

Volition and Advice: Suggesting Strategies for Fixing Problems In Social Situations *

Eric A. Domeshek
Institute for the Learning Sciences
Northwestern University
Evanston, Illinois 60201

Abstract

Just as an abstract causal analysis of a plan's faults can suggest repair strategies that will eliminate those faults [6], so too, an abstract causal account of how a problem arises in a social situation can suggest relevant advice to correct the problem. In the social world, most problems arise as results of agents' actions; the best way to fix such problems is to modify the behavior that produces the problem. The vocabulary of *volition* developed in this paper is proposed as an abstract level of motivational analysis useful for discriminating among strategies for changing behavior.

Volitional analysis focuses on the agents involved in an action. In addition to the *actor*, there is often a *motivator* agent who influences the actor and sometimes a *third-party* agent used as a tool by the motivator. If any of these agents can be swayed, the problematic action may be avoided. By identifying these agents and classifying the influences working on them, volitional analysis can suggest relevant modifications. The influences most often depend on the social context that links agents and establishes goal-generating themes. Behavior, however, is not always directly goal-governed, and volitional analysis recognizes these exceptional cases as well.

1 Problems in the Social Domain

Consider a situation where a man is seen in the company of a woman; he has a wife, but this isn't her. Imagine his wife is the one who sees him with this other woman. If she considers this a problem, what should she do? The answer of course depends on *why* her husband was in the company of this other woman. Her response depends on "*why*" in at least two senses: the first is whether the reason for the observed action, and therefore its meaning, actually signals any threat to her marriage; the second

is that in order to change the behavior, her response should attack its causes. This paper is concerned with the second issue: *the choice of behavior modification strategy*. A reasonable response would be quite different in each of the following circumstances:

1. He was in the middle of one of a long series of secret trysts;
2. He was far from home, lonely, and this woman caught his eye;
3. A friend asked him to keep his sister company;
4. His boss assigned him to entertain this client;
5. He was just holding the door for a random passerby.

Differences between these situations, range from the underlying goals, through the specific actions implied, to the likely effects on the marriage. When the task is to eliminate a problem, however, we can focus on the causes of that problem. To change problematic behavior in the social domain, we can focus on why the actor exhibited the behavior and try to alter that particular causal chain. *Motivational analysis* — the construction of causal explanations for agents' behavior by appeal to goals, plans and other intentional constructs — is a complex and much-studied problem, although within AI, interest seems to have been restricted to the Natural Language Processing community [9, 1, 13, 5, 10].

This paper proposes and justifies a new vocabulary for summarizing complex motivations: *volitional analysis*. The point of this paper is to argue that the vocabulary of volition is particularly useful for discriminating among possible strategies for modifying behavior. The notion of functionally justifying a representational vocabulary by arguing for its fit to some task, (here, counter-planning in the social domain), is discussed in the next section.

This analysis of volition derives from work on the ABBY case-based lovelorn advising system [4]. ABBY's case library is composed of fixed advice packets; when an input problem situation is described in a way that matches the label on a piece of advice, ABBY retrieves the advice and offers it to the user. An extension to the system, currently

*This research was supported in part by the Air Force Office of Scientific Research (AFOSR). The Institute for the Learning Sciences was established in 1989 with the support of Andersen Consulting, part of the Arthur Andersen Worldwide Organization.

being implemented, will retrieve stored advice for inputs whose descriptions include assessments of volition. Sensitivity to volition will help ensure that ABBY's chosen advice will reflect relevant behavior modification strategies.

2 Plan Modification in the Social Domain

Planning systems in AI have historically focused on synthesizing complete and correct plans for specific tasks, building these plans out of simple actions and discarding the results after execution[2]. A newer paradigm — case-based planning [8, 7, 6, 3, 12] — suggests that, when possible, plans are constructed starting from complete solutions to old problems. An important step in case-based planning is *modifying* the old plan to fit the new situation.

To fix a proposed plan, one must be able to characterize what is wrong with it in a way that suggests useful repairs. Hammond [6] has pointed out the importance of understanding the *causal mechanisms* that lead to problems in plans. His CHEF system uses causal analysis to produce abstract descriptions of problem mechanisms which serve as indices to clusters of repair strategies. He demonstrated a class of abstract problem descriptions, derivable from detected problems, and useful for suggesting relevant repairs.

The insight that causal analysis of problem situations can suggest relevant repair strategies applies just as well to *social situations* as to problems centered on physical causality: if we know the social mechanisms underlying a problem, then we can focus on changing aspects of the situation that lead to that problem. Seifert, for example [11], capitalizes on the structure of *mutual goals*, which commonly arise in social situations, to suggest planning strategies.

The sorts of *advice* people offer one another about their everyday social lives can often be viewed as strategies for repairing faulty plans. In giving advice, then, our choice of advice should benefit from sensitivity to causal analysis. In the social world, problems derive from the effects of agents' actions. An especially relevant sort of causality is the mechanisms that determine why agents do what they do. The modification strategies suggested by volitional analysis aim to fix problems by modifying behavior, rather than by changing circumstances to make the same old behavior yield different (unproblematic) results.

Volition analysis is not a substitute for full motivational analysis; the modification strategies it licenses are not substitutes for detailed planning. Volition is a summary vocabulary designed to highlight important behavioral influences and thus suggest ways to effect changes. This paper does not address the difficult issue of how to *generate* these volitional

descriptions of actors' relationships to their actions, nor does it demonstrate how to apply the strategies suggested by volitional analysis to produce specific modifications and finished plans. The former is beyond the scope of this paper; the latter is beyond the scope of the ABBY project.

3 Social Causality

Traditionally, in both Naive Psychology and NLP research, we explain individuals' actions by appeal to *goals*. If someone eats, it is likely because they were hungry. More sophisticated analyses recognize the importance of *goal relationships* [13] and posit still higher motivational entities called *themes* [9]. A theme is a relatively persistent property of an agent that functions as a goal-generator; many themes follow from relationships between agents. We explain a choice to *go out to dinner* by a conjunction of the desires to eat and to socialize (a positive goal relationship). We recognize a choice of an *expensive romantic restaurant* as deriving from the specific nature of a romantic relationship, (an interpersonal theme).

If you want to change this behavior, either because it is problematic for the planner, (requiring plan-modification), or because it is problematic for a bystander, (requiring counter-planning aimed at getting the planner to modify his plan), then knowledge of the underlying goals offers some leverage. In the simplest case, acknowledging the operative goals but pointing out another mechanism for achieving them may be an effective modification strategy. More interestingly, pointing out *other* interacting goals may suggest that the action oughtn't be performed, that the original goals ought to be pursued in some other way, or that there is some better method that will achieve still more goals. Going out to a fancy restaurant may take a big bite out of savings being accumulated to finance a vacation; perhaps, thinking of vacations, you've been wanting to get out of the city; maybe a romantic picnic in the countryside would do better.

In a social context, personal goals are not the only mechanisms that cause agents to act. The romantic interpersonal theme introduced specific interpersonal goals affecting the choice of plan, (dinner together at a romantic restaurant). Alternately we could view this as a goal of the group entity — the "relationship" — which is adopted by an individual member of the group. This sort of analysis is clearer in the case of larger groups like families: the Jones family decides to have a reunion, and many of the individual members figure out how to get to California over Christmas.

Viewing the *group* as an agent with its own goals, the transmission of goals from groups to individuals can be viewed as a subclass of the general phenomena identified by Schank and Abelson [9] as *agency*: getting someone to do something for you. They

proposed a **D**-agency goal and accompanying persuade plans specifying how one agent can influence another to take some action.

Among the standard methods of persuading is **invoke-theme**. **Invoke-theme** might engage a personal theme, as when a charity plays on someone's self-image as a good liberal. Alternately, the theme in question might be the more specific relationship between the solicitor and the target, as when a college asks its alumni for support based on their teary attachment to the old alma mater. The Jones family example hinged on such group membership.

Persuades that establish agency are often important links in the causality accounting for why agents do what they do in social situations. The persuade plans presented in [9] included: **ask**, **invoke-theme**, **inform-reason**, **bargain-object**, **bargain-favor** and **threaten**. We incorporate this fragment of goal/plan motivational analysis into the vocabulary of volition because it indicates when there is a second agent involved in causing an action and because the different forms of persuasion are susceptible to different forms of attack.

4 Types of Volition

In designing a vocabulary for volition, we seek to ensure that it captures distinctions that matter when choosing behavior modifying strategies. We can classify volitions along several dimensions:

- **Source:** The *source* of the impetus to act;
- **Influence:** How the **actor** was *influenced* to act;
- **Choice:** The **actor's** degree of *choice*.

4.1 Source of Impetus

The basic question here is: did the initial impetus to perform the action come from the **actor** (most often in response to one of his own goals), or did it come from someone else? Again, the point is to identify the agents responsible for the action so we can choose strategies to change it. In the first case we consider the action to have been performed under **internal** volition; in the latter, under **external** volition.¹

In the case of *internal* volition there are several other questions to ask. When the law seeks to assess blameworthiness it asks whether the action was **premeditated** or **spontaneous**. If premeditated, we can ask whether the particular action was thoughtfully chosen from among alternatives or whether it was simply the **default** option, adopted without thought or without knowledge of alternatives. If

¹This paper uses boldface type for representational vocabulary items and for the several agent roles identified in volitional analysis (**actor**, **motivator** and **third-party**).

spontaneous, we recognize several sub-classes. Some actions are done for **emotional** reasons that have little to do with rational goal-pursuit; again, the legal world offers a similar distinction in recognizing "crimes of passion." Many actions are thoroughly scripted. It is so **conventional** a part of the normal morning routine to eat breakfast that you need not really think about it. Eating cold cereal for breakfast may be a personally **habitual** routine; again, no thought or decision is required, so you may find yourself eating cereal even on mornings when you don't really want to.

For the most part, agents do things in response to goals — their own or those of others. Sometimes, however, agents are involved in actions for reasons that have nothing at all to do with goals. The prototypical case is **uncontrolled** actions such as sneezing or falling down stairs. Another odd case is when agents do actions **unknowingly**, in the sense of not realizing alternate interpretations of the action; sitting down in the presence of a king may constitute "lese majesty." Finally, agents may *appear* to have done something, but actually have not: they may be **uninvolved** and merely implicated by circumstances. All of these odd cases indicate actions, or construals of actions, that are essentially *non-volitional*.

4.2 Mechanisms of Influence

External volitions are distinguished by the presence of some other agent: the *motivator*. In the case of external volition, the central questions distinguishing among different cases center on how the **motivator** manages to influence the behavior of the **actor**. The persuasion plans mentioned earlier are one set of mechanisms, each suggesting different behavior modification strategies. Consider the differences between changing the behavior of an **actor** who has merely been **asked** to do something, versus one who has been **inspired** or **convinced**. It takes different countermeasures to overcome these varying degrees and sources of commitment to an action.

Many of the more effective plans for persuasion depend on the use of *inducements*, which are defined in terms of the **actor's** goals. As used here inducements may be actual or future, positive or negative; thus they include promises and rewards, threats and punishments. For example, use of the **invoke-theme** plan plays on the notion of **obligation**, and introduces a whole raft of implicit inducements; themes generally subsume many goals, so invoking the theme serves to remind the **actor** of the benefits he can expect from complying, (and thereby maintaining the theme), and also of the loss he can expect if, in refusing, he drives the **motivator** to disrupt the theme. Actions **bought** with money, services or material goods offer the clearest case of inducements. Actions performed in response to arbitrary **threats** illustrate the effectiveness of negative

inducements.

Several other special classes of external volition are worth distinguishing. In most sorts of external volition, the **actor** ends up performing the action because another agent manages to tie it to some goal the **actor** cares about. There are however the external analogs of **uncontrolled**, **unknowing** and **uninvolved** actions defined in the previous section. An **actor** can be **compelled** to do an action — he can be physically manipulated and thereby forced to do, or not to do, almost anything. An **actor** can be intentionally **misled** as to what he is doing; another agent can tell him he is invited to a party when he will actually be crashing it. An **actor** can be **framed** — another agent can intentionally arrange things to appear as though the agent performed some action though he has not.

A final broad class of external volitions are best thought of as **third-party** volitions. These introduce a third agent into the causation of an action. The **third-party** may be recruited as a **surrogate motivator**, or the **motivator** may involve the **third-party** either to **administer** or to **receive** inducements. Examples of **third-party** as surrogate include getting someone influential to make **appeal** to the **actor**, or invoking external authority to **police**, and thus compel, action. An example of getting a **third-party** to administer inducements is **blackmail**: the **motivator** threatens to do something that will cause another agent to make problems for the **actor**. An example of **third-party** as recipient is a **hostage** situation: the **motivator** threatens to punish the **third-party**.

To clarify the various roles in these situations we need a linguistic distinction which does not exist in common usage: an action that is motivated externally will be the result of the someone *inciting* the **actor**, often by setting up an *inducement*. The incitement is the communication intended to influence the **actor**. The inducement is an effect on some goal intended to give teeth to the incitement. The point of this distinction is that different agents can deliver the incitement and the inducement. In the case of **blackmail**, for instance, the prototypical situation involves a blackmailer inciting an **actor** to do something by threatening to reveal some fact. But revealing a fact is not an inducement; it is the role of some **third-party** to respond to the revealed fact and actually deliver the negative consequences implied in the original threat. This **third-party** is effectively delivering the inducement.

4.3 Degree of Choice

Degree of choice is an issue when it comes to assessing credit or blame for an action. This way of classifying volitions works with the internal/external distinction to focus attention on particular agents, and with the varieties of influence to mitigate strategies that rely on negative ways of changing behavior. As

Internal	External	Third-Party
Premeditated	Persuaded	Surrogate
Chosen	Asked	Appeal
Defaulted	Convinced	Police
Spontaneous	Inspired	Administer
Emotional	Obliged	Blackmail
Conventional	Bought	Receive
Habitual	Threatened	Hostage
Uncontrolled	Compelled	
Unknowing	Misled	
Uninvolved	Framed	

Figure 1: The Vocabulary of Volition

in the case of legal distinctions cited earlier: it is considered less justified to execute someone for an action they didn't intend to do, or had no choice about doing.

The vocabulary of volition is outlined in Figure 1. It is intended to be systematic and representative, not necessarily exhaustive. Having introduced all these categories, we must now show how they capture differences in the applicability of strategies for plan modification. In using causal analysis to discriminate among plan repair strategies CHEF capitalizes on the insight that knowledge of how a state was caused suggests ways to disrupt the causal chain and eliminate the effect. In ABBY, the same basic principle translates as: *knowing how an agent came to do an action allows us to work out ways to get him not to do it*.

In both cases, we need to abstract from the details of any particular causal chain, fitting the specific instance to a set of categories that suggest relevant modification strategies. The volition types listed here are those abstract categories. Identifying the degree of choice an **actor** had and the internal goals or external influences that prompted the action tell us where our points of leverage may be in getting the **actor** to stop doing something, do it differently, undo what they've done, or not do it again.

5 Behavior Modification Strategies

The mapping from volition to behavior modification strategies is based on the causal model underlying each volition. This section presents a series of tables sketching the causality underlying each type of volition and showing how each causally significant fact suggests ways to change the behavior. The left column contains the causally relevant facts; the right contains strategies that seek to change those facts (and thus the resulting behavior).

For example, Figure 2 illustrates the differences along the “sources” dimension. Distinguishing internal from external volition and noticing involvement of other agents focuses attention on those agents who help cause problematic actions: the *actor*, the *motivator*, and the *third-party*. Only when you have assessed an **actor**'s volition with respect to a prob-

	MODEL	STRATEGY
Internal	Actor does Act	Change Actor's behavior
External	Actor does Act	Change Actor's behavior
	Motivator influences Actor	Change Motivator's influence on Actor
Third-Party	Actor does Act	Change Actor's behavior
	Motivator affects Third-Party	Change Motivator's effect on Third-Party
	Third-Party influences Actor's behavior	Change Third-Party's influence on Actor

Figure 2: The Three Broad Classes of Volition

	MODEL	STRATEGY
Premeditated	Actor has Goal	Get Actor to drop goal
	Actor does Act as part of Plan	Get Actor to do different Act for Plan
Chosen	Actor considers possible Plans for Goal	Introduce new Plans into consideration Remove Plan from consideration
	Actor adopts Plan for Goal	Get Actor to adopt other plan
Defaulted	Actor retrieves normal Plan for Goal	Tell Actor about new Plan
	Actor adopts normal Plan for Goal	Get Actor to adopt other plan
Spontaneous	Actor in Situation	Change or avoid situation
	— elicits — Actor does Act	Break "Situation-->Act" response
Emotional	Situation evokes Emotion for Actor	Modify emotion felt by Actor
	— elicits — Actor does Act	Break "Emotion-->Act" response
Conventional	Actor in Situation repeatedly	Keep agent from Situation
	— elicits — Actor does Act	Break "Situation-->Act" response
	Actor believes Act is socially common or acceptable	Change Actor's belief in social convention
Habitual	Actor in Situation repeatedly	Keep agent from Situation
	— elicits — Actor does Act	Break "Situation-->Act" response

Figure 3: Models and Strategies for Internal Volitions

lematic action as external, and identified some agent as **motivator**, can you apply modification strategies that seek to influence the **motivator** so he changes his demands on the **actor**. For internal volitions you must focus on the **actor** alone, (although you can, of course, in your counter-planning, act as **motivator** and perhaps introduce a **third-party**). Only when volition assessment identifies a **third-party** can you design interventions that modify the **third-party's**

role and its effect on the **actor's** behavior.

Figure 3 shows each of the internal volitions. They should all be read as specializations of the single table for internal volitions in Figure 2. These volitions illustrate differences on both the dimensions of "influence" and "choice." Lack of choice can result from lack of knowledge; in that case, we can augment the **actor's** knowledge. If the only way John knows to discipline his son is to beat him

MODEL	STRATEGY
Uncontrolled	
Actor in Situation	Avoid situation
— results-in — Actor does Act	Change situation (disenable causation)
Unknowing	
Actor missing information about Act	Inform Actor about Act
Uninvolved	
Actor did not actually do Act	Find out who (if anyone) did
Compelled	
Motivator do some Act	Prevent Motivator's Act
— results-in — Actor do Act	Change Motivator's Act (disenable causation)
Misled	
Motivator do some Act	Prevent Motivator's Act
— results-in — Actor misled	Change Motivator's Act (disable causation)
Actor missing information about Act	Inform Actor about Act
Framed	
Motivator do some Act	Prevent Motivator's Act
— results-in — Advisor misled	Change Motivator's Act (disable causation)
Actor did not actually do Act	Find out who did (suspect Motivator)

Figure 4: Models and Strategies for Non-Volitional Volitions

silly, we might suggest he try restricting privileges instead. Similarly, if Mary knows about mass transit, but chooses to drive to work because she doesn't know about the bus that runs near her house, we can open up new options by giving new specific information instead of giving a completely new plan.

Figure 4 shows the causal models and modification strategies underlying all the unconventional, "non-volitional" volitions. These reflect a total lack of choice on the **actor's** part. Obviously, if there was no choice about the action then it makes no sense to try and change it by appeal to the **actor's** intentional mechanisms: we can't offer an alternate goal, we can't threaten or cajole. Instead we either have to change the circumstances or we have to work on the **motivator** if we can find one. You can't convince someone not to sneeze even if you want them to keep quiet; better to remove the dust. If a baby cries you probably have to look to the parent for a way to get it to stop.

In turning to the external volitions shown in Figure 5, the dimension of "influences" becomes most salient. For external volitions the *strength* of the influences generally increases as you progress through the list of persuasion plans. Simply asking is clearly the weakest. Irrational appeal implicit in **inspiration**, if accepted, will override the rational argumentation of **convince**. Self interest is generally the strongest persuasion. There is a tendency to engage more important goals, (and themes), or invoke stronger impacts as you progress from **obliged** to

bought to **threatened**. To influence any of the agents involved in causing an action you can try to establish any of the forms of external volition. Which ones will actually work, (and with what specific inducements), will depend on the details of the situation, but in general, escalating to a higher level of persuasion is a good strategy [9].

Beyond the general strategy of trying a stronger form of persuasion, many volition classes suggest relatively specific strategies. As the John and Mary examples illustrated, if you classify an **actor's** volition as **default**, you may be able to change the behavior simply by offering another option. If you believe the behavior resulted spontaneously from an **emotional** state, you may be able to set up a countervailing emotion. **Inspiration** may be counteracted by undermining the role-model. Actions performed in response to **threats** may be changed by offering protection. Actions resulting from **compulsion** can be changed by disabling the application of force. **Actors** who **misunderstand** what they are doing can be enlightened; those who are **misled** can be warned of the trickery.

Finally, Figure 6 shows the models and strategies for the third-party volitions. These are simply more specific versions of the external volitions. For example both **hostage** and **blackmail** situations can be viewed as a **threats**. But of course these situations offer more options because of the additional agent involved.

MODEL	STRATEGY
Persuade	
Motivator Persuades Actor:	Block transmission or receipt of Persuasion
“ Actor do Act for Inducement”	Change contents of transmitted Persuasion
Actor believes Inducement valid	Convince Actor Inducement is false Convince Actor Inducement will be blocked
Actor values Inducement	Convince Actor Inducement is unimportant
Actor values Inducement more than consequences of Act	Convince Actor Inducement is less important
Actor agrees to do Act due to Inducement	Use stronger Persuade than Motivator's to get Actor to refuse
— motivates — Actor does Act	Use stronger Persuade than Motivator's to get Actor to renege on agreement
Asked	
Convinced	
Actor believes arguments	Undermine Actor's belief in validity or truth of arguments
Inspired	
Actor respects Motivator	Undermine Actor's respect for Motivator
Obligated	
Actor believes he has unfulfilled obligation to Motivator	Convince Actor there is no obligation to Motivator Convince Actor obligation to Motivator is already fulfilled
Bought	
Threatened	

Figure 5: Models and Strategies for Persuade Volitions

MODEL	STRATEGY
Surrogate	
Motivator influences Third-Party	Change Motivator's influence on Third-Party
Third-Party influences Actor	Change Third-Party's influence on Actor
Administer	
Motivator Incites Actor	Block transmission or receipt of Incitement
Third-Party delivers Inducement	Block execution of Inducement
Receive	
Actor believes Third-Party values Inducement	Convince Actor that Third-Party indifferent to Inducement
Actor believes Third-Party values Inducement more than consequences of Act	Convince Actor that Third-Party thinks Inducement less important
Actor agrees to do Act due to Inducement	Offer extra inducement for Third-Party to get Actor to refuse
— motivates — Actor does Act	Offer extra inducement for Third-Party to get Actor to renege on agreement

Figure 6: Models and Strategies for Third-Party Volitions

6 Using the Strategies

We can now return to the problem that opened this paper and suggest advice for the wife seeking to reclaim her possibly errant husband. Each item of the following list offers instantiations of strategies for the corresponding situation described earlier. The bracketed boldface word is the volitional classification:

1. [**Premeditated**] The wife can threaten her husband with divorce. Alternately, she can try to buy or bully the mistress into surrendering her claim. Of course, she can use any of the third-party techniques, perhaps attempting to bring the mistress to heel through the offices of her husband, or family, or priest.
2. [**Emotional**] The wife can lobby that her husband not go on so many long business trips alone, or arrange that they stay in closer touch when he is on the road.
3. [**Asked**] The wife can simply ask her husband not to spend too much time hanging around with other women, even as a harmless favor to a friend.
4. [**Obliged**] The wife can try to override her husband's obligation, but she has a stake in not ruining his standing at work. She can try to influence the boss somehow, but again, she can't afford to threaten and probably doesn't have much to offer.
5. [**Uninvolved**] There is nothing to do in this case, since nothing really happened.

Notice that detailed advice depends on specifics of the plans and actions chosen to address active goals, and on the specific mechanisms that introduce problems. Experience with specific problems (and past solutions) may often provide surer results than general reasoning through abstract strategies. When available it is preferable to volitional analysis. Of course, if such specific advice is found to need modification in order to fit the current circumstances, we are back in the situation of a case-based planner looking for plan repair rules, and volition may have a role to play.

The classification of volitions proposed in this paper is intended to serve as a useful abstraction summarizing the causation of actions in a way that discriminates among strategies for modifying those actions. Application of the strategies may not be possible without understanding the full motivation underlying these situations. The ABBY system has a model of the social domain that includes knowledge of social units, interpersonal themes, resulting goals, relevant plans, specific social actions, and the effects of those actions on identified goals. Volitional analysis is just one influence on its selection of advice, but this analysis, and the strategies it picks out, hold an important place in the arsenal of case-based reasoning techniques for the social domain.

Acknowledgements

Thanks to Eric Jones for many useful discussions and comments on several drafts of this paper. Thanks also to Kemi Jona for his comments on this paper.

References

- [1] J.G. Carbonell. *Subjective Understanding: Computer Models of Belief Systems*. PhD thesis, Yale University, 1979.
- [2] E. Charniak and D McDermott. *Artificial Intelligence*. Addison Wesley, Reading, MA, 1985.
- [3] C.C. Collins. *Plan Creation: Using Strategies as Blueprints*. PhD thesis, Yale University, 1987.
- [4] E.A. Domeshek. Understanding stories in their social context. In *Proceedings of the Tenth Annual Conference of the Cognitive Science Society*, Montreal, August 1988. COGSCI-88.
- [5] G.D. Dyer. *In-depth Understanding: A Computer Model of Integrated Processing For Narrative Comprehension*. PhD thesis, Yale University, 1982.
- [6] K.J. Hammond. *Case-based Planning: An Integrated Theory of Planning, Learning and Memory*. PhD thesis, Yale University, 1986.
- [7] J. Kolodner. Extending problem solver capabilities through case-based inference. In *Proceedings of the Fourth International Workshop on Machine Learning*, pages 167-178, Los Altos, CA, June 1987. University of California, Irvine, Morgan Kaufman Publishers, Inc.
- [8] J.L. Kolodner, R.L. Simpson, and K. Sycra-Cyranski. A process model of case-based reasoning in problem-solving. In *Proceedings of the Ninth International Joint Conference on Artificial Intelligence*, Los Angeles, CA., August 1985. IJCAI.
- [9] R.C. Schank and R. Abelson. *Scripts, Plans, Goals, and Understanding*. Earlbaum, Hillsdale, N.J., 1977.
- [10] C.M. Seifert. *Mental Representations of Social Knowledge: A Computational Approach to Reasoning About Relationships*. PhD thesis, Yale University, 1987.
- [11] C.M. Seifert. Planning principles specific to mutual goals. In *Proceedings of the Ninth Annual Conference of the Cognitive Science Society*, Seattle, August 1987. COGSCI-87.
- [12] R.G. Simmons. *Combining Associational and Causal Reasoning to Solve Interpretation and Planning Problems*. PhD thesis, MIT, 1988.
- [13] R. Wilensky. *Planing and Understanding*. Addison-Wesley, Reading, MA, 1983.