Learning Planning Heuristic* through Observation

Charles Dolan
Artificial Intelligence Lab
Computer Science Department
UCLA
and
Hughes Aircraft Al Center

Michael Dyer
Artificial Intelligence Lab
Computer Science Department
UCLA

Abstract

This paper diacusses a method for learning thematic level structures, i.e. abstract plan/goal combinations, by observing the bad planning behavior of narrative characters. The learning method discussed is a one-trial, schema acquisition method, which is similar to DeJong's [DeJong, 1983]. The method uses constraint-based causal reasoning to construct a new schema which characterizes a situation. This work is part of the MORRIS project at UCLA [Dyer, 1983b].

1 Introduction

In the real world, tasks cannot always be accomplished by using simple sub-goal partitioning and recursive problem analysis. Both real world agents and narrative characters often must apply plans that require cooperation or information from other agents. A classification of real world plans is found in [Schank and Abelson, 1977]. A taxonomy of goal/plan interactions can be found in [Wilensky, 1978]. In addition, successful real world) planners often use adages to guide them in avoiding bad plans. Adages warn against both specific and general planning errors. *Poor Richard's Almanac* [Franklin, 1733-1758] gives many examples of adages, such as "A stitch in time saves nine" (error avoidance), and "Don't burn your bridges behind you." (error recovery).

Dyer [1983a] showed how a class of planning errors could be represented by Thematic Abstraction Units (TAUs), and how these planning errors might be recognized in stories. This paper will present a representation for planning error recognition that also facilitates the combination of planning descriptions into new thematic structures. The combination method requires an example narrative situation that contains a new planning error. The example is conceptually analyzed to discover whether known planning errors, present in the story, can be combined into new structures.

In addition to the representations for TAUs, goals, and plans mentioned above, the examples here also rely on Schank's Conceptual Dependency theory [Schank, 1972]. Other relevant work on memory organization includes [Schank, 1982], [Kolodner, 1980], and [Lebowitz, 1980].

2 An Example Planning Situation

Here we will see a planning situation that contains three planning errors whose descriptions are already known to the system. From this situation we will see how we can generate two specializations of planning errors, and one novel planning construct. This situation is an Aesop's fable *.

The Fox and the Crow

The Crow was sitting in the tree with a piece of cheese In her mouth. The Fox walked up to the bottom of the tree and said to the Crow, "Crow what a beautiful voice you have; please sing for me." The Crow was very flattered and began to sing. When she did, the cheese dropped out of her mouth. The Fox grabbed the cheese and ran away laughing.

This work was supported In part by a grant from the Hughes Aircraft Al Center and a grant from the Keck foundation. This work was also supported in part by the UCLA Center for Experimental Computer Science.

"Other work dealing with Aesop's fables is TALESPIN [Meehan, 1979] a program which generated stories by simulating a character's planning and problem-solving behavior. TALESPIN, however, lacked a theory of planning errors and had no sense of the moral of the story.

Note that this story can be looked at in two ways: 1) as an instance of bad planning on the part of the Crow and 2) as an instance of good planning on the part of the Fox. The is an instance of planning errors as counter planning as defined in [Carbonell, 1979].

The first of the three planning errors we will discuss is the simplest. When the Crow sings, she does not realize that she is already using her mouth to hold the cheese. This planning error is characterized at an abstract level by TAU-CONF-ENABL (confused enablement). The full representation for TAU-CONF-ENABL is given below.

In each TAU, the representation of a planning error consists of two parts: (1) the binding-spec: a list of conceptual patterns which occur in the story; (2) the constraints: a list of logical constraints among the patterns occurring in the binding-spac and other concepts from the story.

```
TAU name TAU-CONF-ENABL
  binding-spec
   [?standing-goal (P-GOAL actor ?x
                          obj (POSS-BY actor ?x
                                       obj ?y)
                          manner FAIL)]
    [?interfering-goal2 (GOAL actor ?x)]
    [?mistake (ACT actor ?x)]
  constraints
    intention(?interfering-goal,?plan),
    realization(?mistake,?plan),
    resulting(?mistake,?disabling-state),
    achievement(?desirable-state,?standing-goal1),
    disablement(?disabling-state,?desirable-state).
```

The abstract situation this structure characterizes is one where an agent has a goal, ?standing-goal, which has failed and where the goal was to preserve possession of some object. The cause of the goal failure is an act, ?miatake, which attempted to accomplish another goal, ?intarfaring-goal.

The representation presented here enables us to create new TAUs from existing ones. The processes of recognizing and indexing TAUs are covered more fully in [Dyer 1983] and [Dolan 1984]. [Dolan 1984] also covers the comprehension process model which allows the recognition of TAUs in this format when planning errors are encountered in stories.

As we mentioned above, The Fox and the Crow instantiates two other TAUs: (1) TAU-VANITY is the planning error of allowing personal vanity to dictate plan choice; (2) TAU-ULTERIOR is the planning error of not considering another agent's possible motives before acting. Both of these TAUs display an important characteristic of TAUs as planning heuristics; not only do TAUs provide admonitions against bad planning, but they can also be turned around and used as plans to try and force other agents into situations where they will make mistakes.

These TAUs can be combined to form new planning heuristics. There are two key problems in TAU acquisition: (1) How does a program know which TAUs to select and examine for combination attempts; and (2) Once selected, how are TAUs actually combined to form new planning and indexing structures?

Both(1) and (2) are non-trivial. A sophisticated planner will know many stories each indexed by multiple TAUs. Attempting to combine TAUs arbitrarily would lead to combinatoric problems. Fortunately, memorable stories (such as Aesop's fables) are designed to give novel planning advice through illustrating planning errors. Thus, TAU selection can be governed by the following strategy:

> WHENEVER two TAUs share concepts in an *observed* planning situation, TRY to combine them to form a novel planning construct

This heuristic can only be applied after reading a story. The comprehension of the story thus makes available TAUs for combination and indicates which concepts are shared.

There are two ways to combine TAUs based on the way they share concepts: (1) specialization and (2) combination (chunking).

Recent work in specialization learning includes [DeJong, 1983] and [Kolodner, 1980]. Both workers formulate methods for creating new planning knowledge through specialization. Most research in learning by chunking has been in domains where there is no counterplanning [Laird, 1984]. However there are some element of chunking in DeJong's work and in Mitchell's [Mitchell, 1983].

3 Creating New TAUa through Specialization

The type of sharing which exists between TAU-ULTEROR and TAU-CONF-ENABL is called containment. In order to see how containment works, examine the representation for TAU-ULTERIOR given below,

```
TAU name TAU-ULTERIOR
binding-spec
 [?standing-goal (GOAL actor ?x
                       status FAILED)]
 [?desirable-state (STATE)]
 [?mtrans (MTRANS actor ?y
                  to ?x
                  obj ?x)]
 [?mistake (ACT actor ?x)]
constraints
 achievement(?desirable-state,?standing-goal),
 motivation(?x,?interfering-goal),
 achievement(?mistake,?interfering-goal),
 disables(?mistake,?desirable-state),
 goal-conflict(?goal,(GOAL actor ?y
                          obj ?w)),
 achievement(?mistake,(GOAL actor ?y
                            obj ?w)),
 unexpected(?mtrans,?mistake).
```

TAU-ULTERIOR represents the situation in which ?y tells ?x information that motivates x to perform an act which results in the disablement of one of x's goals, while at the same time achieving a goal of ?y's which ?x did not forsee.

One constraint above needs some additional explanation. The last constraint, unexpected(?mtrans, ?mistake), states that there is no short causal inference chain from the information told (?mtrans) to the disabling act (?mistake). In the current implementation of this model a "short" inference chain is one formed by the application of a single inference rule.

By looking at the bind-spec for both TAUs the reader will see that TAU-ULTERIOR contains all the concepts from TAU-CONF-ENABL. Thus we can form a new TAU which is a specialization of TAU-ULTERIOR.

The new TAU is formed by taking the extra constraints from TAU-CONF-ENABL ("extra" meaning those not already in TAU-ULTERIOR) and conjoining them with those in TAU-ULTERIOR. This creates new planning advice specific to possession goals and ulterior motives. For more details on forming TAUs thourgh specialization the reader is directed to [Dolan and Dyer, 1985].

4 Creating New TAUa through Combination

As we saw above, we can get a non-trivial specialization of a TAU by discovering containment in a particular situation. In general, however, neither TAU contains the other. In these cases we must examine the relationships or constraints among the concepts not shared between the two TAUs. For an example, consider TAU-VANITY. The representation for this TAU is given below.

```
TAU name TAU-VANITY
binding-spec
 [?belief (KNOW actor ?x
                obj (APPRAISAL obj (BODY-PART owner ?x)
                                   value GOOD))]
 [?standing-goal (GOAL actor ?x
                       status FAILED)]
 [?interfering-goal (GOAL actor ?x)]
 [?mistake (ACT actor ?x)]
constraints
 intention(?interfering-goal,?plan),
 realization(?plan,?mistake),
 enablement((APPRAISAL obj (BODY-PART actor ?x)
                           value GOOD),
            ?mistake),
 thwarting(?mistake,?standing-goal).
```

TAU-VANITY represents a situation in which ?x believes he has a special skill* and is thus motivated to have a goal (of "showing off" in The Fox and Crow story) which will interfere with pre-existing goals.

TAU-ULTERIOR shares a number of concepts with TAU-VANITY. The concepts which the two TAUs share in The Fox and the Crow are, ?standing-goal, ?mistake, and ?interfering-goal.

A novel TAU we can learn by combining TAU-VANITY and TAU-ULTERIOR is TAU-SUCKERED. The representation for TAU-SUCKERED is given below.

```
TAU name TAU-SUCKERED
binding-spec
 [?standing-goal (GOAL actor ?x
                       status FAILED)]
 [?interfering-goal (GOAL actor ?x)]
 [?act (ACT actor ?y)]
 [?mistake (ACT actor ?x)]
constraints
 intention(?interfering-goal,?plan),
 realization(?plan,?mistake),
 thwarting(?mistake,?standing-goal),
 achievement(?mistake,?interfering-goal),
 enablement(?plan,?sub-goal),
 resulting(?act,?state),
 achievement(?state,?sub-goal),
 goal-conflict(?standing-goal,(GOAL actor ?y
                                    obj ?w))
 achievement(?mistake,(GOAL actor ?y
                            obj ?w))
 not-obvious-result(?act,?mistake).
```

TAU-SUCKERED embodies the planning failure of allowing someone else to take advantage of your dormant goals by providing one of the missing enablement conditions on that goal. In the case of The Fox and the Crow the dormant goal is the Crow's goal to show off. The missing enablement condition is a receptive audience. The Fox provides that audience and so tricks the Crow into defeating her standing goal of keeping the cheese.

Clearly there are many ways (structurally) of combining TAU-VANITY and TAU-ULTERIOR. Only one (or perhaps a few) will turn into an "interesting" TAU such as TAU-SUCKERED. If so, then how is ft that just the right concepts get combined? The process of combining TAU la called "justification". There are a sot of heuristics for justification given in next section.

Tha representation for "skil!" presented here ia greatly simplified.

In order to dynamically construct TAU-SUCKERED, there ara two major •taps. The first is to find the concepts which ara common to tha binding-specs of both TAU-VANITY and TAU-ULTERIOR, namaly. ?standing-goal, ?interfaring-goal, and ?mistake.

These concepts ara included in tha binding-spac for the new TAU, lines 2-4 and 8 in TAU-SUCKERED. Also, constraints that involve only these concepts and free variables are included in the constraint* part of the new TAU, lines 8-11 in TAU-SUCKERED. The second step is to take the remaining concepts from the binding-specs and justify them; that is, make sure they are accounted for either in the new binding-spec or in the constraints.

5 Juatlflcatlon

In this example the concepts which are not shared are ?dwsirable-state and ?belief from TAU-VANITY and ?mtrans from TAU-ULTERIOR. Justification takes a concept from one TAU and finds out where it fits in the causal structure for the instantiation of the other TAU.

First ?desirable-state is found to be already subsumed by constraint number 10 of TAU-SUCKERED. This is so because an ACT which thwarts a goal may negate a state that previously achieved the goal. In this case, the crow's singing thwarts her goal of keeping the cheese because it negates the state of having her mouth holding the cheese.

The achievement condition is found in constraint 10 of TAU-ULTERIOR. The heuristic which is used here is,

IF a constraint can possiblly subsume concepts,
TRY matching the constraints on subsumed concepts
against those in the component TAUs.

All that is left is to justify ?belief and ?mtrans. Intuitively what we want to see is a causal relation established between ?balief, the Crow thinking she has a good voice, and ?atrans, the Fox's flattery of the Crow. Further we want to see ?beiief as an enablement condition on the Crow's dormant goal to show-off.

One thing to note here is that the intention relation is transitive. The rule for intention transitivity is represented as in the figure below.

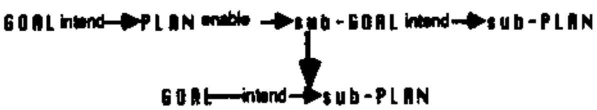

The system knows which relations are transitive, and can look up the particular kind of transitivity for a relation, i.e. what the intermediate relations are. Thus, another heuristic for justifying concepts is:

IF a transitivity rule waa applied in satisfying a constraint of TAU1'
TRY opening up the transitive path and look for concepts of TAU?
that need to be justified

In this case we see that the transitive path is in constraint (line 10) from TAU-VANITY which is reproduced as constraint (line 8) in TAU-SUCKERED, so a new constraint (line 12) is added to TAU-SUCKERED for sub-GOAL In the story this represents the fact that, in order to show off, a character must believe he haa something to show-off; the Fox flattering the Crow provides this enablement condition. The satisfaction of this enablement condition is given in lines 14-15 of TAU-SUCKERED.

The constraints in lines 14-18 are added because they meet various criteria for relating concepts in the binding-spec. For a detailed discussion of these criteria see [Dolan and Dyer, 1985].

The result of this constraint analysis is the creation of a new TAU, TAU-SUCKERED, which represents someone being fooled into having a goal failure by being motivated to satisfy another, currently dormant, goal. This TAU serves as a new indexing structure for Tha Fox and tha Crow story and is now available for use in future planning and comprehension tasks.

6 Progress and Future Work

A program. CRAM, is *under* development as part of this research. Currently, CRAM is able to understand stories that are input as unconnected Conceptual Dependency [Schank 1972] structures. CRAM finds the planning errors in each story and characterizes them in terms of one or more TAUs. These TAUs are then used to index the story in memory for later retrtval.

In the future CRAM will be able to give advice to correct characters' planning errors. Also planned for CRAM are a natural language parser and generator so that CRAM can take in stories as verbatim input and later generate English explanation of new TAUs it has discovered.

7 Conclusions

The approach presented here allows both specialization learning and chunking learning of planning errors in multiple planning agent domains. The structures learned can be used both for critiquing plans and also for generating counterplanning advice.

References

Carbonell, J. G., *Subjective Understanding:* a *Computer Model of Belief Systems,* TR 150 Yale CSD Ph.D Die. 1979.

DeJong, G., "Acquiring Schema Through Understanding and Generalizing Plans", in *Proc. of the Eighth Inter. Joint Conf. on Ai,* 1983.

Dolan, C. P., *Memory Based Processing for Cross Contextual Reminding: Reminding and Analogy Using Thematic Structures,* TR 850010 UCLA CSD, M.S. Thesis, 1984.

Dolan, C. P. and Dyer, M. G. "Encoding Planning Knowledge for Recognition, Construction and Learning", (Submitted for publication), 1985.

Dyer, M. G., *In-Depth Understanding: A Computer Model of Integrated Processing for Narrative Comprehension,* The MIT Press, Cambridge, Mass., 1983a.

Dyer, M. G., "Understanding Stories through Morals and Remindings", in *Proc. of the Eighth Inter. Joint Conf. on AI,* 1983b.

Franklin. B.. *Poor Richard: The Almanacks for the Years 1733-1758,* Heritage Press, New York, 1964.

Kolodner, J. L, *Retrieval and Organization Strategies in Conceptual Memory: A Computer Model,* TR 187, Yale CSD , Ph.D. Dis., 1980.

Laird, J. E., Rosen bloom, P. S., and Newell, A., Towards Chunking as a General Learning Mechnaism", in *Proc. of the Nat. Conf. on AI,* 1984.

Lebowitz, M., *Generalization and Memory in an Integrated understanding System,* TR 186. Yale CSD Ph.D. Dis., 1980.

Meehan, J., *The Metanovel: Writing Stories by Computer,* TR 74, Yale CSD, Ph.D. Dis., 1979.

Mitchell, T. M., Utgoff, P. E., Banerji, R., "Learning by Experimentation: Acquiring and Refining Problem-Solving Heuristics", in *Macine Leaning,* Michalski, Carbonell and Mitchel (Eds), Tioga, 1983.

Schank, R. C, "Conceptual Dependency: A Theory of Natural Language Understanding", *Cognitive Psychology,* Vol 3, No 4,1972.

Schank, R. C. and Abelson, R. P., *Scripts, Plans, Goals, and Understanding: An Inquiry into Human Knowledge Structures,* Lawrence Erlbaum Associates, Hillsdale, New Jersey, 1977.

Schank, R. C, *Dymanh Memory: A Theory of Reminding and Learning in Computers and People,* Cambridge University Press, Cambridge, 1982. CSD, Ph.D. Die., 1979.