

NOTICE WARNING CONCERNING COPYRIGHT RESTRICTIONS:
The copyright law of the United States (title 17, U.S. Code) governs the making of photocopies or other reproductions of copyrighted material. Any copying of this document without permission of its author may be prohibited by law.

Computer Models of Human Personality Traits

Jaime G. Carbonell

**Carnegie-Mellon University
Computer Science Department
Pittsburgh, Pennsylvania 15213**

November, 1979

This research was sponsored in part by the Defense Advanced Research Projects Agency (DOD), Monitored by the Office of Naval Research under contract N00014-75-C-1111, and in part by ARPA Order No. 3597, monitored by the Air Force Avionics Laboratory under Contract F33615-78-C-1551. The views and conclusions contained in this document are those of the author, and should not be interpreted as representing the official policies, either expressed or implied, of the Defense Advanced Research Projects Agency or the U.S. Government.

Table of Contents

1. Why Analyze Personality Traits?	1
2. What Information Does a Personality Trait Convey?	2
3. How Personality Traits May Be Represented	6
4. Goal Trees Representing Personality Traits	8
5. How Personality Traits Constrain the Application of Planning and Counterplanning Strategies	18
6. Combining Personality Traits	23
7. Reactions Upon Failure of Strategies	25
8. Conclusion	28
9. References	29

Computer Models of Human Personality Traits

Jaime G. Carbonell
Carnegie-Mellon University
Department of Computer Science
October, 1979

Abstract

A goal-based analysis of human personality traits is presented with the objective of developing a comprehensive simulation model. It is shown that understanding trait attributions is an integral part of story understanding and therefore much of natural language processing. The model of personality traits is derived from the goal trees of the POLITICS system, the notion of social prototypes, and planning/counterplanning strategies.¹

KEY WORDS AND TOPICS: personality traits, natural language, goal trees, story understanding, inference, heuristics.

1. Why Analyze Personality Traits?

Understanding stories requires information and reasoning about the situation, the causal structure of the events, and the characters in the story. Schank [12], Cullingford [6], Rumelhart [11], and Beaugrande and Colby [1] have analyzed the narrative structure of stories and developed means of automating the analysis process. Schank and Abelson [13], Wilensky [10], and Schmidt and Sridharan [15] developed means of inferring the goals and plans of the characters in a story from their actions. Both the narrative structures and the goals and plans of the

¹This research was sponsored in part by the Defense Advanced Research Projects Agency (DOD), Monitored by the Office of Naval Research under contract N00014-75-C-1111, and in part by ARPA Order No. 3597, monitored by the Air Force Avionics Laboratory under Contract F33615-78-C-1511. The views and conclusions contained in this document are those of the author, and should not be interpreted as representing the official policies, either expressed or implied, of the Defense Advanced Research Projects Agency or the U.S. Government.

characters are crucial in integrating the information contained in stories into a coherent memory representation. Such memory structures are necessary to answer questions about the story in much the same way that people appear to reason about the stories they read.

Character development, however, is an important aspect of story understanding that has been largely ignored by Artificial Intelligence researchers. A person reading a story identifies with one or more characters depending on whether the characters are heroes, villains, compassionate, intelligent, unscrupulous, etc., and depending on how the character's personality relates to the reader's self-image and to other people he knows in real life. Furthermore, knowledge of the characters and their personality helps to interpret their actions and induce their goals. Thus, understanding character development is an integral part of processing natural language stories. Here we deal with the most simple form of character development: the attribution of personality traits to actors in simple stories. We analyze personality traits in terms of personal goal trees and predispositions towards applying certain classes of planning and counterplanning strategies. Goal trees and counterplanning strategies were developed to model ideological beliefs in the POLITICS system [3].

2. What Information Does a Personality Trait Convey?

Consider an example of personality-trait attribution in the following story.

- (1) John is very ambitious.
He abandoned his invalid mother, worked very hard at his job, and badmouthed his coworkers. John was elated when the boss promoted him.

What does it mean for somebody to be ambitious? John's actions in example 1 are characteristic of an over-ambitious person. John's emotional reaction to his promotion also characterizes the type of behavior that one may expect from an ambitious person. What happens if we use different personality traits to define

John's character? Consider the following story:

- (2) John is very compassionate. He abandoned his invalid mother, worked very hard at his job, and badmouthed his coworkers. John was elated when the boss promoted him.

Story 2 is not consistent. Why not? Compassionate people do not abandon invalid mothers. Badmouthing coworkers does not seem to be in character with John being compassionate. The only way we could interpret story 2 is to say that John must have been acting "out of character" for some unknown reason. This suggests that the meaning of words describing personality traits are related to certain characteristic types of behavior. In fact, personality traits often express the deviation between socially defined normative behavior and the particular characteristic behavior of an individual.

Before we analyze the meaning and the subjective nature of personality traits, let us see why this is an important issue that requires our investigation. There are psychological reasons that suggest that the way people talk about personality traits may be an interesting subject of study. Here we focus on the relevance of linguistic descriptions of personality traits to understanding narrative events. Consider the following story.

(3) - Bill was very brave, but his brother John was very cowardly. One night the two brothers were walking by the road when a masked bandit surprised them. The younger brother panicked and ran headlong into the forest where he was lost, never to be seen again. The elder brother fought off the bandit, and, in the process, recovered the long lost royal sapphire, stolen years earlier. The king rewarded him handsomely.

QUESTION: Whom did the king reward?

A person reading the above story has little trouble in answering the question: Clearly, the king rewarded Bill. However, it is not particularly easy to see how one goes about formulating the answer. In order to answer this question without substantial effort, the referent of "him" in the last sentence of 3 must have been

determined while the story was understood. Determining this referent is a very difficult task. Many people have worked on the reference problem (e.g., Charniak [5], Ross [10], Wilensky [16], Cullingford [6]), but resolving this particular referent requires a complete understanding of the story. No simple rule will serve. For instance, the last mentioned character in the story before the word "him" is the bandit, but this is obviously not the correct referent.

The first step in determining the referent is to understand the causal relations among the actions in the story. In order to establish a causal relation between the king rewarding somebody and the sapphire being recovered, one must infer that the sapphire was returned to the king by the elder brother. A story understanding system such as PAM [16] makes this kind of inference and the inference that the king felt indebted to the elder brother. The indebtedness state may have caused the king to reward the elder brother. If the story is thus understood one is able to establish that "him" refers to the elder brother. This, however, is only half of the task. How do we know that Bill is the elder brother who deserves the reward?

In order to determine which brother is which we must use the information contained in their respective character traits. One brother is brave; the other is cowardly. Running away in the face of danger is a characteristic behavior associated with cowardly people. Fighting bandits, or otherwise risking one's life for a worthy cause is the type of behavior characteristic of bravery. Therefore, we determine that Bill, the brave one, must have been the elder brother who fought the bandit and recovered the sapphire. This determination requires knowledge about some types of actions that are characteristic of bravery and other actions that are characteristic of cowardice. Thus, we need to know, or be able to infer, typical behaviors associated with certain character traits. We need to answer the general question: If actor X has character trait P, is he likely to do action A in situation S? It seems, therefore, that an investigation of personality traits and their associated typical behavior ought to be a worthwhile pursuit.

Let us consider a couple of events where knowledge about personality traits is

necessary to understand the actions of the characters. We present two events that differ only in the characterization of the primary actor. The difference in the actor's personality trait accounts for a difference in the probable meaning of the unknown word "tolliked".

- (4) John was a very generous person. When the charity drive asked him for a contribution he put his hand on his wallet and tolliked their request.
- (5) John was a suspicious miser. When the charity drive asked him for a contribution he put his hand on his wallet and tolliked their request.

QUESTION: What does "tolliked" mean?

We cannot be sure of the meaning of "tolliked" in either example, but story 4 suggests a very different meaning for "tolliked" than 5. Generous people are usually willing to share some of their possessions or their time with people in need. Therefore, the reader of 4 might expect John to respond in a positive manner to the charity request. The fact that he put his hand on his wallet can then be interpreted as a precondition to giving money to the charity, thus fulfilling their request. In light of these expectations, we can determine that "tolliked" probably means "complied with" or "fulfilled" Thus, the meaning of "tolliked" is postulated in a large part by the type of behavior one might expect from a generous person under the circumstances of story 4.

What about the meaning of "tolliked" in story 5? Misers do not share their possessions with anybody. Suspicious people distrust the apparent motives of others. Thus, the reader of 5 will expect that John does not want to give money to the charity and that he may mistrust the motives of the person asking for a charity contribution. These expectations may lead the reader to interpret John reaching toward his wallet as a precaution for any tricks that he may suspect on the part of the charity drive person. The fact that John is a miser and the above interpretation of reaching for his wallet suggest the same course of action for John to follow: John

will not comply with the charity request. In this case, "tollikened" takes on the meaning of "denied" or "dismissed". This meaning of "tollikened" is quite different from the meaning suggested by story 4. The only difference in the two stories is that different personality traits were attributed to John. Hence, we see that understanding personality traits is important for generating expectations about probable behaviors and, in some cases, postulating the meaning of unknown or ambiguous words.

3. How Personality Traits May Be Represented

We might consider defining personality traits, such as "ambitious" and "compassionate", by listing the set of behaviors characteristic of that trait. Recall John's behavior in story 1. All his actions are, in a sense, characteristic of an ambitious person, but what do we do with actions that are not characteristic of a particular trait? For instance, neither abandoning one's invalid mother nor working very hard at one's job are characteristic actions of a compassionate person. However, the former action is definitely uncharacteristic of compassion, while the latter action is neutral with respect to being compassionate. Thus, if we are to define character traits by listing characteristic actions, we should also list actions that are typically uncharacteristic of the particular trait.

We must take into account the monumental nature of the task if we are to list all characteristic and non-characteristic actions for each character trait. There are, in essence, infinite numbers of actions that can be classified as characteristic or uncharacteristic for each personality trait. We will try to narrow the problem by only classifying general types of actions. For instance, consider a very incomplete list of jobs characteristically aspired by an ambitious person: president of a company, trial lawyer, real-estate king, shipping magnate, Governor of California, advisor to the President of the United States, and movie star. All of these occupations entail power, wealth, and social respect, to different degrees. Therefore, a useful way of classifying these jobs is by the degree to which they imply high social status, power and wealth. This classification enables us to have only one entry on the list of

typical actions of an ambitious person: He wants a job that maximizes these three qualities. Similarly we can generalize some other actions that characterize ambition. Our list of typical actions and wants remains somewhat cumbersome in length. Table 1 is an incomplete list of characteristic and uncharacteristic actions and wants of an ambitious person.

AMBITIOUS	
CHARACTERISTIC ACTIONS	UNCHARACTERISTIC ACTIONS
1) Neglecting relatives in time of need.	1) Anonymous donations to charity.
2) Wanting a job with as much power, respect and wealth as possible.	2) Stepping down to let a more qualified person assume a position of responsibility.
3) Wanting to constantly improve one's present Job.	3) Avoiding hard work that leads to self betterment.
4) Using friends to further one's own ends, then discarding them.	4) Helping others at cost to self.
5) Badmouthing competitors.	5) Not being concerned with personal appearance in the presence of one's boss or social peers.
6) Wanting social respect and recognition.	6) Being contented with one's past achievements in life.
7) Want an impressive house.	7) Relinquishing social status, wealth or power.
8) Want a luxury or sports car.	8) Placing honesty above self betterment.
9) Want a socially successful spouse.	9) Tolerant of other people's faults.
10) Dishonest business deals.	10) Happy at another's success.
11) Seeking to be in the presence of successful people.	

Table 1: Characteristic actions of an ambitious person.

The set of characteristic actions and wants of an ambitious person is based on underlying personal motivation. Wanting an impressive house, a luxury or sports, car and a job that yields substantial wealth are instances of acquisition goals (A-goals). Being respected and powerful are instances of A-scont goals. (A-social control means desiring an increase in one's social stature. In addition to acquisition goals, there are other types of goals such as preservation goals (P-goals) and enjoyment goals (E-goals). The goal taxonomy is borrowed from Schank and Abelson [13].)

Thus, one way of analyzing personality traits is by associating with each trait the

goals people described by that trait are likely to have. Once these goals are established, certain behaviors can be inferred in particular situations, such as stories 3, 4 and 5, by a story understander applying planning and counterplanning strategies. Since we have developed mechanisms for understanding goal based events (e.g., PAM [16] and POLITICS [3, 2]), it seems quite fruitful to reduce personality traits to the pursuit of certain types of goals.

4. Goal Trees Representing Personality Traits

Consider the process of understanding a story starting with the following initial segment:

- (6) John is a very inquisitive and uncompromising person. He is also rather thrifty in his personal affairs...

There have been no actions thus far in the story, nor any physical or temporal setting that helps the understander establish the situational context. Yet, John's personality traits provide a goal-expectation setting. That is, the understander knows the following information from the above fragment of 6: John's goal of increasing his knowledge about most matters is a goal of very high importance. We denote the acquisition of knowledge goal as A-know(John,X,+). (The "+" means John wants knowledge about X. A "-" would signify that John's goal is to actively avoid knowing about X, and a "0" signifies that John ignores new knowledge about X. Thus, if we know that Mary is apathetic, we mean A-know(Mary,X,0).) The fact that John is thrifty tells us that he also has the goal of preserving his money. In fact, the word "thrifty" states a relationship between the P-money goal and the set of A-goals that can be accomplished by spending money. John holds the goal of P-money to be more important than most such A-goals.

The fact that John is uncompromising is somewhat more difficult to represent in terms of John's goals. No specific goal is defined by a person being uncompromising. This personality trait applies to all of John's goals. An uncompromising person is one

who does not abandon any goal in the face of opposition from another party; that is, a person who will not yield to someone else's goals. Being uncompromising also carries the implication that one holds the goals of others to be less important than is normally the case. Thus, this personality trait modifies the entire set of goals that a person has, rather than establishing a single specific goal.

Since most personality traits describe deviations from a culturally-defined normative person, we know that John's A-know goal is much more important to him than other people's A-know goals are to them. Similarly, we know that his P-money goal is a little more important to him than is generally the case. We may also infer that John's A-goal of things that cost money may be a little less important to him than other people's corresponding A-goals are to them. The trait "Uncompromising" exemplifies an across-the-board deviation from the norm. John will give higher than normal importance to most of his goals.

These importance relations enable us to construct a relative-importance (RI) goal tree for John in the same manner that we constructed goal trees for political ideologies in POLITICS [3]. Here is the fragment of John's goal tree, constructed from the information contained in the personality traits in story 6.




Figure 1: Fragment of John's RI goal tree.

Figure 1 tells us that of the three goals that we know John to have, he considers acquisition of new knowledge as most important, followed by preserving his money, followed by acquiring new material possessions. Since we know that John is a person and a member of western society we know that he has certain normative

goals common to most people in the society. These goals include:

GOAL	EXPLANATION
1) P-health(Self,+)	Self-preservation
2) P-health(Family,+)	Preservation of family members
3) A-possession(Self,+)	Acquire wealth and belongings
4) P-possession(Self,+)	Preserve one's belongings
5) A-social respect(others,+)	Be respected by other people
6) A-know(Self,X,+)	Learn new things
7) E-unpleasant activity(Self,-)	Avoid going through unpleasant experiences (e.g., stay out of jail)
8) E-pleasant activity(Self,+)	Have fun doing enjoyable things
9) P-health(others,+)	Help others survive
10) P-anything(enemies,-)	Wish doom upon one's enemies

Enjoyment goals (E-goals) are a third type of goal in Schank and Abelson's [13] goal taxonomy. These goals are usually of a more fleeting nature, therefore less important than the A-goals and P-goals we discussed earlier.

The set of goals can be ranked in terms of their normative relative importance to a prototypical member of our society. For instance, preservation of oneself and one's family are usually the two most important goals for anybody. However, if we learn that a person is foolhardy, we interpret this as a deviation from the normative goal tree where P-self is a low importance goal to that person. If the person is described as suicidal, the P-health goal is not present in his goal tree; indeed, P-health(self,-) may substitute the normative P-health(self,+) goal. Similarly, wishing doom on one's enemies is, in the normative case, a lower importance goal

than most of the other goals listed above. If a person is described as vindictive, we know that his goal of P-anything(enemies,-) is much more important to him than is normally the case.

The figure below is the relative importance goal network for a prototypical, normative person. Since the network is acyclical, it is conceptually equivalent to a tree. This tree is subject to a set of inference rules developed in [3]. Some of the more fundamental rules are presented later in this section. Figure 2 is an empirical attempt at partially ordering the more common goals that people pursue. The tree includes the goals listed above; it is not meant to be a comprehensive goal tree of all significant goals that may be pursued by people in our society.




Figure 2: Goal tree for a normative person.

In the case of story 6, the personality traits applied to John promote his A-knowledge goal to a higher importance than the other three goals at that level, but keep the higher preservation goals (e.g., P-self and P-family) as most important. The fact that he is thrifty, creates the more specific P-money goal (an instance of the more general P-possession goal) and places this goal as having higher importance than his A-possession goal in the goal tree. Furthermore, his A-possession goal is demoted below the other goals found at that level in the

normative goal tree. The fact that John is uncompromising makes all his goals in the tree more important to him than in the normative case. Thus, we represent the meaning of John's character traits in 6 as a modification to the normative person's goal tree. How can we use this goal tree representation? Let us continue with the story.

- (6) John is a very inquisitive and uncompromising person. He is also rather thrifty in his personal affairs. One day he got an unusually large repair bill for his car...

What is John likely to do about the repair bill? The fact that John's P-money goal has high importance suggests that he may not be willing to part with the large amount of money required to pay the bill. His high importance A-know goal suggests that he may want to discover why the bill is unusually large; he may want to know whether he is being cheated by the repair shop. John will pursue these two goals with more than usual determination; this is, in essence, the meaning of John being uncompromising. We do not think that a reader of 6 would predict any specific actions on John's part, but understanding his goals helps one to understand the nature of later actions that may follow in the story. We believe that it is crucial for the reader to be aware of the fact that two of the goals that John holds in high importance have been violated by the unusually high repair bill. Without this knowledge one is unable to comprehend the entire story. Furthermore, the general expectation should be made by the reader that John's actions in the immediate future will probably be attempts to fulfill either or both P-money and A-know. We illustrate this claim by giving the conclusion of the story.

- (6) John is a very inquisitive and uncompromising person. He is also rather thrifty in his personal affairs. One day he got an unusually large repair bill for his car. John called his brother, a wealthy lawyer, to take care of the problem.

Our interpretation of the last sentence is that John's brother will probably dispute the repair shop on John's behalf to lower the bill. This conclusion is arrived at in light of the goals that John is expected to be actively pursuing, P-money and A-know. How does calling a wealthy lawyer help John achieve these goals? The job of a lawyer is to act as an agent furthering the goals of his client. A lawyer may be more successful at getting the repair bill lowered. The retaining fee associated with hiring a lawyer (violating P-money) may not apply because of the family relationship. The fact that John is uncompromising supports our interpretation of the conclusion; John is pursuing his violated goals as best he can.

Why did we need to generate the expectations that John would pursue P-money and A-know? The answer is: we need to know John's goals in order to correctly interpret the conclusion. John's goals come directly from the personality traits of the first two sentences. In order to illustrate the necessity of determining John's goals before interpreting the conclusion, consider the following similar story:

- (7) John is an apathetic, happy-go-lucky person.
 He is also somewhat of a spendthrift.
 One day he got an unusually large repair bill for his car. John called his brother, a wealthy lawyer, to take care of the problem.

Our interpretation of this story is that John may have wanted some money to pay for his repair bill. His brother, being a wealthy lawyer, was a possible source for the needed money. (In an informal confirmation of our interpretations we gave story 6 to five people and 7 to five other people. Everyone was asked to explain the story they read, and, for the most part, their interpretations concurred with ours.) John, being apathetic, attaches little importance to his A-know goal. The fact that he is a spendthrift means that he does not attach much importance to his P-money goal, but he may give more importance to A-possession or E-things that cost money. Furthermore, someone with a low P-money is likely to not have much money in hand. Therefore, a reasonable expectation is that John may be unable to pay the repair bill. Finally, a happy-go-lucky person does not bother to pursue his goals with much

determination. He is more interested in the quickest solution to the present dilemma. With these goals (or the lack thereof) in consideration when we interpret the last sentence of 7, we conclude that John only cares about dismissing the problem of the repair bill as quickly as possible. A loan or gift from his rich brother fulfills our expectations of John's probable behavior.

If 7 were continued with "John's brother said he had already loaned John too much money." we would confirm our expectation that John was asking for financial assistance. On the other hand, the above continuation following event 6 makes little sense. This illustrates the fact that we did not expect John to ask for money in 6. The continuation is a response to a non-existent expectation, therefore it is not surprising that such a continuation following story 6 is rather puzzling. Thus, different expectations and, therefore, different interpretations are produced by the different personality traits used in the two stories.

The only difference between stories 6 and 7 is the characterization of John's personality by a few personality traits. Therefore, once again, we relied on goal-based information implied by these character traits in order to interpret a story. This suggests that understanding stories where the characters are described by personality traits is inherently a goal-oriented process. The understander considers only certain inferences and certain interpretations of specific behaviors on the basis of the inferred goals and motivations of the primary characters. For instance, in interpreting story 6 one could infer that John would ask his brother for money, for a new car, for moral support, or for a suggestion of a less expensive repair shop. None of these inferences are made in interpreting 6 because the understander already expects John's actions to be in service of particular goals. Thus, goal-orientation is a general method of pruning spurious inferences. In order to carry out the goal-directed inference process, we must first construct the goal trees for the people described by personality traits. The goal trees are constructed from the normative-person goal tree modified by the goals implicit in the character traits.

PERSONALITY TRAIT	GOALS AND THEIR IMPORTANCE (deviations from the socially accepted norm)
1) Ambitious	A-possession (self,+) higher A-scont (self,others,+) higher P-anything (others,+) lower A-know (self,+) slightly higher
2) Curious	A-know (self,+) higher
3) Prudent	P-anything (self,+) higher P-anything (others,+) slightly - higher
4) Spendthrift	P-money (self,+) lower P-possession (self,+) slightly - lower E-things/that/cost/money (self,+) slightly higher
5) Vindictive	P-anything (others who have caused goal failure, -) higher
6) Powerhungry	A-scont (self,others,+) higher
7) Compassionate	P-health (others,+) higher P-anything (others,+) slightly - higher E-unpleasant experience (others,-) higher
8) Playboy	E-sex (self,+) higher E-anything (self,+) slightly - higher A-luxurious poss (self,+) higher P-money (self,+) slightly lower
9) self-centered	<any-goal> (self,+) slightly - higher <any-goal> (others,+) lower
10) Belligerent	A-scont (self,others,+) slightly higher Cause goal-conflicts

Table 2: Goal-oriented personality traits.

Table 2 lists the deviations from the social normative goal tree for each goal implicit in various personality traits. For instance, an ambitious person attributes higher importance to the goals of increasing his social status (i.e., power and prestige), his wealth, and his worldly possessions. Ambition also implies less concern

for the goals of others, especially if any of their preservation goals conflict with the ambitious person's A-goals. Thus the goal tree of an ambitious person is the normative person's goal tree (figure 2) with the above goals raised or lowered in importance according to the entries in table 2.

In [3], we discuss a set of Inference rules that focus the inference process by analyzing the goal tree to determine the subjectively most interesting aspects of a situation. These rules also apply to personality-trait goal trees, as illustrated below. We present some of the more widely-applicable rules:

RULE 1 If progress towards a goal can be achieved by a particular course of action, that course of action should be pursued.

RULE 2: If a possible course of action violates a goal, it should be actively avoided.

RULE 3: If a course of action affects two goals, and no other rules determine which goal to focus on, the effect on the higher-importance goal determines whether the course of action should be pursued.

RULE 4: Relative importance links in a goal tree are transitive.

Let us apply these rules to the interpretation of the following two stories.

(8) John, an ambitious lawyer, had to decide whether to accept the lucrative GM contract or devote his time to the free legal-aid society. It did not take long to make up his mind.

(9) John, a very compassionate lawyer, had to decide whether to accept the lucrative GM contract or devote his time to the free legal-aid society. It did not take long to make up his mind.

The decision confronting John in both stories is whether to pursue the goals of A-money(John,+) and A-scont(John,others,+), or to help others fulfill their P-goals. In the normative person goal tree, the two A-goals are somewhat more important than the P-anything(others,+) goal. The fact that John is ambitious raises the importance of his self-centered A-goals and further lowers the importance of helping others fulfill their P-goals. Applying rule 3 (and 4 if necessary) we conclude that John will pursue his A-money and A-scont goals. Therefore, according to rule 1, John probably made up his mind to accept the GM contract.

On the other hand, if we modify the normative person goal tree by the entry for "compassionate" in table 2, we find that P-goals of others increase in importance. This means that P-anything(others,+) is roughly equal in importance with John's A-money and A-scont goals in story 9. Which way did John make up his mind? We cannot tell unless we have some way to measure the relative increase in importance of the P-goals with respect to the base difference in importance between P-anything(others,+) and the two A-goals in the original normative person goal tree. An alternative solution to this problem involves taking pragmatic considerations of story telling into account when formulating John's goal tree. Why were we told that John is very compassionate? This fact must have some relevance to the rest of the story. The only relevance it could have is to affect John's decision. If we used a normative goal tree for John, we would expect his decision to favor accepting the GM contract. In order to affect John's decision (i.e., reverse it) we must change the relative ranking of his goals with respect to importance. Therefore, the reader of 9 will probably guess that John's goal of P-anything(others,+) takes on greater importance than his A-money and A-scont goals.

The above discussion suggests that *social prototypes* might be a useful concept to have, defined in the same spirit as Rosch's semantic prototypes [9] and Rich's preference stereotypes [8]. As a first-order approximation one can use a single normative-person social prototype, defined by the goals and relative importance relations of figures 1 and 2. However, extending our notion of normative goal tree

one can define other such trees, each corresponding to a well defined social class ("role theme" in Schank and Abelson's notation [13]). Thus, we can say that someone is "ambitious for a bum", "honest for a politician", or "more pious than a priest". In these cases the personality traits are defined as the same deviations from the norm, but the norm itself has been temporarily redefined by the social prototype and its corresponding relative-importance goal tree.

5. How Personality Traits Constrain the Application of Planning and Counterplanning Strategies

It is often useful for an understander to predict the type of planning or counterplanning strategies that a person is likely to use. Let us define what we mean by planning or counterplanning strategies. A planning strategy is a basic planning method applicable to different circumstances, such as bargaining for and object or invoking a social obligation to get another person to do one's bidding are planning strategies. These planning units are discussed in [13] where they are called planboxes. Counterplanning strategies are more complex means of achieving one's goals in spite of other parties actively trying to prevent one's goal fulfillment. Examples of counterplanning strategies include threatening higher goals of an opponent to divert his efforts away from blocking one's goals, and establishing a mutual goal-blockage situation, thereby being in a position to negotiate a resolution to the mutual goal blocking actions. Counterplanning strategies are developed in [3, 4] to model political reasoning and planning in other adversary situations. All strategies have "trick options", such as bargaining or negotiating in bad faith.

Personality traits may suggest that certain strategies are more likely to be applied by a given person that is generally the case. For example, a bully will use threats and overpowers with greater frequency and less hesitation than most people. More often, personality traits constrain the application of counterplanning strategies. A timid person is unlikely to use threats; an honest person will not use trick options.

Consider the following fragment of a conversation overheard on a bus:

(10) Fragment of a conversation.

Person 1 - Do you remember old Ed?

Person 2 - You mean the incompetent salesman who tries to cheat in our card games?

Person 1 - Yeah, he asked me for a raise today. I pulled out his employee record and you can imagine what I told him!

We certainly can imagine that "old Ed" did not get his raise. How do we know this? How does person 1 know that person 2 will understand him? The answer to both of these questions is determined by what we as readers (and person 2 during the conversation) know about "old Ed". Namely, he is dishonest and incompetent. This means that he is willing to use trick-option strategies against his boss, and he is unable to choose or carry out the appropriate strategies in his job as salesman. Thus, Ed's boss has two reasons for denying the raise, corresponding to the following two rules:

RULE 5: MAKING ENEMIES

If an actor X repeatedly counterplans against actor Y, Y will not help X achieve any goals in the future.

REFINEMENT If X is successful in his counterplanning, Y may pursue the goal of terminating any subsumption state that enables X to counterplan against Y.

RULE 2: If a possible course of action violates a goal, it should be actively avoided.

The first reason why old Ed's boss should deny the raise is that Ed had repeatedly counterplanned against his boss by trying to cheat at cards, apparently with little success. Thus, according to rule 5, Ed's boss should not be expected to help Ed by giving him the requested raise. It is interesting to note that if Ed had succeeded in repeatedly counterplanning against his boss then the refinement of

rule 5 is directly applicable. The boss could fire Ed, thus terminating the subsumption state that makes Ed's boss vulnerable to Ed's trick-option strategies. The second reason why the boss should deny the raise is that giving Ed more money violates the A-money goal that all businesses have. Thus, rule 2 vetos any raise to Ed. If Ed was not incompetent he would make more money for the business thus no A-money goal would be violated and rule 2 would not apply. (Businesses have goals too. In fact, they have goal trees, just like nations, individuals and most other institutions in our society.)

Table 3 presents some personality traits and the types of strategies likely to be employed by the person with the respective trait.

PERSONALITY TRAIT	TYPES OF SUGGESTED PLANNING AND COUNTERPLANNING STRATEGIES (deviations from social norm)
1) Ambitious	Higher-order plan boxes (e.g., THREATEN, OVERPOWER) and counterplanning strategies (e.g., BLOCK-HIGHER-GOAL). Trick-options if necessary. No compromises if possible.
2) Trustworthy	No trick-options used. Preference for lower-order strategies and compromises.
3) Dishonest	Trick-options used
4) Unscrupulous	Higher-order strategies, trick-options used disregarding all negative effects on others.
5) Compassionate	Strategies chosen not to harm others and, if possible, to help others fulfill their goals.
6) Capable	Make correct decisions in selecting the proper strategies for each situation. Carrying out strategies without errors.
7) Incompetent	Random or error-prone choice of strategies. Possibly not aware of some strategies.
8) Belligerent	Choice of strategies to maximize plan-conflicts with others. Higher-order strategies used when not necessary.

Table 3: Means-oriented personality traits.

We emphasize that it is important to understand the strategies implied by certain character traits. Without analyzing the strategies in the previous example we would not have been able to invoke rule 5 because we would not have discovered the relevance of the repeated counterplanning on Ed's part. The simple-minded alternative to analyzing the goals and strategies underlying personality traits is to associate all possible outcomes with each trait. For instance, under "Incompetent" one would have to store (and consider each time this trait is mentioned) that one may be denied raises, fired from the job, abandoned by one's friends, scorned by

neighbors, do badly in studies, lose at most games, have an unhappy life, have a rather low intelligence, etc. ad infinitum. This method of directly associating behaviors with personality traits has many drawbacks, such as the sheer size and inefficiency of the memory required to store all behaviors associated with all character traits.

Let us consider a different reason why such a method is insufficient. How do we deal with the following type of characterization? "Millard Fillmore was an incompetent president." Clearly, we do not mean that Fillmore was scorned by his neighbors and did badly on his studies. On the other hand, most of the above characterizations were true of Galileo. (He was fired, scorned, laughed at, and he led an unhappy life.) We would not say that Galileo was incompetent. Quite the contrary, he applied the proper strategies to physics problems, while his contemporaries may have been the real incompetents. Hence, the trait "incompetence" refers to a person's ability to formulate and carry out plans, rather than any specific type of behavior.

The strategy-based personality traits are defined in terms of deviations from the social norm, in the same manner that we defined goal-based personality traits. Thus, asserting that Millard Fillmore was an incompetent president means that he is less competent than other presidents with respect to his planning and counterplanning abilities in his official role as president. We have, however, a much narrower social norm for judging the competence of presidents. Both the comparison set of people is smaller, and the domain of application of the strategies upon which he is judged is much better defined. The same principle applies when we refer to a capable janitor or a belligerent priest. We would not expect a capable janitor to make correct strategy decisions in international politics, nor would we expect a belligerent priest to seek out fist fights.

6. Combining Personality Traits

Personality traits combine with each other and with other personal attributes such as role themes. (A role theme is a characterization of a person's position in society, largely determined by the person's profession - see [13].) The examples we just discussed are interactions of personality traits with role themes. The role theme defines the normative set of people with respect to which the personality trait defines a deviation. As we discussed, the role theme can also define the dimension of applicability of the personality trait. An incompetent president is incompetent with respect to his duties as president. An unscrupulous lawyer is likely to use the higher-order strategies and not worry about the consequences of his actions upon the goals of others only within the confines of his role as a lawyer. The unscrupulous lawyer might be considerate with friends or family outside the courtroom, regardless of how he carries out his professional activities.

Personality traits combine with each other to give a more complete picture of a person's goal tree and the strategies he is willing to use in furthering his goals. Our previous examples included several instances where more than one personality trait was used to describe a person. All of these examples had one important property in common: Each personality trait dealt with different personal goals or different sets of strategies. For instance a description of John as inquisitive, trustworthy, thrifty, and capable is simple to formulate. John has high A-know, high P-money, does not use trick options, and selects and applies strategies correctly. What happens when two personality traits describe the same goal or deal with the application of the same set of strategies?

We developed a set of heuristics for combining personality traits. Consider a person described by traits A and B, where A and B are defined in terms of their deviations from the normative person. A and B consist of a list of attribute-rank pairs. An attribute is either the name of a goal or the name of a planning/counterplanning strategy. In the former case, the rank tells how much more (or less) important that goal is to the person described than to the normative

person. In the latter case, the rank encodes the person's relative predisposition to apply the strategy (again with respect to the normative person). The rules below encode our process for combining traits A and B.

Rule 6: COMPLEMENTARY DESCRIPTIONS.

Take the union of all the attributes that appear in only one of the two traits.

Rule 7: RESOLUTION CRITERIA.

If the attribute appears in both traits A and B, consider the two rankings and apply the following rules:

Rule 8: CONTRADICTORY TRAITS.

If both rankings have a high magnitude, but opposite sign, the two traits cannot be combined. (e.g. A generous miser, and a cowardly brave person are instances of contradictory traits.)

Rule 9: REINFORCEMENT OF EXTREMES.

If both rankings have a high magnitude and the same sign, assert the attribute with a ranking slightly larger than the maximum of the two original rankings. (e.g., An unscrupulous, vindictive person is more likely to violate other people's goals than someone who is merely vindictive, or just unscrupulous.)

Rule 10: DAMPENING MINOR VARIATIONS.

If the magnitude of both rankings is small, but the signs opposite, delete this attribute from the combined trait, as it is of little importance and uncertain consistency.

Rule 11: PREFERENCE TO EXTREMES.

If none of the above rules apply, average the two ratings, but give greater weight to the rating with the higher magnitude.

The heuristic rules were empirically derived by analyzing many personality traits into their component attributes and recombined in different ways. Our rules were

implemented as a simple production system (like PSG [7]) that generates intuitively plausible personality trait combinations. This is only a small part of our larger project, recently underway, to use personality traits as part of an integrated story understanding system.

7. Reactions Upon Failure of Strategies

Many personality traits contain information about people that cannot be encoded in terms of goal trees or preferences for certain types of strategies. However, personality traits in general describe some aspect of the individual that deviates from the socially-defined, normative person. The aspects of personality traits that are outside the scope of our investigation include emotional and attitudinal attributes (but see [14]). For instance there is more to a sensuous person than a person whose goal of E-pleasure is high. Similarly, goals or strategies alone cannot fully describe "meek", "moody" or "outgoing" people.

There is, however, one other aspect to personality traits that can be usefully investigated within our paradigm. People have different reactions toward success or failure of their planning and counterplanning efforts. Some personality traits imply certain types of behavior. A contented or aesthetic person will have a much more restrained reaction to success than an ambitious person, who is likely to be spurred on to further achievements by his past success. Since most stories deal with attempts to fulfill goals that fail repeatedly before (and if) success is ever reached, we focus on reactions to failure situations. The following table includes several personality traits that imply classes of behavior upon failure.

PERSONALITY TRAIT	REACTIONS TO FAILURE OF PLANNING AND COUNTERPLANNING STRATEGIES (deviations from social norm)
1) Persistent	Try plan many times before abandoning. Then, if possible, try new plan to fulfill the same goal.
2) Ambitious	Frustration. Try new plan if possible. Otherwise immediately pursue another goal.
3) Resourceful	Analyze failure to correct the plan or to choose a more appropriate strategy.
4) Happy-go-lucky	Abandon plan and possibly goal if not too important. No frustration reaction.
5) Depressed (or Down-hearted)	Frustration. Probably abandon plan and goal.
6) Vindictive	Try to blame others for failure. Direct counterplanning effort to block the goals of whoever caused the failure. (This often takes greater importance than the original goal.)
7) Patient	No overt frustration. Try same or different plan, possibly after some time has elapsed.

Table 4: Reaction-to-failure classification of personality traits.

In order to see how the information in table 4 may be used in the understanding process, consider the following stories:

- (11) John is a vindictive person. When his vegetable garden was dug up by Bill's dog, he picked up the heavy shovel and went to Bill's house.
- (12) John is a resourceful person. When his vegetable garden was dug up by Bill's dog, he picked up the heavy shovel and went to Bill's house.

QUESTION: Why did John go to Bill's house with the shovel?

Each story suggests a different class of answers to the question. In 11 the answer is that John wants to get back at Bill. We do not know whether he will use the shovel to overpower Bill, dig up Bill's garden in revenge, or some other counterplanning act. In 12 the most logical explanation seems to be that John wants Bill to fix up his garden, informing Bill that it is his responsibility. In any case, whatever action John intends in 12 is focused on the goal of repairing his damaged garden. In 11 the stronger expectation is that John wants revenge for the damage. These general expectations, coming from table 4, help to guide the understander in further interpretation of either story. Consider the following as a possible continuation to 11 and 12:

CONTINUATION: John started digging top soil from Bill's yard.

In story 11 this continuation is interpreted as revenge for what Bill's dog did to his garden, but in 12 the same continuation makes more sense as a part of a plan to repair John's garden. Top soil is necessary for a garden. The reason for the two diverging interpretations is the understander's different expectations about John's currently active goal. In 11 the continuation is first interpreted in light of the expected revenge, and a plausible interpretation is found. Hence, an inference mechanism modeling human understanding need not (should not) look further. In 12 the continuation is interpreted in light of the expectation that John is trying to repair his garden. As before, a plausible interpretation is found for John's action

(stemming from the use of topsoil) and one needs to make no other inferences.

8. Conclusion

Personality traits define a goal-based context in which to interpret further events in a story. Without this context no explanation can be found for many events. If the continuation above occurred as a sentence by itself, we could make any number of inferences as to John's possible intent. He could be digging worms to go fishing, laying the foundation to Bill's house, planting trees, or be building a dam with the soil. Neither these nor other spurious inferences need to be made in interpreting the continuation as part of story 11 or 12.

Some personality traits, such as ambition, encode information about all three aspects discussed above: the relative importance of goals, tendencies towards invoking certain strategies, and reactions to success or failure. Other personality traits focus on one specific aspect with a higher degree of certainty. For instance, dishonesty refers only to a willingness to use trick-option strategies, but the reader is certain of this aspect of a person's personality. Ambition, on the other hand, suggests many more types of goals and strategies, but with a smaller degree of certainty. An ambitious person will probably use the higher-order strategies, but we can easily conceive of an ambitious scientist who does not spend his time threatening, overpowering, or deceiving people.

The usefulness of our analysis of personality traits will, in part, be decided by our current project, where we are implementing a system that infers and applies personality traits in the process of interpreting natural language accounts of human conflict situations. In our system, personality traits are used to help predict the existence, nature and scope of the inter-personal conflicts as well as focusing the attention of the understander on the more promising paths towards resolving these conflicts.

9. References

1. Beaugrande, R. D. and Colby, B. N., "Narrative Models of Action and Interaction," *Cognitive Science*, Vol. 3, No. 1, 1979, pp. 43-66.
2. Carbonell, J. G., "POLITICS: Automated Ideological Reasoning.," *Cognitive Science*, Vol. 2, No. 1, 1978, pp. 27-51.
3. Carbonell, J. G., *Subjective Understanding: Computer Models of Belief Systems.*, PhD dissertation, Yale University, Jan. 1979.
4. Carbonell, J. G., "The Counterplanning Process: A Model of Decision-Making in Adverse Situations," Tech. report Computer Science Dept., Carnegie-Mellon University, February 1979.
5. Charniak, E., *Towards a Model of Children's Story Comprehension*, PhD dissertation, M.I.T., 1972.
6. Cullingford, R., *Script Application: Computer Understanding of Newspaper Stories*, PhD dissertation, Yale University, Sept. 1977.
7. Newell, A., "Productions Systems: Models of Control Structures", In W. G. Chase, *Visual Information processing*, New York: Academic Press, 1973, .
8. Rich, E., *Building and Exploring User Models*, PhD dissertation, Carnegie-Mellon University, April 1979.
9. Rosch, E., "Cognitive Representations of Semantic Categories," *Experimental Psychology*, No. 104, 1975, pp. 192-233.
10. Ross, J. R., "On the Cyclic Nature of English Pronominalization", *To Honor Roman Jakobson*, The Hague: Mouton & Co., 1967, .
11. Rumelhart, D. E., "Notes on a Schema for Stories", In Bobrow, D.G. and Collins, A., *Representation and Understanding*, New York: Academic Press Inc, 1975, pp. 211-236.
12. Schank, R. C., "The Structure of Episodes in Memory", In Bobrow, D.G. and Collins, A., *Representation and Understanding*, New York: Academic Press Inc, 1975, pp. 237-272.
13. Schank, R. C. and Abelson, R. P., *Scripts, Goals, Plans and Understanding*, Hillside, NJ: Lawrence Erlbaum, 1977.
14. Schank, R. C., Wilensky, R., Carbonell, J. G., Kolodner, J. L. and Hendler, J. A.,

"Representing Attitudes: Some Primitive States," Tech. report 143, Computer Science Department, Yale University, 1978.

15. Schmidt, C., Sridharan, N. and Goodson, J., "The Plan Recognition Problem," *Artificial Intelligence*, Vol. 11, No. 2, 1977, pp. 45-83.
16. Wilensky, R., *Understanding Goal-Based Stories*, PhD dissertation, Yale University, Sept. 1978.