

Understanding Stories in their Social Context

Eric Domeshek

Yale University
Department of Computer Science
New Haven, CT 06520-2158

Abstract

Stories concerning multiple agents with interacting goals and plans are difficult to understand; the task can be simplified however, if a program is given sufficient knowledge of social structures. Representations of social aspects of the story may also be necessary components of a satisfactory understanding of the story; here, we consider the distinctions which must be representable to support the task of *advice-giving* in the social-domain.

This paper elaborates on established goal taxonomies in order to capture distinctions among goals embedded in a social context; these distinctions serve both as the basis for choosing advice and as inferential shortcuts. It goes on to explore ways in which such social goals can be predicted from a detailed understanding of the conventions defining social structures linking the various agents. *Social-units*, *social-situations*, and *triangles* are introduced as conceptual structures which organize interpersonal-themes. The goals predicted by these themes provide a focus for motivational and impact inferences which would otherwise be lacking. These and related structures also allow some direct predictions of actions, as when a social-situation provides scriptal specification of action sequences or when a *contract* underlying a social-unit licenses specific *recourses* in response to obligation failures.

I. ABBY the Advice Giver

Story understanding has progressed from understanding scripty stories [2], [3], to simple plan based

sequences [14], to more complex stories of multiple agents with interacting plans and goals. Relatively little attention has been paid to representing the details of the social context in which the agents act out the story; typically, programs either assume a fixed cast of characters with fixed goal structures [1], or limit the range of possible relationships among agents [4]. The first approach is an unrealistic simplification in most cases; the second throws away a potentially rich source of expectations.

Embedding agents in a social context should actually simplify the task of understanding by providing expectations about likely goals and plans; so long as such expectations are not consistently wrong, they will serve to effectively restrict the amount of inference required to make sense of the story. Furthermore, explicit representation of the story's social context may actually form a necessary part of a satisfactory understanding for some purpose [11].

This paper reports work done in the context of the ABBY project. ABBY is a story-telling advice-giver for the social domain. Named in honor of the legendary Abigail VanBeuren of the syndicated newspaper column *Dear Abby*, ABBY reads¹ stories describing lovelorn situations and responds with advice couched in the form of a story. Essentially, ABBY is half of a case-based planning system [9], [5], [6] — modeling memory, problem understanding, and case retrieval. Its ability to find relevant cases in memory depends on a rich representation both of the cases in memory and of the input case.

¹Lacking a parser, the current implementation of ABBY does not actually *read* anything, rather it is given a representation of the input in terms of such abstract categories as could easily be extracted from a natural language version of the story. That is to say *marriage of a man and a woman* as opposed to *rushed marriage of frivolous high-school sweethearts*.

This research was supported in part by the Advanced Research Projects Agency of the Department of Defense and monitored by the Office of Naval Research under contract N00014-85-K-0108.

The program operates by viewing any new input as being like some other situation it has already encountered; it sorts inputs into the categories defined by its experiences. When ABBY tells a story as advice, it is not simply because a story is a good way to communicate a point (which it is), but because that past experience contributed significantly to its understanding of the input.

Consider this example of a story handled by ABBY:

Dear ABBY,

My boyfriend, Joe, and I have been going together for 3 years, and we are just about to graduate from high-school. We love each other more than anything, and we want to get married. My parents, though, say we should wait until Joe gets some kind of job and saves up enough money. We think that since we love each other we should be together. Jobs and money will work themselves out. How can we convince our parents?

- Already Ready

Appropriate advice for this input could come from a reminding of another case of rushed marriage of high-school sweethearts.

Dear Already Ready,

I knew a couple, Brenda and Eddie, who rushed into marriage as soon as they got out of high school.² They had been going together for years and were the most popular couple in their class. Even so, everyone told them they would never make it in marriage. All they knew about was the fun they had together. Neither one had ever had to work for any more than pocket money; earning enough to buy gas and malteds isn't the same as supporting a household. They learned that fast enough. Their marriage lasted only a few months — just about as long as their savings.

- ABBY

ABBY must be judged on the quality of its reminding. Within the paradigm of case-based rea-

²This story is actually adapted from a song by Billy Joel, *Scenes from an Italian Restaurant* off his album *The Stranger*. Though I don't think I've listened to it in years, in the context of this work, I've lately had trouble keeping it out of my head; I mention this to credit Billy Joel, and as evidence of how memorable a good story can be.

soning, expertise results from massive experience; a memory stocked with problems carrying observed or generated solutions provides a reasonable guess at a good plan in any new situation. Though the current implementation has but on the order of a hundred cases in memory, and is only planned to contain on the order of a thousand cases, the intent is to model a human expert's memory with its tens or hundreds of thousands of cases. If we imagine a system with a memory as rich in experience as the real "Dear Abby", or even as the average adult, we can begin to ask what features would differentiate one experience from another. What makes one episode a better reminding than another? This raises the question of representation. If one reminding is really better advice than another, then the differences that make the difference must be representable.

ABBY's representational ontology includes many of the conceptual classes introduced by earlier work in story understanding [10], [1],[8],[4],[9]. There are individuals representing *states*, *actions*, *things*, *agents*, *themes*, *goals*, *preferences*, *affects*, plus concepts for *temporal* and *causal* linkages, as well as larger aggregate structures such as *MOPs*, *TOPs*, and *XP*s. In keeping with ABBY's focus on classification, each of these broad classes dominates a detailed hierarchy of specialized variants.

These representations must facilitate detailed analysis of the story's facts in terms of their impacts on characters' goals, with an emphasis on instances of goal failure; failed and threatened goals are the problems about which ABBY gives advice. The major difficulty for ABBY is that the space of possible goals is nearly unlimited, and the linkage from reported facts to those goals may be arbitrarily indirect; with enough ingenuity, nearly any fact can be shown to impact nearly any goal. The search for relevant goals — those goals that either motivated story actions, fortuitously benefited from those actions, or suffered because of them — must be limited.

Limiting this search requires predictive knowledge of what types of goals the characters are likely to actually hold, and what types of goals are most likely to be affected by reported and inferred facts. This paper focuses on some of the representational machinery developed to supply these expectations. Following [10], ABBY uses several classes of *themes* as explanatory and predictive organizers of goals. After elaborating a taxonomy of goals specifically

for social contexts, this paper discusses how ABBY derives agents' themes from membership in *social-units* and ad-hoc groups, and from participation in *social-situations* and *triangles*.

II. Goals: What do People Want Anyway?

Much work has been done in elaborating goal taxonomies. Our point of departure is the taxonomy proposed in [10]. Among the distinctions made there, the most important for present purposes are those between *satisfaction*, *achievement*, and *preservation* goals.

Satisfaction goals are those that require repeated attention; they may be satiated at one time, but eventually — as originally proposed, on some biologically determined schedule — they will again demand attention. In ABBY, almost all goals are roughly of this form, subject to the generalization that we are no longer just talking about biological drives as for food, sleep, or sex. Under this broadened interpretation, almost all possession goals to fit here. For many things it is not enough to just possess them once; money for instance gets spent, and then you need to get more. This is contrary to the original classification in [10] which counts possession goals among the achievement goals. The prototypical achievement goal though is something like *wanting to marry someone*; that is, its object is a state that you achieve once and then maintain. This is the form in which achievement goals appear in ABBY.

Finally, preservation goals as presented in [10] are reactive goals — they come into being when an agent perceives a threat to some desired state. In ABBY, this notion is generalized by introducing *maintenance* goals. Maintenance goals always exist for a desired state; they account both for the spawning of P-goals in response to a threat, and for anticipatory plans aimed at maintaining the desired state. Consider the goal *maintain-health*; the standard plan of scheduling regular checkups with doctors and dentists is intended to maintain health without needing to experience an imminent threat.

A. Social Goals

In a social context, important goals include what one wants for another, what one can do for another, what one can get from another, and what a group can do together for some members, all members, or

non-members. If we start with some basic biological goal state such as *satiated-hunger* — *X want's X's hunger satiated* — social variants of this goal would include:

- X wants Y's hunger satiated
- X wants Y's hunger satiated by X
- X wants Y's hunger satiated by Y
- X wants Y's hunger satiated by Z
- X wants X's hunger satiated by Y

Such distinctions make sense when we consider beliefs such as: *parents ought to feed their children*. Now a child can go begging in the street and get something to eat — that will take care of the child getting fed — but it does not meet the original goal specification; we still hold that the parent has failed to meet an expected obligation. Considered from the child's perspective, this may simultaneously be a goal success and a goal failure: it didn't just want to eat, it wanted to be fed, and it wanted to be fed by its parents. These distinctions matter to the extent that failures of goals and obligations matter. In the social world they matter quite a bit; we would understand if the child held a grudge against its parents for this perceived failure, and if this had some future consequences.

Not surprisingly, the same social variants appear when dealing with overtly social goals, goals less directly coupled to biology. Consider the goal to *start dating*. It is important to notice that wanting to get a date is different from wanting to date someone in particular. An advice-giver must notice the difference between X wanting to start dating Y and X's mother wanting X to start dating Y. It is reasonable to consider how X feels about other agents' goals for him; a mother with such goals might be a particular type of mother that we know something about — perhaps we have specific advice for how to deal with pushy prying mothers mixing into their children's social lives.

Social goal variants can be further differentiated based on the relationships holding between the X's, Y's, and Z's. We have specific knowledge about the conditions under which certain goals are required, appropriate, expected or forbidden; those conditions often depend on the relationships holding between agents. Consider the mutual goal *having sex with someone*. If the agents involved are husband and wife this is a normal goal, in fact it is an obligation of marriage. If the pair are lovers it

may not be obligatory, but its certainly normal. If the agents are brother and sister — or in fact any close relation — the desire is considered taboo.

There are still other important distinctions among goals in social contexts. When X has a goal with respect to Y, we frequently can predict whether Y wants X to pursue that goal. Parents want their children to be fed; children want their parents to act on that goal. Such goals are *desired* goals; if X satisfies the goal with respect to Y, Y is also having a goal satisfied, if X fails, Y has a failure as well. Desired goals are a shorthand that facilitate noticing impacts, both positive and negative; knowing that a theme-partner was helped or hurt also aids in propogating affect inferences.

Other goals may be classed as *balanced* goals. These are paired goals where X has a goal with respect to Y, and Y has a similar goal with respect to X. As an example consider the goal *ensure-happiness* that normally holds between romantic partners. When ABBY notices X's goal is relevant, it is worth considering if Y's balancing goal is also relevant. For instance, when the goals are both balanced and desired, failure of one may, under the strategy of "tit for tat" lead to failure of the other; if X fails to make Y happy, Y may refrain from making X happy.

Closely related to balanced-desired goals are what Seifert has called *mutual goals* [11], that is, goals whose object states, by definition or by social convention, require the cooperative efforts of multiple agents. Examples include *spending time*, *conversing*, and *starting or maintaining a relationship*. Seifert has elaborated in great detail inferences and strategies specific to mutual goals [12].

B. Subsumption Goals

A large class of social goals can be viewed as goals for abstract possession. This is intended as a generalization of normal possession goals, so that we can speak not only of an agent wanting tangible items like food, clothing, or money, but also of intangibles like company, respect, and love. This notion of abstract possession — of *having* something — suggests the companion social goals of *giving* and *getting* these objects. Again, there are many possible social variants:

- X wants X to have O
- X wants Y to have O
- X wants X to get O

- X wants Y to get O
- X wants X to give O
- X wants Y to give O
- X wants X to give O to Y
- X wants Y to give O to X
- X wants Y to give O to Z

Possession of an object O is generally desired if O is instrumental to some other goal. Having food is a *goal subsumption state* that helps in achieving recurring goal to *satiated-hunger*. A state subsumes another goal-state if achieving the subsumption state makes it trivial to repeatedly achieve the goal-state [13]. In such a situation, it makes sense to pursue the subsumption-state as a goal in its own right. Thus *having-food* becomes a common goal even though milk in the refrigerator is doing no one any immediate good.

The number of possible tangible and abstract entities worth possessing is huge. This section only sketches the outlines of a potentially vast class of common goals; they form the basis for the claim that most of the goals ABBY is concerned with behave like satisfaction goals. Just as it's hard to make it through life on one week's pay, so too is it tough to get by on one unrepeated display of affection. Agents are constantly concerned with ensuring these possession goals get satisfied.

Just as achieving possession of some object can be a subsumption goal, so ensuring a reliable stream of that object — so that you always possess enough of it — can also become a subsumption goal [13]. This is the basis for ABBY's notion of *social-unit*: a long-term contractual arrangements with other agents designed to provide needs on a continuing basis. Joining a company as an employee is the most obvious example of entry into a social-unit serving to subsume possession goals. Less tangible possessions accrue from becoming someone's friend, but companionship, support, and advice are valuable nonetheless. Of course there are social-units that are not usefully explained this way; for example it does not account for why the average conscript joins the army.

Social-units then can be viewed as plans for subsuming common goals. Adopting those plans engenders other goals; here we hark back to the earlier analysis of achievement goals. People spend much of their time pursuing the *initiation*, *maintenance*,

termination, and *restoration* of social-units, or trying to *join* and *quit* other larger social-units. All these (except maintenance goals as described earlier) are achievement goals. In lovelorn situations, it is often these high-level achievement goals that are central to a story.

III. Social-Units and Contracts

We organize the social world by placing people into social-units. They are so pervasive and powerfully predictive that they or the roles and relationships that come with them are almost always explicitly mentioned when discussing people. Who is John? He is Mary's husband, Jimmy's father, Ethel and Eddie's son, an employee of International Widget, a devout church-goer, a good neighbor. That sort of capsule summary can immediately relate John to anyone of importance in his world.

It is interesting to note that for a wide range of social-units we have specific words or phrases to denote the units, the roles, and even the relationships between role-fillers. As suggested, these are very important concepts frequently worth communicating — they define the most important features of any individual's social landscape in powerful shorthand.

Social-units range in size from diads to cultures. On the smallest scale, we live in a world of *acquaintances*, *friendships*, and *romances*. In the midground we define ourselves with respect to our *families*, *neighborhoods*, and *workplaces*. In the larger world, we are members of *professions*, *religions*, and *nations*.

Each member plays some *role* in a social-unit; members relate to each other differently depending on their respective roles. Within a family there are roles for *parents* (*mother*, *father*), and *children* (*son*, *daughter*). The fillers of these roles are related to one another as *parent-of* (*mother-of*, *father-of*) or *child-of* (*son-of*, *daughter-of*).

We know quite a bit about what it means to be someone's parent. Such relationships function as *interpersonal-themes*; that is, they lead to specific goals of one agent involving the other. Every social-unit defines a set of interpersonal-themes. ABBY can use these to predict what social goals are likely to be active between agents. Normal parents want to provide material, social and emotional support

for their children. Children want those things provided to them by their parents.

As suggested before, some goals have a special force in that they are obligations. Even if there are parents who in some sense don't *want* to feed their children, we still expect them to act as if they hold the goal. Claiming that an obligation leads you to want what you don't want is not really a contradiction; goals, as representations, are primarily predictive and explanatory constructs. If we believe an agent has an obligation it is reasonable to try to see actions as attempts to meet the obligation; if we see someone acting in a way that satisfies an obligation it is reasonable to use the obligation as part of the explanation for the action. That is all that a goal means. We do however need to be able to notice when an obligation, though fulfilled is resented — when the goal state is not really wanted.

Reciprocal relationships inherent in social-units, such as parent and child, generally imply *contracts* between the parties; we might take a long-term contractual basis as one of the defining features of social-units. Contracts are what confer special status on obligations. As representational structures, contracts package the reciprocal interpersonal-themes and provide predictions when obligations are not met; contracts license *recourses* — legitimate forms of retribution for obligation failure. This accounts for why obligations may be effective goals even though the goal-state is not in itself desirable; in some cases it is only fear of such recourses that leads a spouse to visit with the in-laws.

Trappings of romance and true love notwithstanding, marriage is a case where a contractual analysis of the relationship is openly adopted in many societies. For example husband and wife are supposed to sleep with each other exclusively; should one violate that part of the contract, the other has recourse to the courts where they can sue for divorce and may be awarded alimony. Violating just about any aspect of the marriage contract leads to the same recourse; in some societies, failure to produce male sons is likewise legitimate grounds for divorce. Dissolution of the underlying social-unit that invoked the contract is the most blunt and sweeping, but also one of the most common recourses. Still there are others; spouses may seek counseling before resorting to divorce.

The prototypical contract is an agreement entered into freely by equals, in which each party

agrees to assume certain obligations such that both parties benefit in some way. Many social-units deviate from this ideal in some respects, yet are still usefully thought of as contracts. This is particularly true when the contract itself is put into a social context; the contract may be enforceable by — that is recourses may be administered through the agency of — a higher authority. We see traces of such outside influences even in relatively prototypical social-contracts like that implied by marriage.

The social contract between parent and child is a good example of one which deviates significantly from the contract prototype. In the first place, the child never asked to enter into that contract. Furthermore, the parent has absolute power over the child. This is similar to the situation of the conscript alluded to earlier. In both cases, enforcement of obligations of the stronger to the weaker party depends on an implicit contract with some higher authority — that of the state. Unfortunately for the conscript, the army and the state are usually the same thing (or worse, the army may be superior to the state).

IV. Scripts as Social-Situations

The notion of social-unit bears an interesting relationship to the older notion of *script* [10]. Many scripts are of a form that might reasonably be called *social-situations*; that is they are spatio-temporally bounded, socially prescribed sequences played out by agents, temporarily related by their entrance into a short-term contract. Along with scripts came the notion of *role-themes* — organizers for the goals agents adopt when they enter into their script roles. Viewed as social-situations, scripts should spawn interpersonal-themes, and in fact it turns out that in social scripts, the goals organized by role-themes are in fact social goals: one agent wants something from another, or wants to do something for another.

In a restaurant, the customer has a temporary contractual relationship to the server; they have goals and obligations with respect to one another, and there are sanctioned recourses should either violate parts of that contract. Beyond the specific actions and their specific sequencing, we know that the server is supposed to provide service in a timely and courteous manner, the customer is supposed to behave civilly and pay his bill. If the service is unsatisfactory, the customer may refuse to leave a tip. If a

customer is offensive or abusive, the server may call the manager to have the customer removed. Again, the parallel to the case of a social-unit with an unbalanced power relationship is striking; a younger child has the same sort of recourse to its parents' authority when frustrated by an older sibling's misbehavior.

What counts as a short-term social-situation and what as long-term social-unit will usually be clear enough — long-term relationships generally lack a pre-specified termination conditions (short of death). Business units might be an exception: employment is generally long-term, but often has well-specified termination conditions. The scriptiness of a social-situation inheres in the stable predictability of its settings, sequence, and specific actions.

Just as several social-units can link the same agents, so several social-situations can be simultaneously active, as can a mixture of social-units and social-situations. Any of these combinations can generate interacting goals. A classic social-unit interaction stems from taking a son into the family business, which forces father and son to become boss and employee. Going to a restaurant where one of your friends works thrusts one among equals into servant status — a social-situation impinges on a social-unit.

V. Efficiency and Adequacy

This paper has presented a variety of representational constructs for use in understanding stories from the social domain. It has argued for particular distinctions and particular classes of concepts on two grounds: efficiency and adequacy. Efficiency implies that having a particular representation makes expectations available which effectively focus subsequent processing. Adequacy implies that without noticing a certain distinction, the program cannot succeed at its task.

In the task of story understanding, the major source of intractability is the need to notice the import of facts. What motives lay behind any given fact? Which agents are affected by that fact and how are they affected? Answering these questions normally requires hypothesizing goals for the agents and using causal knowledge to determine the impact of the fact on those goals. Detailed knowledge of the social domain can simplify the task of understanding complex stories in which multiple agents hold interacting goals. Recognizing the social con-

text which relates characters, provides exactly the kind of expectations that are most needed: expectations about which goals are active and relevant. In addition, there are some situations in which social knowledge can provide even more specific predictions about likely action sequences.

Many of the distinctions introduced in the classification of goals serve not to speed processing, but to make it effective. To succeed in the task of advice giving, a system must be able to detect and represent some very subtle properties of situations. This paper pointed out some of the social variants of goals to which an advice giver must be sensitive. Noticing specific types of goals may hint at particular types of relationships. Specializing the system's understanding of the relationship between agents may tighten the predictions about what types of problems are likely to arise, or allow better predictions about what types of solutions will work — that is, it may lead to better advice.

This paper did not discuss the system's understanding of an agent's character, but specialized stereotype information can yield the same sorts of advantages. The same holds true with specialized types of goals, and interpretations of actions. Clusters of such specialized concepts hang and work together. A family may be like the Waltons: the mother may be a saint, the father stern but fair, their relationships to the children close and loving; then again, a family may be more like some nightmare version of the Bunkers: the mother a dizzy simpleton, the father a bigoted tyrant, and their relationships to the children based on manipulation and domination. People know about all these types and know which ones tend to go together; pinning down one, helps to pin down the other. Getting the whole picture suggests specific advice suited to the situation. The representations presented here form part of that sort of complete picture, and as such are a part of what ABBY must understand to give good advice.

References

- [1] J.G. Carbonell. *Subjective Understanding: Computer Models of Belief Systems*. PhD thesis, Yale University, 1979.
- [2] R. Cullingford. *Script Application: Computer Understanding of Newspaper Stories*. PhD thesis, Yale University, 1978.
- [3] G.F. DeJong. *Skimming Newspaper Stories by Computer*. PhD thesis, Yale University, 1979.
- [4] G.D. Dyer. *In-depth Understanding: A Computer Model of Integrated Processing For Narrative Comprehension*. PhD thesis, Yale University, 1982.
- [5] K.J. Hammond. *Case-based Planning: An Integrated Theory of Planning, Learning and Memory*. PhD thesis, Yale University, 1986.
- [6] J. Kolodner. Extending problem solver capabilities through case-based inference. In *Proceedings of the Fourth International Workshop on Machine Learning*, pages 167–178, University of California, Irvine, Morgan Kaufman Publishers, Inc., Los Altos, CA, June 1987.
- [7] N.I. Adams Rees, J.A. and J.R. Meehan. *The T Manual*. Computer Science Department, Yale University, 1984.
- [8] R.C. Schank. *Dynamic Memory*. Earlbaum, Hillsdale, N.J., 1982.
- [9] R.C. Schank. *Explanation Patterns: Understanding Mechanically and Creatively*. Lawrence Erlbaum Associates, Hillsdale, NJ, 1986.
- [10] R.C. Schank and R. Abelson. *Scripts, Plans, Goals, and Understanding*. Earlbaum, Hillsdale, N.J., 1977.
- [11] C.M. Seifert. *Mental Representations of Social Knowledge: A Computational Approach to Reasoning About Relationships*. PhD thesis, Yale University, 1987.
- [12] C.M. Seifert. Planning principles specific to mutual goals. In *Proceedings of the Ninth Annual Conference of the Cognitive Science Society*, COGSCI-87, Seattle, August 1987.
- [13] R. Wilensky. *Planing and Understanding*. Addison-Wesley, Reading, MA, 1983.
- [14] R. Wilensky. *Understanding Goal-Based Stories*. PhD thesis, Yale University, 1979.