
A Model of Planning, Action and Interpretation with Goal Reasoning

Michael T. Cox

MICHAEL.COX@WRIGHT.EDU

Wright State Research Institute, Wright State University, Dayton, OH 45435 USA

Abstract

Many algorithms have been presented in cognitive systems research for problem solving and planning. Given a goal, these algorithms search for solutions that achieve a goal state by interactions with some environment. However a major assumption is that goals are given, usually by a user directly as input or as part of the problem representation. Furthermore, once given, the goals do not change. Here we formalize the notion that goal formulation and goal change are themselves major parts of the problem-solving process. We include in our model not just planning and plan execution but also interpretation of the environment as plans are executed and exogenous events occur. We apply this model in a goal reasoning context.

1. Introduction

In virtually all AI systems, goal states are predefined and exogenously provided by an external user. But to have a continuing existence in time, agents must be flexible to survive and to continue to be useful. Recent work on *goal reasoning* (see Aha, Cox, & Munoz-Avila, 2013; Hawes, 2011) has started to examine how intelligent agents can reason about and generate their own goals instead of always depending upon a human user directly. Much of this previous work has been situated within the context of the automated planning community and has borrowed some of their formal notations as a theoretical framework. This paper will further extend the standard planning formalism to account for mechanisms of goal-reasoning and goal-driven autonomy and to begin to integrate the vocabulary of cognitive systems with that of the larger artificial intelligence field. The result integrates notations for planning, action, and interpretation within the scope of goal reasoning. The claim is that the result provides a uniform theoretical basis for expressing and communicating our work and explicitly highlights the contribution of the goal reasoning community. That is goal reasoning agents interpret a changing world and express their desires in terms of dynamic goals; they do more than just plan for and achieve a set given by an outside party.

As an example, consider a package delivery domain (e.g., see Figure 1) having the following characteristics. There is a network of locations connected by roads; and from time to time, vehicles may need to transport objects from one location to another as requests come in. Deliveries may be accomplished, for example, by picking up and delivering packages (e.g., in Figure 1, delivering pallets of bottles to a cola bottling plant). AI planning research has considered similar but simplified domains in which the world is static (i.e., no states change unless the agent performs an action), there is a given fixed goal to achieve (e.g., the delivery of certain specified packages), and the planning problem ends at any world state in which the goals are satisfied. In some sense, routine delivery of packages on a route is hardly a problem requiring intelligence. These problems omit any consideration of why those goals need to be achieved and what the agent should do after they have been achieved. In a more realistic model of a dynamically changing world, bottles need to be

delivered to the bottling plant as the inventory changes, and it might be necessary for an agent to explain and hence understand this and to generate a new goal when the inventory becomes unexpectedly low.

We begin this paper with a section that provides our formal representations for goal reasoning and shows how it frames the research using a simple blocks world example. The subsequent section examines the formalism within a kind of goal reasoning called goal-driven autonomy. This section is followed by related research. A brief summary concludes the article.

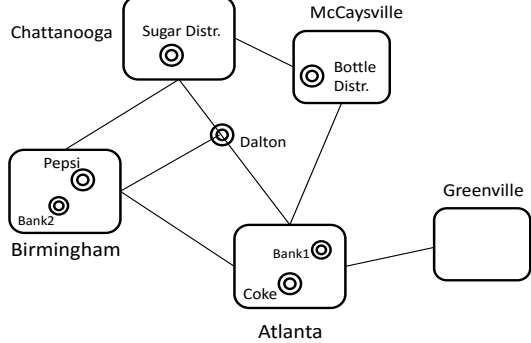


Figure 1. Simple package delivery domain.

2. A Formalism for Goal Reasoning

Much of the research in AI planning has focused on a restricted case called *classical planning*, in which all actions have deterministic effects, and the task is to generate a plan that reaches any of a predefined set of goal states.

2.1 Classical Planning Theory

A *classical planning domain* is typically defined (Ghallab, Nau, & Traverso, 2004) as a finite state-transition system in which each state $s \in S = \{s_1, \dots, s_n\}$ is a finite set of ground atoms of a function-free, first-order language \mathcal{L} . A *planning operator* is a triple $o = (\text{head}(o), \text{pre}(o), \text{eff}(o)) \in O$, where $\text{pre}(o)$ and $\text{eff}(o)$ (o 's *preconditions* and *effects*, respectively) are sets of *literals* (logical atoms and negated logical atoms), and $\text{head}(o)$ is a syntactic term of the form $\text{name}(\text{args})$, where name is the operator's name and args is a list of the variables in $\text{pre}(o)$ and $\text{eff}(o)$. Each *action* is a ground instance of a planning operator.

An action $\alpha \in A$ is *executable* in a state s if $s \models \text{pre}(\alpha)$, in which case the resulting state is $(s - \text{eff}^-(\alpha)) \cup \text{eff}^+(\alpha)$, where $\text{eff}^+(\alpha)$ and $\text{eff}^-(\alpha)$ are the atoms and negated atoms, respectively, in $\text{eff}(\alpha)$. A plan $\pi = \langle \alpha_1 \alpha_2 \dots \alpha_n \rangle$ is executable in s if each α_i is executable in the state produced by α_{i-1} .

For a classical planning domain, the *state-transition system* is a tuple $\Sigma = (S, A, \gamma)$, where S is the set of all states and A is the set of all actions as above. In addition, γ (γ) is a state transition function $\gamma: S \times A \rightarrow S$ that returns the resulting state of an executable action given a current state. Thus given a state and action, one infers the subsequent state $\gamma(s, \alpha) \rightarrow s'$ that follows after the action is executed.

A *classical planning problem* is a triple $P = (\Sigma, s_0, g)$, where Σ is a state transition system, s_0 is the initial state, and $g \in G$ (the *goal formula*) is a conjunction of first-order literals.¹ A goal state s_g satisfies a goal if $s_g \models g$. A *plan* (π) represents a (possibly empty) sequence of plan steps (i.e.,

¹ A goal is often represented as a ground literal or conjunction of ground literals (Ghallab, Nau & Traverso, 2004), and the literal is an arbitrary state predicate. For the purposes of this paper we will not assume the goal to be grounded (i.e., it can be existentially quantified and thus include non-ground arguments). Further the goal is a member of the set of possible goals, not the wider set of all possible states. More generally it could be a complex expression (see related research). We otherwise limit goals in this paper to conjuncts.

actions $\langle \alpha_1 \alpha_2 \dots \alpha_n \rangle$) that incrementally changes the state of the world. Here we will use a notation that enables indexing of the individual steps or sub-sequences within the plan. In equation (1), we use the subscript g to indicate a plan that achieves a specific goal. A plan is composed of the first action α_1 followed by the rest of the plan.

$$\pi_g[1..n] = \alpha_1 \mid \pi_g[2..n] = \langle \alpha_1 \alpha_2 \dots \alpha_n \rangle \quad (1)$$

Now we recursively redefine γ as mapping either single actions or plans to states. Hence π_g is a solution for P if it is executable in s_0 and $\gamma(s_0, \pi_g) \models g$. Recursively from the initial state, execution of the plan results in the goal state (see equation 2).

$$\gamma(s_0, \pi_g) = \gamma(\gamma(s_0, \alpha_1), \pi_g[2..n]) \rightarrow s_g \quad (2)$$

2.2 Planning with Nondeterminism

In the research literature on planning, uncertainty about the possible outcomes of actions has been dealt with mainly in two different ways: using Markov Decision Processes (MDPs) (Kaelbling, Littman, & Cassandra, 1995; Dean, Kaelbling, Kirman, & Nicholson, 1995; Boutilier, Dean, & Hanks, 1999) and using model-checking techniques (Cimatti, Roveri, & Traverso, 1998; Aiello *et al.*, 2001; Bertoli, Cimatti, and Traverso, 2006) like those used for program verification. Which approach is best depends on the situation: MDPs are useful when the transition probabilities are important (e.g., when it is desired to maximize expected utility), and model-checking is useful when the transition probabilities are unknown or unimportant (e.g., if it is necessary to achieve a goal regardless of which nondeterministic outcome occurs). In both approaches, an action with multiple possible outcomes is often (though not always) represented as a *nondeterministic classical operator* $o = (\text{head}(o), \text{pre}(o), \text{eff}_1(o), \text{eff}_2(o), \dots, \text{eff}_k(o))$ that has multiple possible sets of effects. When this representation is used with MDPs, each set of effects $\text{eff}_i(o)$ will have a probability $p_i(o)$, where $p_1 + p_2 + \dots + p_k = 1$. Instead of a sequential plan π that achieves a goal, MDPs learn a policy π (i.e., a mapping from states to actions) that maximizes expected utility. In this paper we will use the classical definition for clear illustration of the extension to interpretation and goal reasoning.²

2.3 Interpretation and Goal Reasoning

Broadly construed, the topic of *goal reasoning* concerns cognitive systems that can self-manage their goals (Vattam, *et al.*, 2013). Goal reasoning has recently extended the classical formulation by relaxing the assumption that an initial goal originates independently of the agent. But even if the planning process starts with an exogenous goal, a dynamic environment may present unexpected events with which the system must contend. In response, a goal reasoner must be able to generate new goals or change old ones at execution time as situations warrant.

Formally, the function *beta* (β ; see expression 3a) returns a (possibly) new goal formula g' given some state s and a current goal g . As such, β is a state interpretation process that perceives the world with respect to its goals. It is central to goal formulation and goal management.

$$\beta(s, g) \rightarrow g' \quad (3a)$$

² We ignore for now the set of exogenous events E that are outside the control (and possibly the observation) of the reasoning system. Note also that the use of the word “interpretation” in this paper does not correspond to the notion of interpretation that defines the logical semantics of formal representations. Rather it denotes a high-level perceptual-like process.

We posit a simple model of goal change $\Delta = \{\delta \mid \delta: G \rightarrow G\}$ that represents the set of potential operations or transformations on goals an agent may select. Then given s and g , the agent makes a decision $\delta_1, \delta_2, \dots, \delta_n \in \Delta$ so that the result of β becomes $\delta_n(\dots \delta_2(\delta_1(g))) = g'$. Here we do not commit to how the operations might be selected. Cost and expected benefit information might be used of course, but these details are at the discretion of the application (see section 3 for a possible approach to a complex goal-reasoning operation). Theoretically β represents the choice of an agent given a new state and current goal expression to determine anew what it wants (i.e., its goal) in the future relative to its observations.³

2.3.1 Goal Transformation

Unlike classical planning models that assume goals to be static and given externally, a goal reasoning perspective views goals as malleable and subject to change. Among the elements of Δ is the *identity transformation* $\delta^i(g_i) = g_i$ for all $g_i \in G$. The decision not to change a goal given some state s is thus an agent's selection of the identity transformation in β . But the decision to change an agent's goals in the face of an uncooperative environment represents a powerful adaptive alternative to plan modification or re-planning. Goals can undergo sequences of transformations including priority shifts and, in extreme cases, goal abandonment (Cox & Veloso, 1998).

As an example, a chess player may start out with the goal to achieve checkmate (g_{ch}). But given a series of unsuccessful opening moves (i.e., $\pi_{g_{ch}}[1..k]$ where $k < n$), the player may change the goal to draw (g_{dr}). We represent this application of the goal transformation $\delta(g_{ch}) = g_{dr}$ with a specialization of (3a) as shown in (3b).

$$\beta(\gamma(s_0, \pi_{g_{ch}}[1..k]), g_{ch}) \rightarrow g_{dr} \quad (3b)$$

Cox and Dannenhauer (in press) enumerate a specific set of transformations Δ that represent many of the potential operations on goals that exist. Furthermore over time, goals follow arcs or trajectories through a space of goals (Bengfort and Cox, 2015). These trajectories result from repeated goal changes through different instances of β interpretation. But most importantly for goal reasoning, goals can be created or formulated given a particular problem state.

2.3.2 Goal Formulation

From some initial state s_0 and an empty goal state, an agent formulates a new goal as shown in expression (3c). The operation is termed goal a *goal insertion transformation* $\delta^*(\emptyset) = g$.

$$\beta(s_0, \emptyset) \rightarrow g \quad (3c)$$

In one sense, this can still involve user-provided goals. If the input state is one resulting from a speech act whereby a human requests a goal to be achieved, the function of β is to interpret the intention of the human and to infer the goal from the utterance. In another sense, however, this significantly differs from the classical formulation of a problem. For goal reasoning in its simplest form, a planning problem can be cast as the tuple $P = (\Sigma, s_0)$. Given a state transition system and an initial state, the goal-reasoning task is to formulate a goal (if a problem indeed exists in the initial

³ The language we use here implicitly equates the current state with an observation and thus assumes full observability. This simplification is further reinforced by the use of the terms interpretation and perception within the description of β . But we could (and if space allowed should) distinguish the perceptual observation of partial states from the conceptual interpretation of the resulting percepts. We leave this exercise for future work.

state) and create (and execute) a plan to achieve it. Under classical planning (indeed under most planning schemes), the system halts when the goal state is achieved (or even when a plan is simply found). In goal reasoning, an agent can search for new problems once all goals are achieved by interpreting the final goal state s_g . In this case, (3c) becomes the expression in (3d).

$$\beta(s_g, \emptyset) \rightarrow g' \quad (3d)$$

Of course, and as we will see, goals can potentially be formulated from any current state.

2.4 A Model of Planning, Action and Interpretation

A plan to achieve a goal $\beta(s, \emptyset)$ can now be written as $\pi_{\beta(s, \emptyset)}$. Using this notation, we combine planning, action (plan execution), and interpretation in equation (4).

$$\gamma(s_0, \pi_{\beta(s_0, \emptyset)}) = \gamma(\gamma(s_0, \alpha_1), \pi_{\beta(\gamma(s_0, \alpha_1), \beta(s_0, \emptyset))}[2..n]) \quad (4)$$

Now when goals change (or new ones are added) through β , plans may need to change as well. So the problem with the above formalization is that, in the recursive right-hand side of (4), the plan is not static as defined in (1). That is, it is not necessarily of size $n - 1$. Instead, because the goal may change, the goal reasoner may need to re-plan and alter the length and composition of the remainder of the plan. To cover this contingency, we define a (re)planning function *phi* (φ ; see expression 5) that takes as input a state, goal, and current plan.

$$\varphi(s, g', \pi_g[1..n]) \rightarrow \pi_{g'}[1..m] \quad (5)$$

Note that in the general case g' may or may not be equal to g . Inserting the re-planning function into (4), we obtain equation (6) below that resolves the anomaly indicated above.

$$\gamma(s_0, \pi_{\beta(s_0, \emptyset)}) = \gamma(\gamma(s_0, \alpha_1), \varphi(\gamma(s_0, \alpha_1), \beta(\gamma(s_0, \alpha_1), \beta(s_0, \emptyset)), \pi_{\beta(s_0, \emptyset)}[2..n])) \quad (6)$$

Given φ , the formalism is general across different variations of goal reasoning and (re)planning. More fully we have equation (7) showing the initial planning on the left-hand side.

$$\begin{array}{c} \pi_{g_1} \qquad \qquad \qquad \pi_{g_2} \qquad \qquad \qquad (7) \\ \underbrace{\hspace{10em}} \\ \gamma(s_0, \varphi(s_0, \beta(s_0, \emptyset), \emptyset)) = \gamma(\underbrace{\gamma(s_0, \alpha_1)}_{s_1}, \underbrace{\varphi(\underbrace{\gamma(s_0, \alpha_1)}_{s_1}, \underbrace{\beta(\underbrace{\gamma(s_0, \alpha_1)}_{s_1}, \underbrace{\beta(s_0, \emptyset)}_{g_1})}_{s_1}, \underbrace{\varphi(s_0, \beta(s_0, \emptyset), \emptyset)}_{\pi_{g_1}}[2..n])}_{g_2}) \end{array}$$

When β generates an exogenous initial goal $g_1 = g_{user}$ from the initial state s_0 and simply returns the input goal from all other states (i.e., $g' = g$ in 3a using δ^j), the formalization reduces to classical planning with a user-given goal. That is, equation (4) is equivalent to (2) because (3a) represents a trivial boundary case. Equation (7) can also be reduced to classical planning (2) by assuming a special case of a goal insertion transformation $\delta^*(\cdot) = g_{user}$ and assuming $\Delta = \{\delta^*, \delta^j\}$. Functions β and φ are as follows.

$$\beta(s, g) = \begin{cases} \delta^*(\cdot), & (s = s_0) \wedge (g = \emptyset) \\ \delta^j(g), & otherwise \end{cases} \quad \varphi(s, g, \pi) = \begin{cases} \langle \alpha_1 \alpha_2 \dots \alpha_n \rangle, & (s = s_0) \wedge (\pi = \emptyset) \\ \pi, & otherwise \end{cases}$$

2.5 A Blocks World Example

Many details are buried within the β interpretation function. Just consider a situation where goals are given to the agent by a human user in natural language. In this case, β does not just input a goal in first-order predicate form. Instead it must translate an utterance or the results of a speech act into a predicate representation. That is, it must infer the intent of the user from the current state of the world and what was said. Indeed when a user both states imperatives and asks questions, the illocutionary act intended is often a request for the agent to assume a goal.⁴ Consider the utterance “Put block-B on block-C.” Although spoken as an action, the intent is actually for the agent to assume the goal on the behalf of the speaker to achieve the state of the block B on top of C. Although we will not discuss details here, Figure 2 shows how the formalism organizes the context of a dialog between human and agent so that intent of the last utterance “Add one more” is possible.

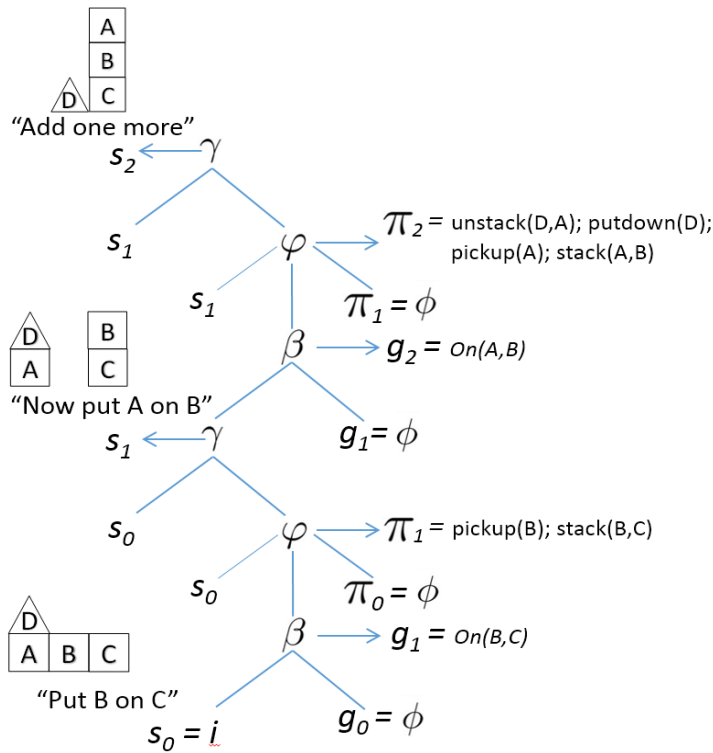


Figure 2. Partial context for blocks example (read from the bottom up)

Figure 2 hides details as well. π_1 has two actions that execute sequentially, not concurrently in one invocation of γ as the figure suggests. A more realistic depiction of the plan execution is shown in Figure 3. Here we see that the execution of the first action results in an intermediate state. From that state no new goal is formulated, and subsequently no new re-planning is performed. The goal is achieved in the next state s_2 , and β evaluates the goal achievement and returns the null goal.

⁴ For questions, the intended goal is a *knowledge goal* (Ram, 1990) as opposed to an attainment goal.

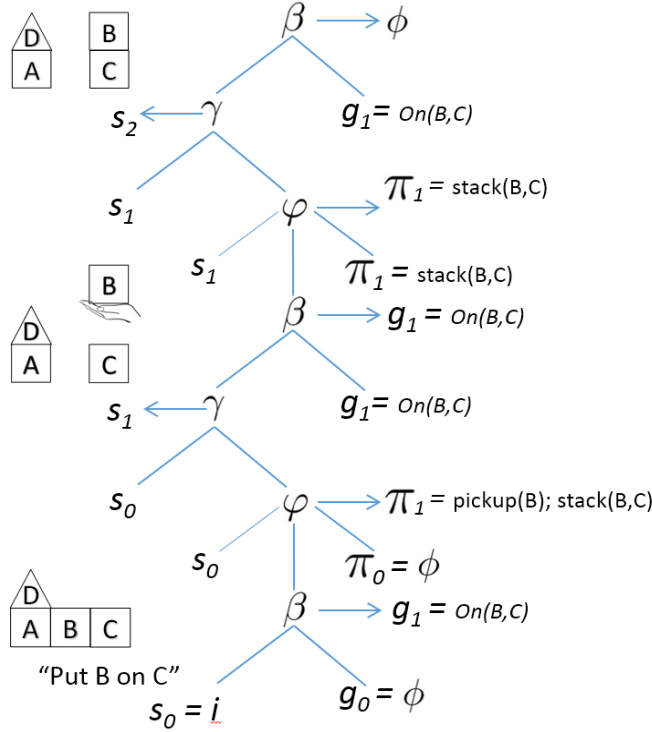


Figure 3. The execution of π_1 . Here we drop the user imperative in Figure 2 to put A on B.

3. Goal-Driven Autonomy (GDA)

Goal-driven autonomy (GDA) (Aha, et al., 2010; Cox, 2007; Klenk, Molineaux, Aha, 2013; Munoz-Avila, et al., 2010) is a kind of goal reasoning that focuses much of the research on goal formulation and causal explanation. In the GDA framework, goal management is significant and represents a key computational feature. A common framework partitions goal reasoning into a pipeline of (1) discrepancy detection; (2) explanation generation; (3) goal formulation; and finally (4) goal management (Klenk, Molineaux, & Aha, 2013). In alternate frameworks, goal formulation and management are seen as part of the same computational process, i.e., goal management can be considered various combinations of goal transformations with formulation simply an instantiation of the goal insertion transformation (Cox, 2013). Regardless, the common assumption is that formulation of goals need not depend directly upon an outside human user and represents a central feature of autonomy for intelligent agents. Here we examine GDA to illustrate concrete aspects of the theory put forth in this paper.

The current goal g_c of an agent is a distinguished member of the set of all goals the agent currently desires (\hat{G} ; see equation 8), but it differs in that the agent has selected it for planning.

$$\hat{G} = \{g_1, g_2, \dots, g_c, \dots, g_n\} \subseteq G \quad (8)$$

Also in the GDA framework, the goal reasoner produces not only a plan, π , but also a set of *expectations*, $EX = \{x_1, \dots, x_k\}$ that represent sets of anticipated states associated with each of the

$i=1..k$ actions α_i . When any currently observed state, s_c , diverges from the system's expected state, $s_e = x_c \in EX$, a *discrepancy*, $d \in D$, is said to occur (Dannenhauer, Munoz-Avila, & Cox, in press).

Given a discrepancy, a new top-level goal may be warranted, so the system will apply an *abductive explanation* χ of d to determine a cause within the context of s_c . The explanation is performed to resolve the discrepancy. Finally given d , χ , and s_c , goal generation seeks to determine a new top-level goal (or subgoal), g_c , that seeks to remove the discrepancy d by either changing the state s_c to s_e or by learning a new expectation that justifies s_c . As such we now define a GDA planning problem as a 6-tuple in equation (9).

$$P_{gda} = (\Sigma, s_c, g_c, s_e, \hat{G}, \Delta) \quad (9)$$

In the case where all plan execution outcomes equal expectations (i.e., $s_c = s_e$) and thus no explanation is necessary (i.e., $\chi = \phi$) and where the current goal is given (i.e., $g_c = g$) and the initial current state is s_0 , P_{gda} devolves into a classical planning problem $P = (\Sigma, s_0, g)$.

The explanation⁵ $\chi: \omega \rightarrow \dots \rightarrow s_c$ contains a *salient antecedent* ω that represents the root cause of the problem signaled by the discrepancy (Cox, 2007). The goal then is to remove the cause of the problem, hence $g_c = \delta() = \neg\omega$. Goal insertion is thus a search for an explanation that operationalizes the problem in terms of root causes. See Table 1 for a procedural algorithm and Cox (2013; Cox et al., 2016) for a cognitive architecture that instantiates the algorithm to implement the abstract representation of a goal insertion transformation. See Dannenhauer, Munoz-Avila & Cox (in press) for implementation details with respect to both *the expectation generator* $X(s_c, \alpha_i)$ and the *check* function that detects the discrepancy.

Table 1: GDA algorithm for adding new goal to current goal set \hat{G}

```
function GDA-goal-insertion ( $s_c : S$ ;  $\alpha_i : O$ ;  $\hat{G} : S$ ;  $X: S \times \Pi \rightarrow S$ ):  $S$ 
1.  $s_e \leftarrow X(s_c, \alpha_i)$  ;; determine expectations
2.  $d \leftarrow \text{check}(s_e, s_c, \alpha_i)$  ;;  $d: s_c \neq s_e$  see Dannenhauer et al. (in press)
3. if  $\nexists d$ , then return( $\hat{G}$ )
4. Find an explanation,  $\chi: \omega \rightarrow \dots \rightarrow s'$ ,
   where covers( $s', s_c$ )
5.  $\sigma \leftarrow \text{unify}(s', s_c)$  ;; unification returns a substitution set
6. for  $y \in \text{precond}(\chi)$  do subst( $\sigma, y$ )
7. if  $\forall y | y \in \text{precond}(\chi) \wedge \text{satisfied}(y)$ 
   then  $\hat{G} \leftarrow \hat{G} \cup \neg \text{subst}(\sigma, \omega)$  ;; insertion transformation, i.e.,  $\delta()$  inserted into  $\hat{G}$ 
return( $\hat{G}$ )
```

So consider again the blocks world example from Figures 2 and 3. When told to put A on B, the system easily creates a plan to clear A, pick it up, and place it on B. During the planning it had to solve the subgoal of having A clear before picking up the block. But B is already clear and it expects it to stay that way. Now suddenly assume that block B catches on fire and is no longer clear (see Figure 4 and Paisner, Cox, Maynord & Perlis, 2014 for an extended domain description).

Figure 4 shows state s_2 after the agent unstacked pyramid D from block A during the execution of plan π_2 . The agent is still holding D, and B is on fire and hence no longer clear. Here the expectation s_e is clear(B), whereas the current observed state s_c is $\neg\text{clear}(B)$. A very simple explanation is that a block being on fire causes it not to be clear (see expression 10).

⁵ The explanation is actually a graph $\chi = (V, E)$ with $\omega \in V$ an element of the source nodes and $s_c \in V$ a distinguished element of the sink nodes. See also Cox (2011).

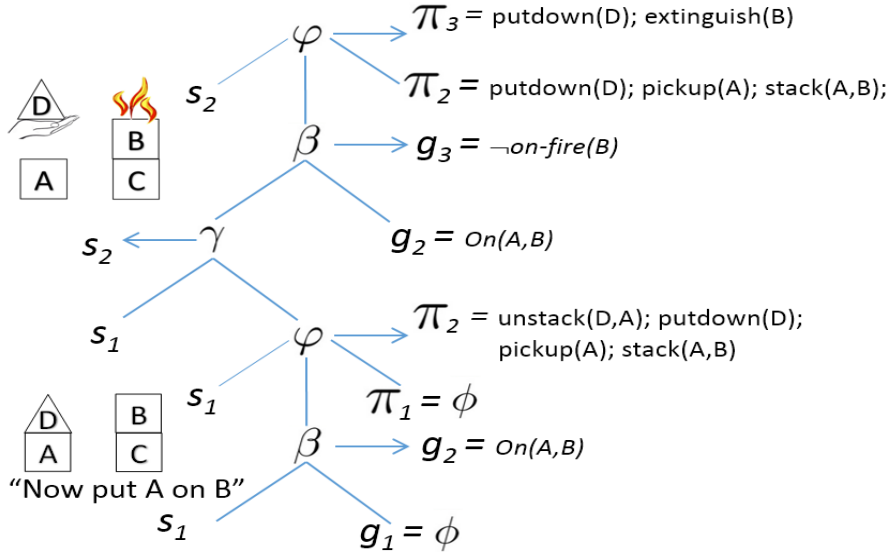


Figure 4. Explanation-based goal formulation

$$\chi: on_fire(y) \rightarrow \neg clear(y) \quad (10)$$

Given the *substitution set* $\sigma = \{y \mapsto B\}$, the goal $\neg subst(\sigma, \omega)$ becomes the achievement of $\neg on_fire(B)$ and is added to the goal set $\hat{G} = \{g_2\}$. Planning then can subsequently create a sequence of actions to put out the fire before continuing. The planning function represented by φ needs to be intelligent enough to infer that the putdown action is duplicated in both plans π_2 and π_3 . However for β to be equally intelligent, it should have been informed of the current plan. Therefore, we redefine (3a) as follows.

$$\beta(s, g, \pi, \Delta) \rightarrow g' \quad (3a')$$

In (3a'), β can use the current plan for expectations. Indeed, we stated in the example above that s_e , the expectation that B will remain clear, occurred during planning for π_2 . With such input, β can also generate a goal to retry a failed plan step during execution (e.g., vacuum cleaner did not remove all the dirt during a particular sweep). Furthermore, goal reasoning can be clear when trying to determine under unexpected and dynamic contexts (e.g., sudden resource reductions) whether plan adaptation or goal adaptation (using Δ) is the most rational choice. In this way too, planning is not only informed by goals from interpretation, interpretation is informed by planning (see also Alavi & Cox, in press, for further influence planning has upon interpretation).

3.1 GDA and the Learning of Goals

The function β is less about planning per se than about understanding the larger context within which a planning agent finds itself. If planning and interpretation are to be successful, learning must also be taken into consideration. Technically learning in this context is about inducing a set of goal states and applicability conditions from observed behavior of an execution system; and the main role for planning is to enable a GDA agent to infer which of the possible goal states are achievable within the expected costs and rewards of achieving those goals (see also Pozanco, Fernandez, & Borrajo, in press).

In the simplest sense, goal formulation can be cast as a novel classification task. An agent is given a state transition system, Σ , and a current state, s_c , and it must infer the current goal, g_c . Using a naïve supervised learning approach, one can present the system with many positive and negative examples of the tuple (s, g) . From these examples, a learning algorithm can then learn a mapping in the form of a decision tree with goals at the leaves (see Maynard, Cox, Paisner, & Perlis, 2013 for an implementation of this approach). However the time it takes to learn such relations and the effort it takes to prepare suitable examples may not be available. Instead another approach is to exploit failure in performance by explaining why expectations do not hold in new situations and then using these explanations to learn appropriate goal orientations (i.e., to learn $G_{\subseteq S}$, the set of useful goal states).

But we do not limit ourselves to goal generation alone, because if so the human burden of providing an external goal is shifted to the human specification of the complete range of goal representations. Instead if it is to be effective, the agent should also learn the classes of states in the world that constitute useful goals. This general approach represents a new conception of autonomous behavior.

Reconsider the planning domain from the introduction section and Figure 1. If a GDA-based agent interacts with other agents performing actions, it may observe a sugar distributor make a delivery to an Atlanta cola plant and then watch the manufacturer produce Coca-Cola at that location. The Coca-Cola is then delivered to Greenville where an advertising campaign has concluded. Sales then are observed to produce profits. Given these exogenous-event observations from the execution of π_{Coke} and π_{SDistr} , the GDA-agent can understand the causal relations by performing a goal regression (Veloso, 1994; Waldinger, 1981) upon the goal-subgoal graph entailed by Σ and the observations (see Figure 5). Each node in the graph is an instantiated action-operator whose preconditions must be satisfied before the action can be executed. A tree of conjuncts can then be extracted to give the causal explanation structure shown in Figure 6. For example the state $at(Cola-5, Coke)$ exists because the conjunct $has(Coke, \$) \wedge at(Bottles-2, Coke) \wedge at(Sugar-1, Coke)$ has been satisfied. If any term is unsatisfied, then the explanation structure collapses.

Now a subsequent series of observations may violate the expectations present in this causal structure. For example the system may expect Coke to have further profits in Greenville, but Coke might lack the sales. Sugar may be delivered and promotions may occur, but without bottles, no cola is produced and hence no deliveries arrive for sale. In this case, the goal insertion algorithm (Table 1) would detect the discrepancy between the expectation $has(Coke, sales)$ and the observation $\neg has(Coke, sales)$. An explanation process would then use the prior explanation (from Figure 6) to derive a relational sequence responsible for the impasse. The resulting causal chain is shown in Figure 7.

The goal insertion process would use the extracted failure chain to generate a goal and to add it to the agent’s current set of desired goals \hat{G} . In Figure 7 the salient causal antecedent ω of the explanation is $\neg at(Bottles-2, Coke)$, that is, the bottles are not at the plant (and that is ultimately why Coke lost profits). The new goal, g_c , is therefore to obtain the state of bottles being at this location (equation 11).

$$g_c = \neg (\neg at(Bottles-2, Coke)) = at(Bottles-2, Coke) \tag{11}$$

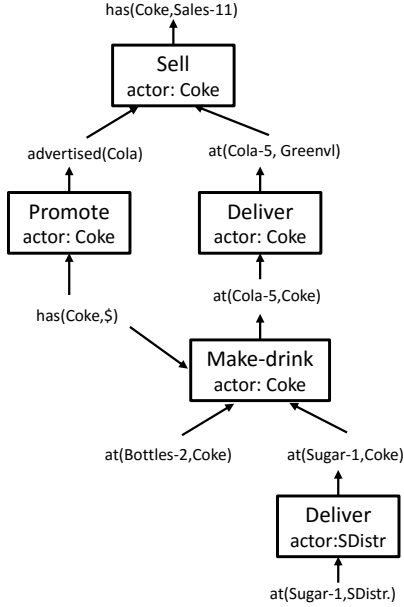


Figure 5. Goal-subgoal graph

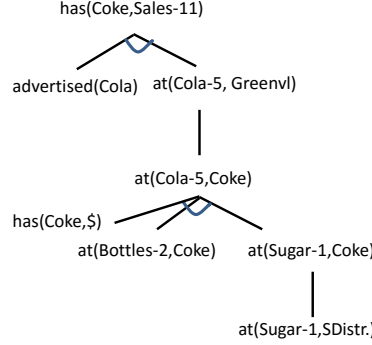


Figure 6. Causal explanation structure

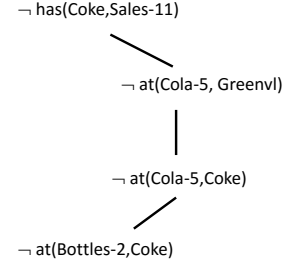


Figure 7. Failure causal chain structure

Given a cooperative agent domain where the GDA agent represents a bottle distributor (i.e., BDistr) that coordinates with the sugar distributor and the cola manufacturers, the GDA agent can solve the new goal by generating a plan to perform the delivery method. That is, it will load the bottles at the McCaysville location, drive the load to Atlanta, go to Coke, and unload the shipment there. However a number of technical problems remain. Two in particular that we will discuss now concern issues of goal generalization and of goal anticipation.

3.1.1 Goal Generalization

First the goal above must be *generalized*. It is not useful to learn an overly specific goal of getting a particular shipment to a destination. Rather the agent must generalize the shipment to any available instance of type **bottles**. Furthermore, and this is a bit more subtle, the agent must eventually generalize the destination to any cola manufacturer in the domain (and not to just any location). That is Pepsi as well as Coke require bottle stock. By doing so, this simple explanation-based generalization rules out many irrelevant potential goal states such as $at(bottles,bank)$ or $at(\$,Dalton)$. Here we intend to induce the goal expression g_c in equation (12) and to add $g_c \in G$.

$$g_c = at(b,m) \wedge bottles(b) \wedge cola-manufacturer(m) \quad (12)$$

3.1.2 Goal Anticipation and Deciding Which Goals Are Worth Generating

Second the agent needs to *anticipate* the conditions under which the causal chain in Figure 7 might re-occur and plan to prevent it in the future. It should not wait until Coke sales fail each time. Thus in this example the GDA agent learns to generate the goal, g_c , when the bottles are low in inventory, not waiting until bottles run out and sales plummet. The agent learns a characterization of the state,

s_c , in tuples (s_c, g_c) , that is a rule $g_c \leftarrow s'$. This is the problem of learning goal-selection criteria (Powell, Molineaux, & Aha, 2011).

Unlike Powell and colleagues who use human guidance to acquire goal-selection knowledge, we suggest that an agent can learn this through unsupervised means. For example the state of inventory at manufacturers may include the predicates *surplus*, *low*, *very-low*, and *out-of*. The GDA agent should learn that bottles should be delivered when the supplies are low or very low, rather than waiting until they are out or trying to deliver when a surplus exists. In another instance of failure driven learning, the GDA agent would learn to rule out the expression $\text{surplus}(m,b) \wedge \text{bottles}(b) \wedge \text{cola-manufacturer}(m)$ after it attempts to deliver bottles to a manufacturer having a bottle surplus. In this case we assume that the *unload* operation would not be given permission at the destination. Essentially the learned rule asserts the goal g_c when bottles at the manufacturer are low. This is what we meant in the introduction of this paper when we stated that the agent should be able to explain that bottles need to be at the bottling plant *because the inventory is low*.

Just because a goal state s_g would be a desirable state to be in, there are several situations in which g may not be a useful goal for a goal reasoning agent to formulate. For example, s_g may be impossible to achieve; or in a nondeterministic domain, it may be impossible to guarantee that s_g will always (or usually) be achieved. Even if s_g is achievable, there may be another goal state $s_{g'}$ that is nearly as desirable as s_g and can be achieved at much lower cost than s_g . In such a situation it may be better for GDA to insert g' into its agenda rather than g . In general it is an open research question as to the best way to learn which goals are worth formulating and which to commit to first.

4. Related Research

An alternative formal model (Johnson, Roberts, Apker, & Aha, this volume; Roberts et al., 2015; 2014) treats goal reasoning as *goal refinement*. Using an extension of the plan-refinement model of planning, Roberts and colleagues model goal reasoning as refinement search over a *goal memory* M , a set of *goal transition operators* R , and a transition function *delta* that restricts the applicable operators from R to those provided by a fundamental *goal lifecycle*. Unlike the formalism here that represents much of the goal reasoning process with the single function β , Roberts proposes a detailed lifecycle consisting of goal formulation, goal selection, goal expansion, goal commitment, goal dispatching, goal monitoring, goal evaluation, goal repair, and goal deferment. Thus many of the differential functionalities in β are distinct and explicit in the goal reasoning cycle. However the model here tries to distinguish between the planning and action side of reasoning (i.e., φ and γ) and the interpretation and evaluation components inherent in goal reasoning (i.e., β). We can roughly note, however, that φ relates to Robert's goal expansion and repair, whereas β relates to goal formulation, monitoring, and deferment. Missing in my current analysis is the processes of goal dispatching, commitment, and evaluation.

Additionally Roberts proposes a more complex goal structure. A *goal node* includes not only the desired state but also super-ordinate and subordinate goal linkages, goal constraints, quality metrics, and pointers to the current plan associated with the node. We have been more circumspect regarding the syntactic structure of a goal. In section 2.1 on classical planning, we defined the goal formula to be a conjunction of first-order literals. More generally we would also allow universally quantified predicate conjuncts as in Cox and Veloso (1998; Veloso, Pollack & Cox, 1998). Thus in a package delivery domain we allow $\forall p | (\text{package } p) \wedge \exists d | (\text{destination } p d) \wedge (\text{location } p d)$ as a goal to deliver all packages to their destination. Under the scope of this goal expression, new goals $(\text{location } p d)$ may arise during planning or plan execution time as

additional packages arrive at the warehouse for delivery. See also Talamadupula, Benton, Kambhampati, Schermerhorn, & Scheutz (2010) for their concept of *open world quantified goals*. But overall, much exists in common with the representations here and with Roberts including an association with learning (Roberts & Aha, 2015).

The two problems of goal generalization and goal anticipation make the learning task as we discuss it here considerably different from related research. Planning research is concerned with generating an action sequence to carry out a goal that is given to a system by an external user (see survey in Ghallab, Nau, & Traverso, 2004). In the larger scope of integrated cognitive systems, we intend to better understand autonomy as the capacities (1) to recognize novel problems and (2) to independently form the desire to remove these problems. Goal anticipation and goal generalization are instrumental processes associated respectively with these two aspects of autonomy. However unlike bottom-up, data-driven learning approaches, the top-down explanation of anomalies is a central pivot in our learning approach (Cox, 2011).

The explanation-based approach to GDA and to learning we take should not be confused with the older, mainly deductive approach used in explanation-based learning (DeJong & Mooney, 1986; Mitchell, Keller, & Kedar-Cabelli, 1986). We use explanation techniques (Cox, 2011; Cox & Ram, 1999; Roth-Berghofer, Tintarev, & Leake, 2011) that depart significantly from earlier work, being more abductive, knowledge rich, and case-based. The idea is to retrieve a graph structure called an explanation pattern given a discrepancy and a context, to instantiate it and to bind it to the discrepancy, and then to adapt it to the context. This also departs from other GDA systems (e.g., ARTUE, Molineaux, Klenk & Aha, 2010) which use abductive causal inference in the form of an assumption-based truth maintenance system (de Kleer, 1986) for explanation.

As mentioned previously, goal formulation or generation has been an instrumental task in research on goal reasoning and goal-driven autonomy. In the ARTUE system, goal formulation occurs as a response to discrepancy detection and is based upon knowledge structures called principles (Klenk, Molineaux, & Aha, 2013). Other approaches rely upon links between specific states and a priori goal candidates (Dill & Papp, 2005) or triggering schemas based on the current state (Talamadupula, Benton, Schermerhorn, Kambhampati, & Scheutz, 2009). In the EISBot interactive game-playing system, association rules called goal formulation behaviors link discrepancy-explanation types to goal types (Weber, Mateas & Jhala, 2010). Here we discussed the use of the negation of explanation root causes to generate new goals.

5. Conclusion

This research attempts to reconcile some of the existing research on goal reasoning in the cognitive systems community with theoretical frameworks in the automated planning and the intelligent agents communities. The resulting goal-reasoning formalism augments notions of planning and plan execution with formal models of re-planning and both goal formulation and goal change. In situations where goals are given at initialize time and do not change throughout the process and where plans execute as expected, the model is equivalent to a classical planning formalism. We have illustrated the model with a kind of goal reasoning called goal-driven autonomy, and we have applied the model to ideas about explanatory learning of those goals worth pursuing.

Acknowledgements

This material is based upon work supported by ONR grant N00014-15-1-2080. I also wish to acknowledge the previous input of David Aha, Michael Maynard, Dana Nau, Don Perlis and Mark Wilson on an earlier formulation of these ideas.

References

- Aha, D. W., Cox, M. T., & Munoz-Avila, H. (Eds.) (2013). *Goal Reasoning: Papers from the ACS workshop* (Tech. Rep. No. CS-TR-5029). College Park, MD: University of Maryland, Department of Computer Science.
- Aha, D. W., Klenk, M., Munoz-Avila, H., Ram, A., & Shapiro, D. (Eds.) (2010). *Goal-driven autonomy: Notes from the AAI workshop*. Menlo Park, CA: AAI Press.
- Aiello, L. C., Cesta, A., Giunchiglia, E., Pistore, M., & Traverso, P. (2001). Planning and verification techniques for the high level programming and monitoring of autonomous robotic devices. *European Space Agency Wkshp on On-Board Autonomy*, Noordwijk, Netherlands. ESA.
- Alavi, Z., & Cox, M. T. (in press). Rationale-based visual planning monitors. To appear in M. Roberts (Ed.), *Working Notes of the 4th Workshop on Goal Reasoning*. New York, IJCAI-16. <http://makro.ink/ijcai2016grw>
- Bengfort, B., & Cox, M. T. (2015). Interactive reasoning to solve knowledge goals. In D. W. Aha (Ed.), *Goal reasoning: Papers from the ACS workshop* (pp. 10-25). Tech. Rep. No. GT-IRIM-CR-2015-001. Atlanta, GA: Georgia Institute of Technology, Institute for Robotics and Intelligent Machines.
- Bertoli, P., Cimatti, A., Roveri, M., & Traverso, P. (2006). Strong planning under partial observability. *Artificial Intelligence*, 170(4): 337–384.
- Boutilier, C., Dean, T. L., & Hanks, S. (1999). Decision-theoretic planning: Structural assumptions and computational leverage. *Journal of Artificial Intelligence Research (JAIR)*, 11:1–94.
- Cimatti, A., Roveri, M., & Traverso, P. (1998). Strong planning in non-deterministic domains via model checking. In R. Simmons, M. Veloso, & S. Smith (Eds.), *Proceedings of the Fourth International Conference on Artificial Intelligence Planning Systems* (pp. 36–43). Menlo Park, CA: AAI Press.
- Cox, M. T. (2013). Goal-driven autonomy and question-based problem recognition. In *Second Annual Conference on Advances in Cognitive Systems 2013, Poster Collection* (pp. 29-45). Palo Alto, CA: Cognitive Systems Foundation.
- Cox, M. T. (2011). Metareasoning, monitoring, and self-explanation. In M. T. Cox & A. Raja (Eds.) *Metareasoning: Thinking about thinking* (pp. 131-149). Cambridge, MA: MIT Press.
- Cox, M. T. (2007). Perpetual self-aware cognitive agents. *AI Magazine* 28(1), 32-45.
- Cox, M. T., Alavi, Z., Dannenhauer, D., Eyorokon, V., Munoz-Avila, H., & Perlis, D. (2016). MIDCA: A metacognitive, integrated dual-cycle architecture for self-regulated autonomy. In *Proceedings of the 30th AAI Conference on Artificial Intelligence, Vol. 5* (pp. 3712-3718). Palo Alto, CA: AAI Press.
- Cox, M. T., & Dannenhauer, D. (in press). *Goal transformation and goal reasoning*. To appear in M. Roberts (Ed.), *Working Notes of the 4th Workshop on Goal Reasoning*. New York, IJCAI-16. <http://makro.ink/ijcai2016grw/>
- Cox, M. T., & Ram, A. (1999). Introspective multistrategy learning: On the construction of learning strategies. *Artificial Intelligence*, 112, 1-55.
- Cox, M. T., & Veloso, M. M. (1998). Goal transformations in continuous planning. In M. desJardins (Ed.), *Proceedings of the 1998 AAI Fall Symposium on Distributed Continual Planning* (pp. 23-30). Menlo Park, CA: AAI Press / The MIT Press.

- Dannenhauer, D., Munoz-Avila, H., & Cox, M. T. (in press). Informed expectations to guide GDA agents in partially observable environments. To appear in *Proceedings of the 25th International Joint Conference on Artificial Intelligence*. Menlo Park, CA: International Joint Conference on Artificial Intelligence, Inc.
- Dean, T., Kaelbling, L. P., Kirman, J., Nicholson, A. E. (1995). Planning under time constraints in stochastic domains. *Artificial Intelligence* 76(1-2): 35-74.
- DeJong, G., & Mooney, R. (1986). Explanation-based learning: An alternative view. *Machine Learning*, 1(2), 145-176.
- de Kleer, J. (1986). An assumption-based TMS. *Artificial Intelligence*, 28(2), 127-162.
- Dill, K., & Papp, D. (2005). A goal-based architecture for opposing player AI. In *Proceedings of the First Artificial Intelligence for Interactive Digital Entertainment Conference* (pp. 33-38). Menlo Park, CA: AAAI Press.
- Ghallab, M., Nau, D., & Traverso, P. (2004). *Automated planning: Theory and practice*. San Francisco: Morgan Kaufmann.
- Hawes, N. (2011). A survey of motivation frameworks for intelligent systems. *Artificial Intelligence*, 175(5-6), 1020-1036.
- Johnson, B., Roberts, M., Apker, T., & Aha, D. (this volume). Goal reasoning with information measures. In *Fourth Annual Conference on Advances in Cognitive Systems 2016*. Palo Alto, CA: Cognitive Systems Foundation.
- Kaelbling, L. P., Littman, M. L., & Cassandra, A. R. (1995). Partially observable Markov decision processes for artificial intelligence. In *Proceedings of Reasoning with Uncertainty in Robotics* (pp. 146-163).
- Klenk, M., Molineaux, M., & Aha, D. (2013). Goal-driven autonomy for responding to unexpected events in strategy simulations. *Computational Intelligence*, 29(2), 187–206.
- Maynard, M., Cox, M. T., Paisner, M., & Perlis, D. (2013). Data-driven goal generation for integrated cognitive systems. In C. Lebiere & P. S. Rosenbloom (Eds.), *Integrated Cognition: Papers from the 2013 Fall Symposium* (pp. 47-54). Technical Report FS-13-03. Menlo Park, CA: AAAI Press.
- Mitchell, T. M., Keller, R., & Kedar-Cabelli, S. (1986). Explanation-based generalization: A unifying view, *Machine Learning*, 1(1), 47-80.
- Munoz-Avila, H., Jaidee, U., Aha, D. W., Carter, E. (2010). Goal-driven autonomy with case-based reasoning. In *Case-Based Reasoning. Research and Development, 18th International Conference on Case-Based Reasoning, ICCBR 2010* (pp. 228-241). Berlin: Springer.
- Molineaux, M., Klenk, M., & Aha, D. (2010). Goal-driven autonomy in a navy training simulation. In *Proceedings of Twenty-Fourth AAAI Conference on Artificial Intelligence*. Menlo Park, CA: AAAI Press.
- Paisner, M., Cox, M. T., Maynard, M., Perlis, D. (2014). Goal-driven autonomy for cognitive systems. In *Proceedings of the 36th Annual Conference of the Cognitive Science Society* (pp. 2085–2090). Austin, TX: Cognitive Science Society.
- Powell, J., Molineaux, M., & Aha, D. W. (2011). Active and interactive discovery of goal selection knowledge. In *Proceedings of the Twenty-Fourth International Florida Artificial Intelligence Research Society Conference*.

- Pozanco, A., Fernández, S., & Borrajo, D. (in press). On learning planning goals for traffic control. To appear in M. Roberts (Ed.), *Working Notes of the 4th Workshop on Goal Reasoning*. New York, IJCAI-16. <http://makro.ink/ijcai2016grw>
- Ram, A. (1990). Knowledge goals: A theory of interestingness. *Proceedings of the 12th Annual Conference of the Cognitive Science Society* (pp. 206–214).
- Roberts, M., & Aha, D. W. (2015). Cleaning efficiently: A case study in modeling goal reasoning and learning. To appear in *Proceedings of the 2015 Annual Conference on Advances in Cognitive Systems: Workshop on Goal Reasoning* (IRIM Tech Report). Atlanta: Georgia Institute of Technology.
- Roberts, M., Vattam, S., Alford, R., Auslander, B., Apker, T., Johnson, B., & Aha, D. W. (2015). Goal reasoning to coordinate robotic teams for disaster relief. In A. Finzi, F. Ingrand, & Andrea Orlandini (Eds.) *Planning and Robotics: Papers from the ICAPS Workshop*. Jerusalem, Israel: AAAI Press.
- Roberts, M., Vattam, S., Alford, R., Auslander, Apker, T., Johnson, B., & Aha, D.W. (2015). Goal reasoning to coordinate robotic teams for disaster relief.
- Roberts, M., Vattam, S., Alford, R., Auslander, B., Karneeb, J., Molineaux, M., Apker, T., Wilson, M., McMahon, J., & Aha, D. W. (2014). Iterative goal refinement for robotics. In *Proceedings of ICAPS 2014*.
- Roth-Berghofer, T., Tintarev, N., & Leake, D. B. (Eds.) (2011). *Proceedings of the IJCAI-11 Workshop on Explanation-aware Computing ExaCt 2011*, July 2011.
- Talamadupula, K., Benton, J., Kambhampati, S., Schermerhorn, P., & Scheutz, M. (2010). Planning for human-robot teaming in open worlds. *ACM Transactions on Intelligent Systems and Technology*, 1(2).
- Talamadupula, K., Benton, J., Schermerhorn, P., Kambhampati, S., & Scheutz, M. (2009). Integrating a closed world planner with an open world robot: A case study. In M. Likhachev, B. Marthi, C McGann, & D. E. Smith (Eds.) *Bridging the Gap between Action and Motion Planning: Papers from the ICAPS Workshop*. Thessaloniki, Greece.
- Vattam, S., Klenk, M., Molineaux, M., & Aha, D. (2013). Breadth of approaches to goal reasoning: A research survey. In D. W. Aha, M. T. Cox, & H. Munoz-Avila (Eds.), *Goal Reasoning: Papers from the ACS workshop* (pp. 111-126). Tech. Rep. No. CS-TR-5029. College Park, MD: University of Maryland, Department of Computer Science.
- Veloso, M. (1994). *Planning and learning by analogical reasoning*. Berlin: Springer.
- Veloso, M. M., Pollack, M. E., & Cox, M. T. (1998). Rationale-based monitoring for continuous planning in dynamic environments. In R. Simmons, M. Veloso, & S. Smith (Eds.), *Proceedings of the Fourth International Conference on Artificial Intelligence Planning Systems* (pp. 171-179). Menlo Park, CA: AAAI Press.
- Waldinger, R., (1981). Achieving several goals simultaneously. In N. J. Nilsson & B. Webber (Eds.), *Readings in Artificial Intelligence* (pp. 250-271). Palo Alto, CA: Tioga.
- Weber, B. G., Mateas, M., & Jhala, A. (2010). Applying goal-driven autonomy to StarCraft. In *Proceedings of the Sixth AAAI Conference on Artificial Intelligence and Interactive Digital Entertainment*. Menlo Park, CA: AAAI Press.