

A PARTIAL TAXONOMY OF KNOWLEDGE ABOUT ACTIONS

Eugene Charniak

Institute for the Study of Semantics and Cognition
6976, Castagnola, Switzerland

Abstract

An incomplete taxonomy is proposed for the knowledge of actions a computer must have if it is to understand stories about people performing these actions. The classification is along two dimensions, Force - why a person should follow the rule, and Form - what the rule looks like. Four kinds of force are distinguished, Strict, Social, Suggested, and Regulatory, while six forms are distinguished, Subactions, Substates, Side Conditions, Methods, Time Orderings, and Do-Whiles. For the most part it seems that any force type may combine with any form type in producing a rule. One example, how to use an umbrella, is examined in detail although facts about many human activities, from playing bridge to washing one's hair, are used as examples.

I Introduction

It is commonly accepted that a computer which answers questions about a natural language text must have roughly a human reader's knowledge of the subject matter. Restricting ourselves to simple stories (like children's stories) the knowledge needed is that "common sense" knowledge which all members of a particular culture share. Since, by all indications even this knowledge is quite extensive, it will have to be well organized and classified before a computer will be able to make much use of it. In this paper I will be concerned with classifying the kinds of facts we have about actions. Some typical facts of the sort I will be concerned with are:

- (F1) In giving a present the giver should give something the receiver would like to get.
- (F2) When painting some object it is a good idea to put paper on the floor under the object so as not to get paint on the floor.

Some facts which would not be included *in* this study, as they are intuitively not about actions, are:

- (F3) The Earth has one moon.
- (F4) Few people like modern music.

Furthermore, this study will be restricted to facts like (F1)-(F2) which tell us something about how to perform an action, as opposed to facts like:

- <F5) Many people break their legs while skiing.
- (F6) At playing chess, nobody is much better than Bobby Fisher.

Intuitively it seems to me that the "how to do it" type facts form a natural subclass of knowledge, so that by restricting the study to them I will not be hopelessly distorting the results. Naturally it would be preferable to have some firm evidence on this point, but short of already

having a better taxonomy, it is hard to imagine what the evidence could be. Since one has to start someplace, one starts by backing one's intuitions.

II What Should the Taxonomy Look Like?

To say that a particular fact is "about" some action is, at the very least, to say that the fact states some relation between the action and something else which might be another action, or perhaps some state. What I would like to find then is a list of relations such that for any action all the "how to do it" facts about it will fall under one or the other of the relations in the list. Furthermore I want the relations to be "interesting" in the sense that knowing the relation-type of a fact will give a language comprehension program information about what can be inferred from the fact, and perhaps even under what circumstances it should be used to make an inference.

To make all this concrete, consider the concepts "pre-requisite" and "result". It should be clear that "pre-requisite" and "result" qualify as relation-types as defined in the last paragraph. After all, when one says that an action A has a result state S, one means that after A is performed S will hold. So result is a relation between A and something else, in this case the state S. Also both relations are "interesting" since to some degree they characterize inferences which can be drawn from the facts. In the case of pre-requisite (henceforth abbreviated as "pr") we have inference rules like:

- (R1) If state S is a pr of action A, then S must hold just prior to A's being done.
- (R2) If animate being P desires action A, and S is a pr of A, then that desire would constitute a reason why P might desire S to hold.

Naturally, in an actual language comprehension system (R1) and (R2) would be expressed in an appropriate formalism. However, the points I wish to make in this paper are hopefully independent of particular formalisms, so I will use informal natural language expressions, both in the statement of facts, and in inferences which we might draw from them.

Of course, pr type facts and result facts hardly exhaust the kinds of facts we have about actions, but with one or two exceptions nobody has bothered to extend the list. And while this paper will not produce a complete list, it will expand the list quite a bit.

III Justifying a Putative Fact

Ideally, a paper of this sort should proceed along the following lines: first produce a puta-

tive piece of common sense knowledge, then show that language users in a particular culture {middle class American in this paper} indeed have this piece of knowledge, and finally show how it fits, or fails to fit, in the taxonomy as developed so far, and if the latter, extend the taxonomy so it does fit. In this paper, however, I will not show that the putative pieces of common sense knowledge are indeed real, but I will rather rely on the reader's intuitions. To show what such justification would look like, let me give an abbreviated version for the first example in the paper, repeated here as (F7).

(F7) In giving a present the giver should give something the receiver would like to get.

In general the justification goes in two parts. First we show that something like (F7) is needed. Consider:

Fred was thinking of getting a kite to give to Bob for his birthday. Then Fred learned that Bob had a kite and did not want another.

If we were to ask whether Fred will still buy a kite, the answer would be something like "Probably not". If asked why, the response would be "Because Bob does not want a kite". If we assume a fact like (F7) these responses are easy to account for. Indeed there seems to be no other option than to assume something on the order of (F7). But need the fact be <F7> itself? That is, perhaps there are more basic facts which combine to produce (F7), and what people know are these more basic facts rather than (F7). Showing that this is not the case is the second part of (F7)'s justification and it is in general quite difficult. All one can usually do is to examine a few possible alternatives, show them false, and then put the burden of proof on anyone who objects to the putative fact. So, for example, we can see that (F7) is a fact about present giving, rather than one about giving in general, since it would not be true for a policeman's giving somebody a parking ticket. I could go further and show that (F7) cannot even be generalized to giving "useful" things (unlike parking tickets), but will not do so here. (See (Charniak 74) for a deeper analysis of this particular fact.)

IV Strict versus Non-strict Pre-requisites

Given then that (F7) is part of our knowledge repertoire, and recognizing that it tells us something about how to go about giving a present, where should it go in our taxonomy? While it is clearly not a pr of present giving, it shares many properties of pr's. To see this, let us recast (F7) in the form of a pr.

<F8> Receiver liking present is a pr? of giving a present.

First, as (F8) indicates, (F7) describes a relation between the action of present giving and a particular state, namely "receiver wanting the present". Second, (F8) obeys rule (R2) which is a rule about pr's. That is, wanting an action to happen often leads to wanting the pr's to hold. So in the case of (F8) we have examples like:

Fred bought Bob a kite for his birthday. "I hope Bob wants a kite" thought Fred.

Of course, what prevents (F8) from really being a pr is that (R1) (pr's must hold for the action to take place) does not hold for (F8). That is, it is perfectly possible to give somebody a present he does not like. However, a rule very similar to (R1) does apply for (F8), namely:

(R3) If state S is a non-strict pr of action A, then it is a good idea for S to hold just prior to A's being done.

What we are doing then is to distinguish between strict pr's, which obey (R1), and non-strict pr's, which obey (R3). So (R1) would now read:

(R4) If state S is a strict pr of action A, then S must hold just prior to A's being done.

Another distinction between strict and non-strict pr's is:

(R5) If an action A is done although a strict pr is broken we ask "how" it was done that way. If a non-strict pr is broken we ask "why" it was done that way.

Some examples of strict pr's are:

(F9) Having an apple is a strict pr of eating it.

(F10) Not being at a location is a strict pr of going there.

(F11) There being a wind is a strict pr of flying a kite.

(F12) The umbrella restraining strap being unhooked is a strict pr of opening the umbrella.

Some other examples of non-strict pr's are:

(F13) Having permission is a non-strict pr of using an object which does not belong to you.

(F14) Being invited is a non-strict pr of attending a party.

V Logical versus Physical Pre-requisites

We will be spending most of our time looking at various kinds of non-strict pr's, but let us first take a brief look at the strict variety. In particular, there might be a case for saying that there are really two kinds of strict pr's, which we might call "logical" and "physical". To see the difference, consider the difference between (F9)-(F10) versus (F11)-(F12). If somebody told you that he went to Chicago last year, but that he had never, ever, been outside of Chicago, you would probably accuse him of talking riddles. One might say that it simply made no sense for the speaker to say that he had gone to Chicago, if he had never been outside the city. In some sense he had contradicted himself. On the other hand, there would certainly be no contradiction in terms if someone said that he had flown a kite without any wind, although you might not believe him, and ask skeptical questions about how he did it. That, roughly, would be the distinction between logical and physical pr's.

Now philosophers have said a lot about the

existence or non-existence of distinctions like this one. However, our problem in this paper is a different one. Simply, do ordinary people in everyday conversation make such a distinction? I myself tend to doubt it, but my opinions either way are not very strong, so I will list this distinction when I sum things up later, but put a question mark beside it.

VI Social versus Suggested Pre-requisites

So far we have only investigated one type of non-strict pr which, for reasons which will soon be apparent, we will call "social" pr. Another kind of non-strict pr is a "suggested" pr. For example:

- (F15) Having a tail on a kite is a suggested pr of flying it.
Reason: it makes flying the kite easier
- (F16) Having paper, or otherwise uncared about material around and if possible under an object is a suggested pr of painting that object. Reason: one is less likely to get paint on other things.
- (F17) Being sober is a suggested pr of driving.
Reason: less likely to have an accident this way.

One difference between social and suggested pr's is immediately obvious. Suggested pr's are backed up with reasons for the suggestion, while social pr's, like (F14) or (F8) do not have reasons - things are simply "done that way". Having a reason for a suggestion allows us to relate the breaking of a suggested pr to certain difficulties experienced *in* performing the action. For example:

- | | |
|---|----------------------------------|
| (a) Jack did not put paper under the chair before he painted it | (c) He got paint over everything |
| (b) Jack did not clean the chair before he painted it. | (d) The paint did not coat well. |

If we combine (a) with (c), or (b) with (d) the stories make sense, but the two other combinations do not make connected stories. Our comprehension model could account for such cases with a rule like:

- (R6) If a suggested pr F of an action A is broken, then associate the negation of the reason for F, if it occurs, with the breaking of the suggested pr.

Another difference between social and suggested pr's is that social pr's are much more subject to opinion than suggested pr's, and probably any representation of social pr's will have to allocate space for an indicator which shows which class of people believed in this particular pr and at which period in time. So some social pr's which at the current time divide people are:

- (F18) Having no foreign matter in the mouth is a social pr of talking.
- (F19) Wearing whites is a social pr for playing tennis.

Another distinction between the two is that social pr's have a "moral" component which we do not find in suggested pr's. For example:

Tom felt guilty using Bill's baseball without asking him first.

Mr. Jones was embarrassed playing tennis in his street clothes.

I am making the implicit claim here that embarrassment and guilt spring from breaking the same kinds of rules, and the distinction is due to the particular circumstances and the strength of the feeling. This would seem to be supported by examples like:

- (F20) Being correct is a social pr of giving somebody information.

When Mr. Jones realised that he had given the driver wrong directions he felt guilty/embarrassed.

In this last example either emotion would be reasonable, but presumably in both cases it would spring from the breaking of rule (F20). On the other hand there are clear cut cases where between the two (guilt and embarrassment) only one is appropriate. So whether one should distinguish between social and moral pr's or not is by no means clear. I have somewhat arbitrarily chosen the latter alternative.

Note that it is perfectly possible to feel guilty about breaking a suggested pr. For example, it is possible, one might even say preferable, for a person to feel guilty about driving while drunk. However, it is surely the case that this is not due simply to breaking the suggested pr, but rather that in breaking the suggested pr one is also breaking a social rule of the form "it is bad to put lives or property in danger, especially through carelessness or thoughtlessness". To see that it must be a rule like this which is responsible for guilt after breaking suggested pr's, note that it is hard to imagine feeling guilty over not putting a tail on one's kite. On the other hand, if for some strange reason, somebody's life depended on your getting the kite up quickly, it would be easy to imagine feeling guilty over forgetting to put on the tail.

Hence, another rule which distinguishes social from suggested pr's is:

- (R7) Breaking a social pr which one accepts is often a cause of guilt or embarrassment.

VII Pre-requisites and their Force

It is perhaps a good time to step back for a second and look again at what we have been doing. So far we have accumulated three kinds of pr's - social, suggested, and strict. In each case these facts mentioned some state which ought to be satisfied prior to an action taking place. What differed from type to type was why the state should be satisfied. To put this slightly differently, what differed was what would happen if we disobeyed, so to speak, the pre-requisite. What we have been classifying then is "pre-requisite

forco", that is, the "driving power" behind those various pr rules. As we will see later, these same kinds of forces apply equally well to other kinds of knowledge about actions, so really then we have been making a catalogue of "forces". That in each case we saw the force applied to a pr relation was simply for ease of presentation.

Sppn in this light, it is logical to think in terms of a single pr relation but with each rule "tagged" with a force. Furthermore, we might now ask if each pr is limited to one such "tag". The answer seems to be no. So a society may take some suggestion (in my use of the word) and make it into a social rule - making virtue out of necessity, as it were. For example

(F21) Having clean hands is a pr of eating.

Force: (Social)
(Suggestion (Reason: this helps to stay healthy))

VIII Regulatory Pre-requisites

Regulations are probably the most complicated kind of force we have looked at so far, so we will start with an example which, while typical, does not exhibit all the complexities we will eventually associate with the form.

(F22) Having a driver's license is a pr of driving a car.

Force: (Regulation (Authority : legal)
(Punishment : fine))

Clearly (F22) is not a strict pr. Nor could it be a social pr as social pr's do not need to indicate authority (since the authority is always "society" in some sense) nor punishment (since either there isn't any, or it is just very subtle, depending on how you look at it). As for suggested pr's, (F22) would not fit there very easily either. For one thing, regulatory pr's need not have reasons behind them. (In theory, of course, a regulation should have a reason behind it, but in practice that reason is not why the regulation is obeyed.)

A more interesting regulatory pr is the following:

(F23) Not being married is a pr of getting married.

Force: (Regulation (Authority : legal))
(Social)

What makes this rule interesting is the effect of breaking it. What happens, of course, is that the action is declared "invalid" in the sense that the end result, the two people being married, is declared no longer to be the case. In some strange sense then (F23) acts like a strict pr, but there are many differences. For one thing, breaking a regulation only affects certain results. Those resulting states which are affected we will call "conventional" states. That not all results are affected is not easily seen in the case of (F23), for getting married does not have any particular end results other than being married, but we can see this difference in:

(F24) The receiver being a legal player is a pr

of completing a forward pass in football.

Force: (Regulation (Authority : Referee,
also participants))

The completion of a forward pass includes the end result that the "receiver" has the ball. However, should it turn out that (F24) was broken, say an extra player sneaked out on the field, while we would still say that the forward pass was invalid, it would nevertheless be the case that the receiver had the ball at the end of the pass. What would differ is that the other results of a forward pass, like the ball being placed at the location where the receiver was finally stopped, would not hold. (Where the ball is placed at the start of a play in football is of course part of the rules of football and hence conventional.)

Another difference between the effects of a broken regulation and a broken strict pr is that with the former the appropriate authority must notice the infraction before the conventional result is declared invalid. So we have the following rule:

(RB) If F is a fact about action A with regulatory force, should F's being broken come to the attention of the appropriate authority the conventional results of A will usually be declared not to hold.

Notice that (RB) only says the results will usually be declared invalid. To some degree this depends on the "importance" put on the regulation. For example, regulation (F23) will be strictly enforced in our culture, but at the other extreme we have a regulation like:

(F25) The ship's being at sea is a pr of the ship's captain marrying a couple.

Force: (Regulation (Authority : legal))

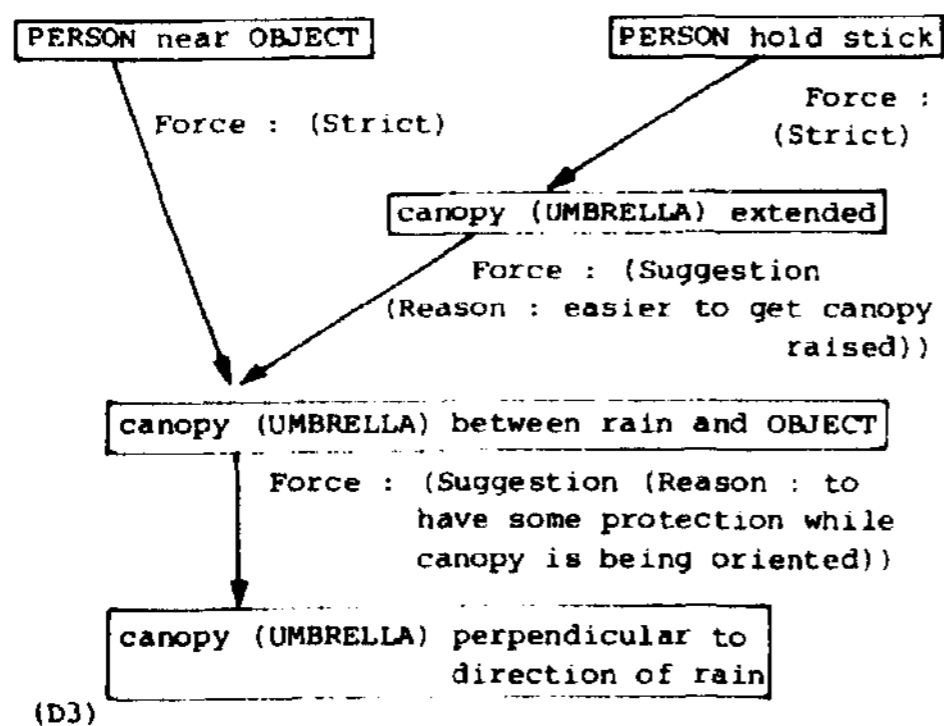
If, perchance, the ship had pulled into port before the ceremony started, one would not be surprised that upon later learning of this fact everybody decided to ignore it. Even formulating this notion of importance seems hard enough, but as we will see in a moment, there are cases where in spite of the regulation being important (R8) still may not be applied. So, for the time being, the circumstances under which broken regulations will be ignored is an open problem.

One regulatory pr which is interesting because it is a common one, especially in the Artificial Intelligence literature, is the following:


(F26) Owning an object (or having the owner's permission) is a pr of selling or otherwise transferring ownership of the object.

Force: (Regulation (Authority : owner,
participants or law))

The fact that you must own an object in order to sell it is often presented as a strict pr in the A.I. literature, but it seems more plausible that it is a regulation. For one thing, (F26) obeys rule (R8) concerning the invalidation of its results, rather than the corresponding rule for strict pr's, (R4). So notice that it is not necessarily



The major thing left out of both (D2) and (D3) is information about how to get the umbrella canopy extended if necessary. As we have already noted, we cannot simply add new substates to (D3) since if the canopy is already up there would be no point in putting it up again, or worrying about the various substates which must hold if the canopy is to be raised. What we will do instead is to include in our information about using an umbrella a "method" statement to the effect that if it is necessary to get the canopy up, the way you must do it is as follows ... So our statement about "canopy up" will now look like:



And the raise-canopy method will go as illustrated in (D5).

XI Force and Form

Sections four through eight of this paper were devoted to a classification of the different kinds of force, in each case as applied to the pr form (and hence to the substate form). In the last two sections we have introduced three new forma, but have not shown which forces they take, except that our analysis of umbrellas showed that most all of the forms occur with both strict and suggested force. In fact, all of the forms we have seen so far combine with all the forces. To illustrate this let us look at some social and regulatory force examples (F27)-(F38).

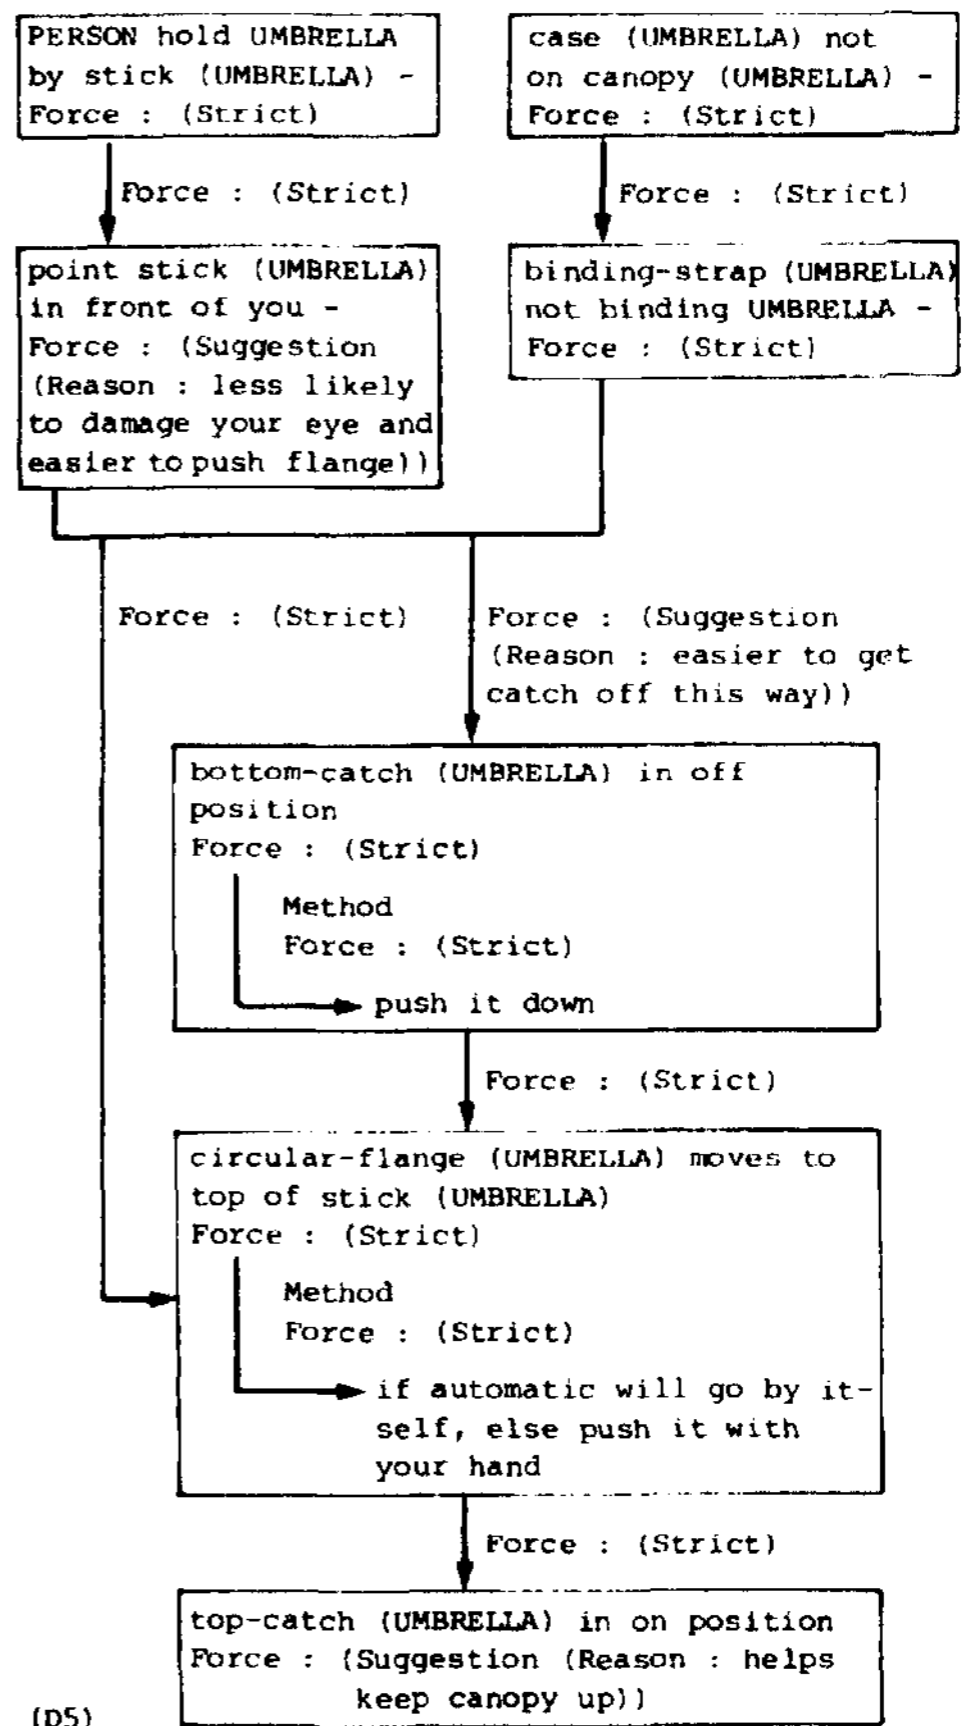
SOCIAL RULES

Subactions

- (F27) When you get married, give a party.
- (F28) When you make a phone call, announce your name.

Methods

- (F29) To get a waiter's attention in a restaurant, make a small gesture at him while he is looking in your direction.



- (F30) To cut something while eating, use your knife held in your right hand.

Time Orderings

- (F31) Have permission to use object belonging to someone else prior to using it.
- (F32) Announce your name on the phone first thing.

REGULATIONS

Subactions

- (F33) Bidding is part of bridge.
- (F34) When making a turn while driving one must signal.

Methods

- (F35) In some shops it is necessary to have a numbered slip indicating your place in the serving queue; one must obtain this

slip from the normal dispenser.

- (F36) In many Italian bars one pays by first telling the cashier what one wants, giving her the money, and getting a receipt which is in turn given to the barman.

Time Orderings

- (F37) It is necessary to obtain the slip mentioned in (F35) after entering the shop on the current buying trip.
- <F38) In bridge the dummy is layed down after the first card is played, but before the second.

XII Other Forms

The four forms we have investigated so far hardly exhaust the possibilities. In this section we will look at two more, but more in the spirit of ongoing inquiry, rather than finished products.

Our four established forms can be put into two classes. One class, which I will call "occurrence statements" consists of those statements which describe some action or state which should occur. Naturally, substates and subactions fall under this heading. Another form which is also in the category is "side condition". A side condition is an event or state which in some sense describes a second state or action. For example, a side condition of moving the flange up the stick of the umbrella is that the canopy raises at the same time and to the same degree. Note that, on one hand this does not qualify as a separate subaction, that is one does not extend the canopy separately from raising the flange, yet we cannot simply ignore this direct relation between the two events. For example, it enables us to understand how a canopy could be half extended. Another example of a side condition is:

- (F39) Paint being left on the paint brush is a side condition of transferring paint from the brush to the object being painted.

Force : (Strict)

Notice that both of these side conditions have strict force. At the moment I do not know if other kinds of force are possible.

We have said that side conditions are "occurrence statements" along with substates and subacts. Time ordering statements and method statements on the other hand fall into the category of "control statements". Control statements tell us not what needs doing, but rather when or how to do it. The obvious analogy is with control information in programs, with ordering relations corresponding to the implicit ordering of program statements, and method statements corresponding to subroutine calls. This analogy suggests that there ought to be other kinds of control information needed to describe actions, and indeed there are. One particular is the "do-while" statement. So in painting (using a brush) we know to dip the brush in the paint and then transfer the paint to the object. However, it is necessary in general to

repeat these two steps until the job is done. This is a strict do-while statement, but do-whiles occur with other forces also.

- (F40) In washing one's hair one rubs soap into the hair and then rinses it off. Repeat these two steps until soap lathers easily.

Force : (Suggestion (Reason : makes hair cleaner))

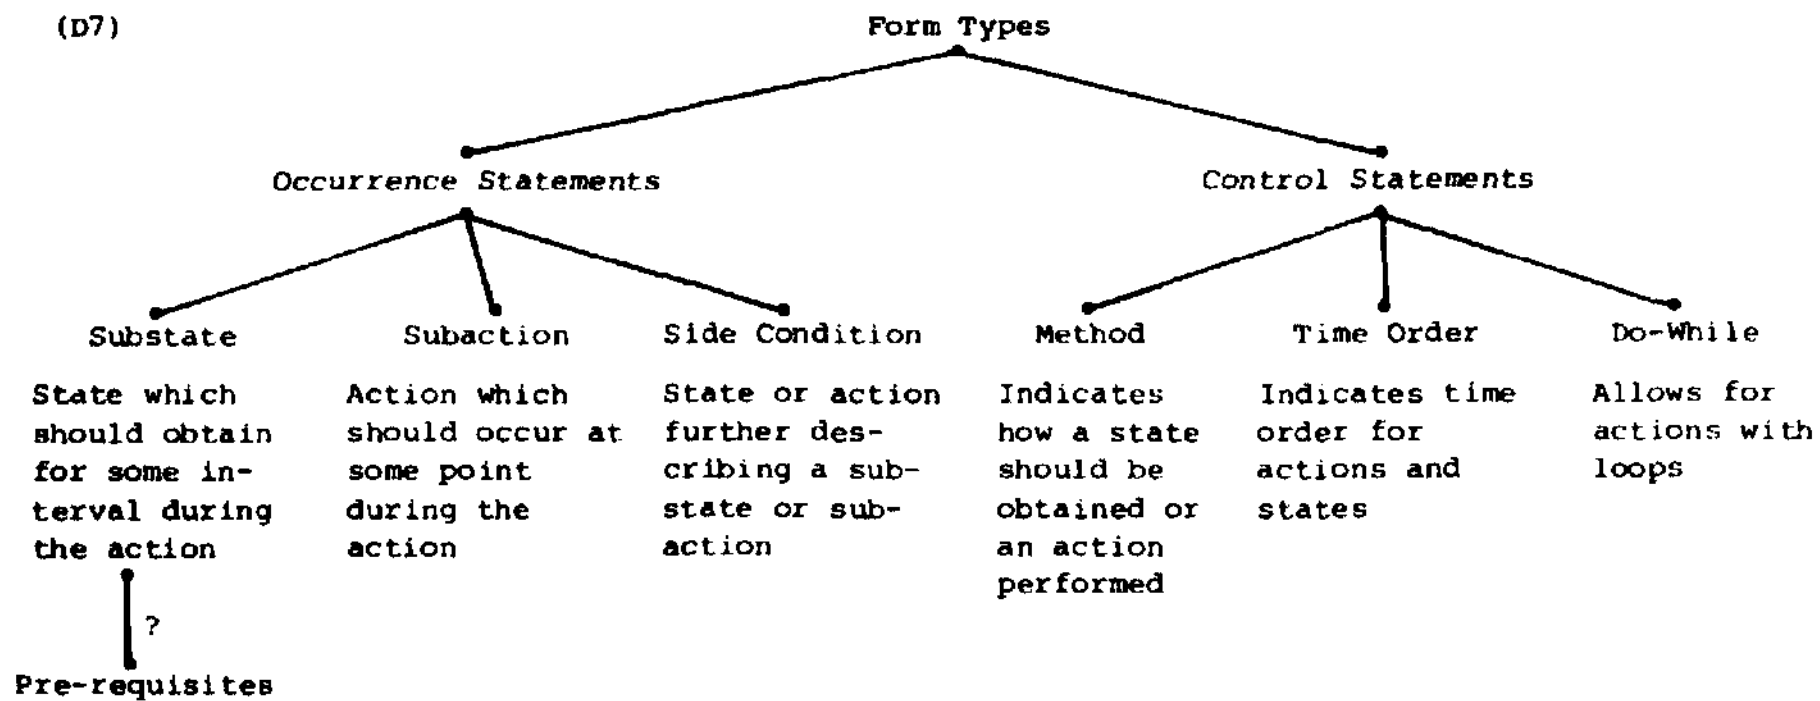
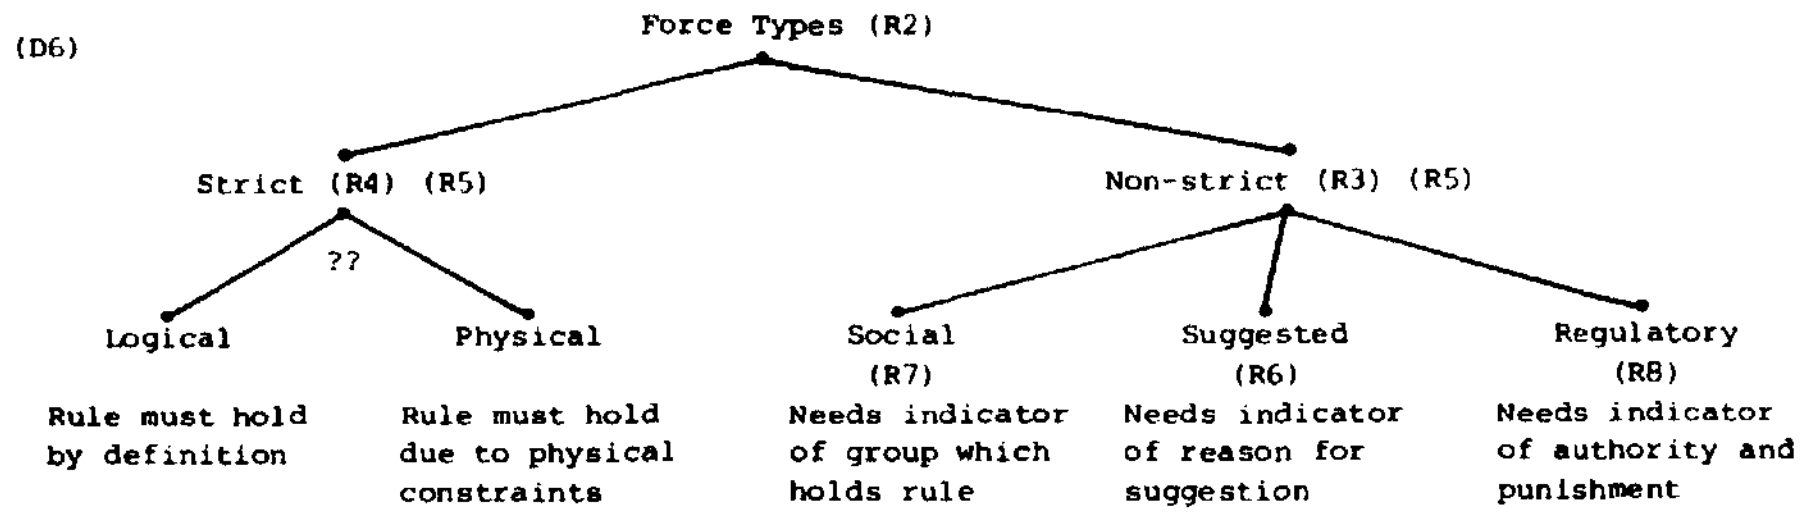
- (F41) In bidding in bridge, one bids after the person to one's right. This is repeated until either four passes or a bid followed by three passes.

Force : (Regulation)

XIII Conclusion

What I have done in this paper then is to classify knowledge about actions according to two criteria, the force behind the rule (that is, why you should follow the rule) and the form of the rule (or what the rule looks like). As such, the paper can be summarized in diagrams (D6) and (D7). The question marks in the diagrams indicate distinctions which I am not sure need be made. In (D6) the numbers following the force types indicate rules in the paper which the force type obeys.

I have so far not mentioned other people's work in relation to the results presented in this paper. Space limitations prevent any detailed discussion, but obviously this work was not done in a vacuum, and there are several papers which have influenced it. Some particular influences are: the "absolute" vs "reasonable" conditions on actions found in (Schank 74), the classification of inference types of (Rieger 74), the "infelectivities" of (Austin 62), the idea of "institutional facts" found in (Searle 69), and the "frames" representation of (Minsky 74).



References

1. Austin, J. L. How to do things with words. Oxford: Oxford University Press, 1962.
2. Charniak, E. "He will make you take it back": a study in the pragmatics of language. Working paper 5, Institute for Semantic and Cognitive Studies, 1974.
3. Minsky, M. M. A framework for representing knowledge. AI Memo 306, MIT Artificial Intelligence Laboratory, 1974.
4. Rieger, C. J. Conceptual memory. Ph.D. Thesis, Department of Computer Science, Stanford University, 1974.
5. Schank, R. C. Understanding paragraphs. Working paper 6, Institute for Semantic and Cognitive Studies, 1974.
6. Searle, J. R. Speech acts. Cambridge: Cambridge University Press, 1969.