliefs and intentions that S will, for the purposes of forming her response, consider A actually to have. See [19, pp. 126-130].

This rule says that, if S 's belief that A has some eplan includes a belief that A intends to do an act α_n , and S also believes that act-type α_n conditionally generates some γ under condition C , then S may infer that A has the additional intention of doing α_n in order to do γ —i.e., that he intends to do $by(\alpha_n, \gamma)$. A 's having this intention depends upon his also having the supporting belief that α_n conditionally generates γ under some condition C , and the further belief that this C will hold at performance time.¹⁸

One way to view Rule (PI1) is as an explication of Condition (3) of Figure 2 in the language of plans as mental phenomena. If a system made use only of Rule (PI1) (and its symmetric partner), it would be implicitly committing to the assumption that A has the same beliefs about conditional generation as it does. Additional plan inference rules are needed to move beyond this assumption. In one obvious form of reasoning, S may attribute to A beliefs that are slight variations of her own, for example as encoded in (PI2):

$$\begin{array}{l}
 \textbf{(PI2)} \quad \text{BEL}(\text{R}, \text{EPLAN}(\text{Q}, \alpha_n, [\alpha_1, \dots, \alpha_{n-1}], [\rho_1, \dots, \rho_{n-1}], t_2, t_1), t_1) \wedge \\
 \quad \text{BEL}(\text{R}, \text{CGEN}(\alpha_n, \gamma, C_1 \wedge \dots \wedge C_m), t_1) \\
 \quad \rightarrow \\
 \quad \text{BEL}(\text{R}, \text{EPLAN}(\text{Q}, \gamma, [\alpha_1, \dots, \alpha_n], [\rho_1, \dots, \rho_n], t_2, t_1), t_1)
 \end{array}$$

¹⁸A rule symmetric to (PI1) is also needed since S not only can reason about what acts might be generated by an act that she already believes A plausibly intends, but also about what acts might generate such an act.

$$\text{where } \rho_n = CGEN(\alpha_n, \gamma, C_1 \wedge \dots \wedge C_{i-1} \wedge C_{i+1} \wedge \dots \wedge C_m) \wedge \\ HOLDS(C_1 \wedge \dots \wedge C_{i-1} \wedge C_{i+1} \wedge \dots \wedge C_m, t_2)$$

What (PI2) expresses is that R may ascribe to Q a belief about a relation between act-types that is a slight variation of one she herself has. It asserts that, if there is some CGEN relation that R believes true, she may attribute to Q a belief in a similar CGEN relation that is stronger, in that it is missing one of the required conditions. As another example, if S believes that two act-types α and β are quite similar, she thereby has reason to believe that it is plausible that A has confused them, or has made a bad analogy from one to the other. Such reasoning is encoded in Rule (PI3):

$$\begin{aligned} \text{(PI3)} \quad & BEL(S, EPLAN(A, \alpha_n, [\alpha_1, \dots, \alpha_{n-1}], [\rho_1, \dots, \rho_{n-1}], t_2, t_1), t_1) \wedge \\ & BEL(S, SIMILAR(\alpha_n, \delta), t_1) \wedge \\ & BEL(S, CGEN(\delta, \gamma, C), t_1) \\ & \rightarrow \\ & BEL(S, EPLAN(A, \gamma, [\alpha_1, \dots, \alpha_n], [\rho_1, \dots, \rho_n], t_2, t_1), t_1) \\ & \text{where } \rho_n = CGEN(\alpha_n, \gamma, C) \wedge HOLDS(C, t_2) \end{aligned}$$

(PI1), (PI2), and (PI3) are merely meant to be suggestive of the sort of plan inference rules that can be stated within the analysis of plans as complex mental attitudes; a number of other such rules, along with examples of their use, can be found in [19, Chap. 6]. Such rules can enable a system to infer a wide range of plans, including plans that are constructed out

of recipes-for-action that are not stored in the system's operator library. Adopting a view of plans as mental phenomena also makes it possible to reason about whether a particular set of actions is likely actually to be the object of an agent's plan. Consider once more the sample dialogue presented at the beginning of Section 2. Imagine, however, that when the query is asked, *S* has reason to believe that *A* believes that Kathy is at home. In these circumstances *S* cannot use (PI1) to reason to the set of beliefs shown in Figure 2; indeed in this case *A*'s query may be deemed to be incoherent. To distinguish between these two cases it is necessary to have both a concern with the beliefs and intentions that are entailed by having a particular plan and a framework that distinguishes between the beliefs of the inferring agent and those that she attributes to the actor.

5 Reconsidering the Traditional Model

As I have emphasized throughout this paper, the traditional models in AI are models of the object of a plan, while my concern has been with the mental state of having a plan. But of course, when an agent has a plan, that plan does have an object—to wit, the set of acts he intends. What is the correspondance between the structure imputed to the actions that constitute a plan under the traditional approach and the structure imputed to the actions that are intended by an agent that has a plan under my approach? More

specifically, what correspondances, if any, are there between the relations between acts used in the traditional models—*causes*, *is-a-precondition-of*, and *is-a-way-to*—and those used in the account of plans I have developed here—*generation* and *enablement*?

The relations used in the traditional view involve some redundancy, as illustrated by the following two operators:

Header: flip switch

Effect: light on

Header: turn on light

Body: flip switch

These operators encode the same information, provided we equate the act of turning on the light with the act of achieving that the light is on. In general, whenever some act α *causes* some proposition P, it is also true that α *is-a-way-to* achieve(P). It will thus be sufficient to restrict our attention to the two relations *is-a-precondition-of* and *is-a-way-to*.

Consider then the following generic operator, in which α *is-a-way-to* β , and P *is-a-precondition-of* β :

Header: β

Preconditions: P

Body: α

Is there a way to map the relations expressed in it into the two relations *generation* and *enablement*? In fact, there are a number of plausible mappings.

Figure 3 shows some of these.

Sentence (1) Figure 3 encodes the interpretation, apparently implicit in certain operators, that the preconditions P are sufficient for the performance of the header β by performance of the body α , but without necessarily being sufficient for the performance of α itself (and consequently, without necessarily being sufficient for the performance of β itself). This is the interpretation that most straightforwardly translates into the relations between act-types used in this paper: under this interpretation, the entire generic operator is associated with the sentence $\text{CGEN}(\alpha, \beta, C)$. However, it is not always the interpretation implicit in existing planning work. Sentences (2a) and (2b) encode the interpretation of the preconditions P as necessary for the occurrence of β : these two sentences differ from one another in the time interval during which P is meant to hold. Consider a typical planning operator in which a proposition representing that the power is on is included as a precondition for an operator with a header representing the act of turning on the light. It is not always obvious whether the former is intended to be a necessary, or merely sufficient, condition for performance of the latter. (It might be merely sufficient, if, say, the light is attached to an emergency generator.) The preconditions in an action operator are sometimes also meant to be related to the performance of the act-type in the body of the operator—either to be necessary for it, as encoded in (3a) and (3b), or to be sufficient for it, as encoded in (4a) and (4b). One illustration of the preconditions being

Figure 3: Interpretation of the Standard Relations

related to the body act-type α instead of the header act-type β would be an operator with a header representing turning on the light, body representing flipping the switch, and precondition representing both the power being on and the agent standing near the switch. The agent may well be able to turn on the light without standing near the switch, say by throwing something at it. Standing near the switch is meant, in an example like this, to be necessary for the body action to occur. Combinations of the sentences shown in the figure are also possible: the interpretation underlying Sacerdoti’s work [20] seems to be a combination of (1) and (4a).

Unfortunately, much of the existing planning literature has been vague about the intended interpretation of action operators, and has used them at different times to mean different things. It is because of the resulting ambiguity of interpretation that I have avoided using the relations *causes*, *is-a-precondition-of*, and *effects* in this work, and have instead made use of *generates* and *enables*. Any particular operator that makes use of the former set of relations can be expressed in terms of the primitive relations OCCURS and HOLDS—the two relations that were themselves used to define generates and enables. The translation must proceed on an operator-by-operator basis, however; because of the variety of ways in which the traditional relations have been used, there is no one translation that will apply uniformly to all the operators that have been proposed in the literature.

6 Conclusion

For nearly two decades, research on cooperative conversation has entailed a concern with the notion of plans and with the process by which one agent can infer the plans of the other agents with whom she is conversing. In this paper I have argued that in understanding plan inference it is important to go beyond studying the structure of the object of an agent's plans: it is necessary to analyze the nature of the mental state of the plan itself. I presented such an analysis, in which having a plan amounts to having a particular configuration of beliefs and intentions. I then showed how this analysis can be put to use in a model of conversation that avoids certain limitations of any approach that does not consider the nature of plans as complex mental attitudes. In particular, I showed how an analysis of plans as complex mental attitudes makes it possible to reason about the plans of an actor even when those plans are invalid: how to reason about whether it is likely that the actor has an invalid plan, how to reason about the ways in which the plan may be invalid, and how to reason about what information to consider including in a cooperative response to a query that arises from an invalid plan.

When plans are viewed as complex mental attitudes, the process of plan inference can be seen as that of attributing a collection of beliefs and intentions to an actor. Under this analysis there can be discrepancies between

an agent's own beliefs and the beliefs that she ascribes to an actor when she thinks he has some plan. I associated such discrepancies with a judgement that the actor's plan is invalid. Then I showed that the types of any invalidities judged to be present in a plan inferred to underlie a query affect the content of a cooperative response. I further suggested that, to guarantee a cooperative response, an inferring agent must attempt to ascribe to a questioner more than just a set of beliefs and intentions sufficient to believe that he has some plan: she must also attempt to ascribe to him beliefs that explain those beliefs and intentions. The *eplan* construct was introduced to capture this requirement. I described the process of inferring eplans—that is, of ascribing to another agent beliefs and intentions that explain his query and can influence a response to it. Finally, I compared the representation used in traditional AI models of plans with the representation used here for the objects of an agent's plans, showing that the latter can be used to express and make more precise the intended meaning of uses of the former.

References

- [1] James F. Allen. *A Plan Based Approach to Speech Act Recognition*. Technical Report TR 121/79, University of Toronto, 1979.
- [2] James F. Allen. Recognizing intentions from natural language utterances. In Michael Brady and Robert C. Berwick, editors, *Computational*

Models of Discourse, pages 107–166, MIT Press, Cambridge, Ma., 1983.

- [3] James F. Allen. Towards a general theory of action and time. *Artificial Intelligence*, 23(2):123–154, 1984.
- [4] Michael E. Bratman. *Intention, Plans and Practical Reason*. Harvard University Press, Cambridge, Ma., 1987.
- [5] Michael E. Bratman. Taking plans seriously. *Social Theory and Practice*, 9:271–287, 1983.
- [6] John Seely Brown and Richard R. Burton. Diagnostic models for procedural bugs in basic mathematical skills. *Cognitive Science*, 2:155–192, 1978.
- [7] M. Sandra Carberry. *Pragmatic Modeling in Information System Interfaces*. PhD thesis, University of Delaware, 1985.
- [8] Allan Collins, Albert Stevens, and Sarah E. Goldin. Misconceptions in students’ understanding. *International Journal of Man-Machine Studies*, 11:145–146, 1979.
- [9] Donald Davidson. Intending. In Donald Davidson, editor, *Essays on Actions and Events*, Oxford University Press, New York, 1980.

- [10] R. E. Fikes and Nils J. Nilsson. Strips: a new approach to the application of theorem proving to problem solving. *Artificial Intelligence*, 2:189–208, 1971.
- [11] Michael R. Genesereth. The role of plans in automated consultation. In *Proceedings of the 6th International Joint Conference on Artificial Intelligence*, pages 311–319, Tokyo, 1979.
- [12] Michael P. Georgeff and Amy L. Lansky. Procedural knowledge. *Proceedings of the IEEE, Special Issue on Knowledge Representation*, 1383–1398, 1986.
- [13] Alvin I. Goldman. *A Theory of Human Action*. Prentice-Hall, Englewood Cliffs, N.J., 1970.
- [14] H.P. Grice. Intention and uncertainty. *Proceedings of the British Academy*, 57:263–179, 1971.
- [15] Henry A. Kautz. *Toward a Theory of Plan Recognition*. Technical Report TR162, University of Rochester, 1985.
- [16] Henry A. Kautz and James F. Allen. Generalized plan recognition. In *AAAI-86, Proceedings of the National Conference on Artificial Intelligence*, Philadelphia, Pa., 1986.

- [17] Diane Litman. *Plan Recognition and Discourse Analysis: An Integrated Approach for Understanding Dialogues*. PhD thesis, University of Rochester, 1985.
- [18] Nils J. Nilsson. *Principles of Artificial Intelligence*. Tioga Publishing Co., Palo Alto, Ca., 1980.
- [19] Martha E. Pollack. *Inferring Domain Plans in Question-Answering*. Technical Report 403, SRI International, Menlo Park, Ca., 1986. Also appears as a University of Pennsylvania PhD thesis.
- [20] Earl D. Sacerdoti. *A Structure for Plans and Behavior*. American Elsevier, New York, 1977.
- [21] Candace L. Sidner. Plan parsing for intended response recognition in discourse. *Computational Intelligence*, 1(1), 1985.
- [22] Beverly Woolf and David McDonald. Human-computer discourse in the design of a pascal tutor. In *Proceedings of the CHI'83 Conference on Human Factors in Computer Science*, 1983.