

# Problem Recognition, Explanation and Goal Formulation

Sravya Kondrakunta, Venkatsampath Raja Gogineni, Matthew Molineaux, Michael T. Cox

Wright State University, Dayton OH 45435, USA  
{kondrakunta.2, gogineni.14, matthew.molineaux,  
michael.cox}@wright.edu

**Abstract.** Goal reasoning agents can solve novel problems by detecting an anomaly between expectations and observations, generating explanations about plausible causes for the anomaly, and formulating goals to remove the cause. Yet, not all anomalies represent problems. This paper addresses discerning the difference between benign anomalies and those that represent an actual problem for an agent. Furthermore, we present a new definition of the term “problem” in a goal reasoning context. This paper discusses the role of explanations and goal formulation in response to the developing problems and implements it; the paper also illustrates the above in a mine clearance domain and a labor relations domain. We also show the empirical difference between a standard planning agent, an agent that detects anomalies, and an agent that recognizes problems.

**Keywords:** Goal reasoning, cognitive architecture, case-based reasoning, explanation, goal formulation.

## 1 Introduction

An intelligent, autonomous agent in a partially observable world should formulate its own goals, make plans to achieve those goals and successfully execute those plans. An agent can formulate its own goals based on an anomaly, i.e., the difference between an expected state and an observed state, by generating a hypothesis that explains this anomaly and generating new goals that respond to the hypothesis. However, many anomalies that arise in the real world do not represent a problem to an agent’s mission or goal. For example, the playing of unexpected loud music may or may not be a problem for roommates. If a roommate is preparing for an upcoming exam, the music is a problem. If, on the other hand, she is doing her laundry, it is not a problem. Generally speaking, an agent does not need to respond to every observed anomaly; an intelligent agent should be capable of distinguishing between those that signal a problem and those that do not.

The *Goal Driven Autonomy (GDA)* approach [5, 6, 20, 22] to agency represents an appropriate response for an autonomous agent’s anomalies. However, the existing research does not formally address the issue of which anomalies should be considered problems; for example, ARTUE [20] motivations implicitly consider certain situations

to require new goals, but the system provides no formal basis for determining whether a certain situation is problematic. In this paper, we consider the task of recognizing whether an anomaly should constitute a problem for an agent. Performing this efficiently will improve both the effectiveness and robustness of the agent. We use the term *problem* in this paper to refer to anomalies that require a response in order to meet the agent's goals. Then, we present such an approach and further discuss the role of explanations and goal formulation in recognizing and responding to problems. We show empirical results illustrating the effectiveness of this approach as part of a GDA agent in two uncertain, dynamic environments.

The paper continues as follows. Section 2 defines the problem recognition task and presents a formalism to represent the problem. Section 3 describes *explanation patterns (XP)* and their role in understanding problems as well as the goals formulated when an agent detects a problem. Section 4 discusses the implementation of the problem recognition task in two domains: mine clearance domain and labor relations domain. The evaluation of GDA agents and the presentation of the results follows in Section 5. We consider related research in Section 6 and conclude with Section 7 with a discussion about some ideas for future research.

## 2 Formalities and Notation

An anomaly occurs when the expected state of the agent does not match its current observed state, but a problem only occurs when the agent needs to address the above anomaly. As such, problem recognition refers to reasoning about the anomaly and deciding whether it is something that the agent needs to handle. There might be different types of problems as well as different ways to recognize and address them. One such problem is the planning problem.

### 2.1 Classical Planning Problem Representation

A classical planning domain is defined as a finite state-transition system in which each state  $s \in S$  is represented by a finite set of ground atoms [14]. A classical planning problem is a triple  $P = (\Sigma, s_0, g)$ , where  $\Sigma$  is a planning domain,  $s_0$  is the initial state, and the goal  $g$  is a conjunction of first-order literals. A state  $s_g$  satisfies a goal if  $s_g \models g$ ; in this situation we refer to  $s_g$  as a goal state. A plan  $\pi_g \in \Pi$  represents a solution to  $P$  if it consists of a sequence of plan steps  $(a_1, a_2, \dots, a_n)$  that incrementally changes the world, starting from the initial state  $s_0$  and ending in a goal state. That is, it is a solution if  $\gamma(\pi_g, s_0) = \gamma(\dots \gamma(\gamma(s_0, a_1), a_2) \dots, a_n) \models g$ .

### 2.2 Extended Planning Problem Representation

Our extended definition of planning problems is intended for usage during execution. New problems arise during problem solving and execution, and necessitating updated solutions. These iterated online planning problems therefore incorporate an agent's prior expectations and knowledge about the execution context. The extended problem also considers whether an agent should formulate new goals in response to the changing

world. Solutions to the extended problem comprise, in addition to a plan, an explanation of an encountered anomaly and an updated goal agenda. Critically, the plan need not solve every goal in the goal agenda. Formally, we define an extended planning problem  $\mathcal{P}_x = (s_c, s_e, Bk, H_c)$  where:

- $s_c \in S$  is the current state of the environment,
- $s_e \in S$  is the state the agent previously expects to hold at this point in time,
- $Bk$  is the agent's background knowledge,
- $H_c$  is an episodic history.

A solution  $\Psi_x$  to a problem  $\mathcal{P}_x$  is of the form  $(\hat{G}_c, g_n, \chi_c, \pi_c)$  where:

- $\hat{G}_c$  is an updated goal agenda (set of pending goals for the agent to accomplish),
- $g_n$  is a goal from the agenda (i.e.,  $g_n \in \hat{G}_c$ ) chosen to address next,
- $\chi_c$  is an causal explanation that accounts for any discrepancy between  $s_c$  and  $s_e$ , and
- $\pi_c$  is a plan that will accomplish  $g_i$ , at least one goal,  $g \in \hat{G}_c$  in at least one possible world.

Note that the subscript “ $c$ ” refers to the current iteration throughout this definition and later in this section; as the extended planning problem is iterative, this explicitly links present outputs to future inputs.

The agent's background knowledge  $Bk$  is a tuple  $(\Sigma, \Delta)$  consisting of the planning domain  $\Sigma$  as defined in the classical planning problem along with a set  $\Delta$  of goal operation models  $\delta = (head(\delta), parameter(\delta), pre(\delta), res(\delta))$ . In a goal operation  $pre(\delta)$  and  $res(\delta)$  are a set of preconditions and a result. The transformation's identifier is  $head(\delta)$ , and its input goal argument is  $parameter(\delta)$ . Any goal operation model  $\delta$  with no input goal (written  $parameter(\delta) = \emptyset$ ) models a *goal formulation* operation. Otherwise,  $\delta$  represents *goal change* operations. Collectively, the agent's goal operation models  $\Delta$  define an interpretation function  $\beta: S \times G \rightarrow G$  that transforms an earlier goal into a desired goal [5, 6]. See prior work for additional detail on the function  $\beta$ .

The episodic history  $H_c$  is a tuple  $(\hat{G}_h, \pi_h, \varepsilon_h, \chi_h)$  that includes the agent's memory of past goal agendas  $\hat{G}_h = (\hat{G}_1, \hat{G}_2, \dots, \hat{G}_{c-1})$ , plans  $\pi_h = (\pi_1, \pi_2, \dots, \pi_{c-1})$ , and causal explanations  $\chi_h = (\chi_1, \chi_2, \dots, \chi_{c-1})$ , as well as the execution history containing states and actions  $\varepsilon_h = (s_0, a_1, s_1, a_2, \dots, s_{c-1}, a_c)$ .

### 2.3 Problem Recognition Subproblem

We describe three subproblems we address: problem recognition, goal formulation and change, and replanning. Problem recognition requires the agent to determine what problem, if any, occurred. Specifically, given the extended planning problem tuple  $(s_c, s_e, Bk, H_c)$ , problem recognition outputs a problem, root cause, and explanation pair  $(d, \omega_c, \chi_c)$  such that:  $\exists d, g_i: s_c \models d \wedge s_e \not\models d \wedge restricts(g_i, d, \Sigma) \wedge g_i \in \hat{G}_{c-1} \wedge \omega_c \xrightarrow{\chi_c} d \wedge \omega_c \notin \hat{G}_{c-1}$

This is read: There is some discrepancy, a literal conjunct  $d$ , that was observed (i.e., entailed by  $s_c$ ) but not expected (i.e., not entailed by  $s_e$ ). The discrepancy  $d$  restricts how the agent can address a goal  $g_i$  in its legacy goal agenda  $\hat{G}_c$ . The discrepancy  $d$

was caused by a root cause  $\omega_c$  (a literal conjunct) according to explanation  $\chi_c$ . Finally, the root cause  $\omega$  was not intended (i.e., not in the legacy agenda  $\hat{G}_c$ ). In the above, a goal  $g_i$  is *restricted by* a discrepancy  $d$  if there is no plan  $\pi$  that can accomplish the goal  $g_i$  without first eliminating discrepancy  $d$ . The explanation  $\chi_c$  is a logical derivation or proof tree. When this condition is met, the discrepancy  $d$  is considered a problem.

## 2.4 Goal Formulation and Change Subproblem

Once a problem is recognized, the agent must update its goal agenda to respond. Given a current state  $s_c$ , problem cause  $\omega_c$ , explanation  $\chi_c$ , and legacy goal agenda  $\hat{G}_{c-1}$ ; goal formulation and change must find a new goal agenda  $(\hat{G}_c, g_n)$  that responds “appropriately”. Currently, no recognized definition of what constitutes a correct or strong goal agenda outside of a particular target problem exists.

## 2.5 Replanning Subproblem

The replanning subproblem finds a revised or new plan to accomplish the newly formulated goal  $g_n$ . This problem is much the same as the classical planning problem, adding only a legacy plan. It is defined by the tuple  $(\Sigma, s_c, g_n, \pi_{c-1})$ , and the solution is the new plan  $\pi_c$ . We do not further consider replanning within the scope of this paper, choosing to reuse results of other research.

# 3 Explanation and Goal Formulation

In this work, an *explanation* is a causal structure that represents a hypothesis about the cause of an anomaly. Problem explanations hypothesize the cause of an anomaly that limits an agent’s goals. In a nutshell, explanations help an agent to decide whether an anomaly is a problem or not, while goal formulation helps an agent to resolve a problem.

To implement our ideas, we modified Meta-AQUA [7] an open-source story understanding system, to generate grounded explanations. In prior work, Meta-AQUA was supplied with a case-base of explanation patterns that explain anomalous actions performed by actors in a story. In our work, we extended Meta-AQUA to retrieve an explanation from memory and adapt it for goal formulation whenever an agent encounters a problem.

## 3.1 Explanation Pattern

In our work, we use a case-base of problem explanations engineered manually to fit the domain. Approaches to learning such explanations are outside the scope of this paper, but [25] has sketched out a detailed approach on learning explanation patterns. Each explanation in our case-base is an abstract *explanation pattern (XP)* [26]. An XP is a data structure that represents a causal relationship between multiple states and/or actions; variables adapted during or after case retrieval abstractly define each action/state.

### 3.2 Retrieving an Explanation Pattern

When an agent observes an anomalous state or action, a retrieval process outlined in [15] gets a problem explanation from the case-base. In general, the agent might retrieve zero, two, or more explanations. However, for the purpose of this paper we assume the agent retrieves exactly one problem explanation if an anomaly is a problem. However, in the future, we would like to address this assumption by select one problem by evaluating multiple explanations.

### 3.3 Goal Formulation

Goal formulation is essential for an intelligent agent to respond to unpredictable events. We perform formulation by chaining backward on each of the antecedents of a retrieved explanation until we reach all the causes to which the agent can respond [15, 16]. The causes in an explanation pattern include actions and states; therefore, when the agent wishes to prevent an undesired consequent from recurring, it considers as potential goals the elimination of the actors or objects that cause the anomalous state. The removal mapping function that takes in the agents or objects and outputs the goals that eliminate them create potential goals.

## 4 Domains

To illustrate these concepts in this paper and to assess the performance of the resulting GDA agents, it will be useful to consider them in the context of the following concrete examples from the mine clearance domain and the labor relations domain.

### 4.1 Mine Clearance Domain

To prepare a harbor for use during maritime operations, it is essential to conduct mine clearance activities to ensure that ships can operate safely as they transit between the open sea and the port in the harbor. As searching and clearing mines in the entire harbor is likely to be a time-consuming and expensive undertaking, a network of safe shipping lanes is typically established to reduce the size of the area within the harbor that needs clearing. Such a system is known as a Q-route [19].

For experimentation, we created simulated scenarios with a fixed Q-route that consists of a single shipping lane. In simulation, an Autonomous Underwater Vehicle (AUV) controlled by an agent performs both mine detection and clearance. In each scenario, the agent knows of two previously identified areas within the Q-route – green area one (GA1) and green area two (GA2) – where mines are expected. As such, any mines encountered which do not lie within GA1 or GA2 constitute anomalies. However, only anomalous mines within the Q-route are classified as problems, because mines outside the Q-route will not pose a hazard to shipping. It is the role of the agent to determine how to respond to these anomalous mines in each scenario.




Figure 1: Simulation of the mine clearance domain in Moos IvP.

Figure 1 illustrates the mine clearance domain. The cylindrical object in the top left corner represents the AUV named Remus and each triangle represents a ground truth mine position, not given to the agent a priori. The area between the two horizontal lines represents the Q-route, and the octagons on the left and right represent GA1 and GA2 respectively. In each scenario, the mines are uniformly distributed throughout the transit area (used by Remus to enter and exit the Q-route) and the Q-route.

**Example Mine Clearance Problems** In this section, we give a series of example extended planning problems in the Mine Clearance domain and their solutions. In the Mine Clearance domain, the initial goal agenda ( $\hat{G}_c$ ) includes the following goals: clear the mines in GA1, clear the mines in GA2, and head back to its initial position (home). Let us assume that the agent selects the goals in the provided order they so the current goal of the agent would be to clear the mines in GA1. Therefore, initially  $g_0 = \text{cleared-mines}(\text{remus}, \text{ga1})$  is the goal. The expectation of the agent is that the mines are present only in GA1. The initial plan ( $\pi_0$ ) is comprised of several steps to achieve the first of its initial goals ( $g_0$ ); some of which are the following:

$$\pi = \{ a_1 = \text{move\_to\_the\_location}(\text{remus}, \text{home}, \text{location-a}), a_2 = \text{move\_to\_the\_location}(\text{remus}, \text{location-a}, \text{location-b}), a_3 = \text{move\_to\_the\_location}(\text{remus}, \text{location-b}, \text{ga1}), a_4 = \text{survey}(\text{ga1}), a_5 = \text{identify-mines}(\text{ga1}), a_6 = \text{clear-mines}(\text{ga1}) \}$$

Here location-a and location-b are locations of waypoints outside the Q-route that describe a path to GA1 from “home”, the agent’s launching point. After reaching GA1, the plan directs the Remus to survey locations in GA1, locate/identify any mines and clear them. In the initial state ( $s_0$ ), GA1 is expected to (likely) contain one or more unknown mines. After  $g_0$  is satisfied, however, the agent has cleared all mines in GA1.

The agent detects a mine at location-b after it completes the action  $a_1$ . This is an anomaly because in the observed current state ( $s_c$ ), there is a mine at location-b, but not in the expected state ( $s_e$ ). This triggers problem recognition. However, no explanation pattern is found therefore the problem remains false as the detected mine is outside of the Q-route. The anomalous mine is therefore not considered a problem, and goal formulation does not occur. As such, the new solution is  $\Psi_x = (\hat{G}_0, \emptyset, \emptyset, \text{tail}(\pi_0))$ , containing the unchanged initial goal agenda, no new goal or explanation, and the remaining actions in  $\pi_0$ . The history  $H_c = (\hat{G}_h, \pi_h, \varepsilon_h, \chi_h)$  is updated with the unchanged goal agenda, updated plan, state, action, and explanation.

After the agent achieves  $g_0$ , it travels towards GA2 to achieve the next goal,  $g_1 = \text{cleared-mines}(\text{remus}, \text{ga2})$ . Its new plan visits intermediate locations location-d and location-e in the route to clearing mines in GA2. In the expected state ( $s_e$ ), the Remus is at location-e between GA1 and GA2, where there are no mines. However, in the current state ( $s_c$ ), a mine is observed. As described previously, this anomaly triggers problem recognition. This time, the anomaly is a problem, as the situation matches the problem pattern from explanation, producing an explanation  $\chi_c$  with a root cause  $\omega_c = \text{mine-at}(\text{location-e}) \wedge \text{in-qroute}(\text{location-e})$ . Goal formulation then generates a new goal  $g_n = \neg \text{mine-at}(\text{location-e})$ , and replanning produces a plan  $\pi_c$  that achieves  $g_n$ ; the updated goal agenda adds this new goal:  $\hat{G}_c = \hat{G}_{c-1} \cup \{g_n\}$ . This completes the solution  $(\hat{G}_c, g_n, \chi_c, \pi_c)$ , which is used to update the history  $H_c$  as before. The agent then switches to the new plan  $\pi_c$ .

#### 4.2 Labor Relations Domain

This domain describes a virtual institution, consisting of an institute head, employees, and customers. The institution starts with an initial reputation and a budget represented by numeric values. The head of the institute enacts policies; implementing a policy takes a known fixed amount of budget and increases the institute's reputation value by a fixed amount. Disagreements about enacted policies may occur between the head and the employees with certain intensity. Intensity is a numeric value; high intensity disagreements may lead to a strike. Intensity values vary for each disagreement and are unpredictable. Disagreements can be resolved by negotiating. Negotiating to solve a disagreement also requires a budget, which varies with respect to the intensity of the disagreement. Moreover, negotiations also decrease the reputation value by a function of intensity.

The agent acts as the head of the institute with a goal to increase the reputation of the institute, and, in order to achieve that goal, the agent must implement some policies. The expectation is that the employees agree to the policy, but, if the employees disagree then this is considered an anomaly, and only those disagreements that can lead to a strike are considered as problems. Therefore, the intelligence of the agent lies in identifying the anomalies that might lead to strikes and working on them. This domain is not related to trading and marketing agents in artificial intelligence, and is not simulated using third-party software.

### 5 Evaluation of the Implementation in the Domains

In both domains, in order to perform the evaluation we have introduced two other agents along with our GDA agent, namely an eager agent and a baseline agent. They each respond differently to anomalies. The GDA agent detects all anomalies, but only works on those perceived as problems. In contrast, the eager agent addresses all anomalies that it encounters, i.e., it tries to fix every anomaly it encounters. The baseline agent plans only for its original goals and ignores all anomalies. The performance of the three agents is assessed and the results are presented by varying the environment and averaging the results for 10 different runs for each mine scenario.

### 5.1 Empirical Results in Mine Clearance Domain

The performance of the agents is calculated based on the number of the ships that reach the other side of the harbor safely in various mine density scenarios that include a total of six ships and three mine densities: low, medium, and high. We also introduced deadlines ranging from 0 to 2 seconds in the domain with increments of 0.5 seconds. Please note that the seconds indicate the simulation time, not a real world time. These deadlines specify the time gap between the agent starting from home to clear the mines and the ships starting their journey from one end side of the shore to other.



Figure 2: Scores obtained by the agents in mine clearance domain

Figure 2 shows the scores achieved by the three different agents in all mine density (average of low, medium, and high density) scenarios for the varied deadlines. The X-axis depicts the delay with which the ships start and the Y-axis indicates the number of ships that safely traverse the Q-route. Here, when we start looking at the values from the left side of the graph, at the delay of 0, very few ships were able to traverse the Q-route successfully, for all three agents. Those that were able to reach the other side were able to cross the Q-route in the low mine density scenarios, while very few or no ships made it across in the medium and high mine density scenarios.

To understand what a delay of 0.5 seconds means, consider what each agent can accomplish within that timeframe in a characteristic scenario. After 0.5 seconds, the baseline agent clears the mines in GA1 and is on its way to clear the mines in GA2. At the same point in time, the GDA agent has cleared the mines in GA1 as well as some mines within the Q-route; the eager agent has cleared mines outside of the Q-route and in GA1. After 1 second, the baseline agent has cleared the mines in both green areas and is headed towards home, the GDA agent has cleared some mines within the path from GA1 to GA2 and some mines within GA2, whereas the eager agent is working on the mines within the Q-route after clearing the ones in GA1. In these conditions (delay of 0.5 and 1 seconds), the difference between the performance of the GDA and eager agents does not seem very large. This is because the average of the various mine density scenarios also contains low mine density fields where the two agents perform almost identically, since the eager agent only has a few mines to clear outside of the Q-route. However, the difference in performance for medium and high mine density scenarios



is two ships. At a delay of 1.5 seconds or greater, all agents have performed all clearance tasks intended; thus, performance does not change for delays greater than 1.5 seconds.

### 5.2 Empirical Results in Labor Relations Domain

To assess the performance of the three agents in this domain, we compare the reputation values of all agents after they implement a certain number of policies. There are some numerical values in this domain: Initial reputation is 500, Total budget is \$4000. Implementing any policy reduces the budget by \$25.

The intensity value of a disagreement is a random number between (1, 100). When a disagreement arises, the employees can demand a budget amount which is a random number between (1, 25). Providing the budget amount of any amount within 40% to 60% of the amount demanded by employees solves the disagreement.

If there is no disagreement when a policy is implemented then the reputation of the institution is increased by five. However, if the agent encounters a disagreement, then it has two options ahead of it: to solve the disagreement or to ignore the disagreement and strictly adhere to its initial policy. In the first option that addresses the disagreement, the reputation is neither decreased nor increased, i.e., the change in reputation is zero. If the agent does not address the disagreement, then the reputation value is decreases as a function of intensity value. So, if the intensity is  $\leq 34$  then  $Rep = -Int/100$  and, if the intensity is  $\geq 35$  then  $Rep = -[(n+2)*Int]/100$  where  $n = \text{integer}((I-35)/5)$ . The integer() acts as a rounding function. There is an interest value added to the budget after implementing every 50 policies with a rate of 2.5%. Reputation values can become negative if the agent does not address disagreements.




Figure 3: Scores obtained by the agents in labor relations domain

Figure 3 shows the reputation achieved by the three different agents over 200 policies. The X-axis depicts the number of policies implemented and the Y-axis indicates the reputation value scaled to 10. All the agents have an initial reputation of five. Each point on the lines contains average of 20 policies and the reputation value is cumulative and can reach a max value of 10. In this experiment if any agent is out of budget then it starts to behave as a baseline agent as the debt should be as minimum as possible when the agent implements all policies.

Starting with the baseline agent, at the completion of 50 policies the baseline agent already has a negative reputation because it does not address any of the disagreements, so the reputation value drops and keeps on decreasing monotonically.

The GDA agent gets a little behind the eager agent from 50 to 120 policies because the GDA agent will not address the anomalies with lower intensity values, whereas, the eager agent addresses all the anomalies and spends its budget on every anomaly. However, this means the eager agent is out of its budget much sooner than the GDA agent.

The eager agent is out of budget at 140 policies and starts behaving as a baseline agent, and begins to drop its reputation, while the GDA agent preserves its budget and continues on its increment streak of the reputation value for around 150. The baseline agent's reputation value is still sinking to much lower values but the scale is adjusted such that the negative values are only visible up to -10. These results indicate that the GDA agent should perform better overtime and maintain a higher reputation than the eager agent by a significant amount. The GDA agent underperformed by a negligible amount for a period of time when compared to eager agent because of the higher amount of resources. A smart agent is not needed if the amount of resources present are infinite, but this is not very realistic and as long as resources have a limit, then there will be a need for the GDA agent to use them sustainably.

## 6 Related Research

Statistical anomaly detection has been the subject of extensive research because of its applications to a variety of detection tasks such as network intrusion [18], credit card fraud [2], and malignant tumors from MRIs [27] among many others [4]. Those works rely on large volumes of data to build statistical models of expected patterns. In that context, anomalies correspond to outlier patterns deviating from expected patterns. In our work, our models are planning models and anomalies correspond to deviations of those models. One of the most challenging problems of statistical anomaly detection is the potentially large number of false positives, which trigger unnecessary alarms. In our work, in contrast to the previously mentioned works, explanations for an anomaly is generated to determine the nature of the anomaly and decide if the agent must deal with it. Cox, [8] also presents similar, work on the importance of distinguishing an anomaly from a problem.

The concept of anomaly detection has played a central role in GDA research. In this work, we are focusing on environmental failures since the anomalies are the result of the partial observability in the environment. In Munoz-Avila et al. [22], observed that not all anomalies require triggering a new goal. In that study, the GDA agent is operating in an adversarial environment with a reward function (i.e., the score of the game). A reward function is also used in Jaidee et al [17] which uses reinforcement learning techniques to show GDA knowledge. With both of these, when the current plan is resulting in a positive reward rate, the agent will ignore anomalies. In contrast, in our work, we do not assume a reward function; instead, we generate a causal linkage to determine if a problem underlies the anomaly.

ARTUE [20] is a GDA system that was used to provide control in a Naval strategic simulation of an adversarial and partially observable environment. In this work, explanations were generated using a truth maintenance system that identifies plausible worlds that are consistent with the observations made by the agent and triggers a new goal as a result. ARTUE explains all anomalies, whether problematic or not; goal formulation is responsible for determination of whether the agent should respond. The initial version of ARTUE used rule-based knowledge; extended versions incorporated learning of goal selection knowledge [24] and domain-independent motivations [28] responsible for identifying situations that require response. However, these techniques modified the goal formulation process, rather than incorporating a separate problem recognition step prior to explanation generation.

Other kinds of explanations also exist; external explanations describe an agent's anomalous behavior to others [12], while internal explanations hypothesize the cause of an anomaly for its own needs [1, 21]. In our work we use a variant of the internal explanations called problem explanations.

More recently, the notion of GDA agent's expectations has been extended to consider only the necessary effects of the plan executed so far as opposed to considering the whole state [9]. Our work can use this form of expectations.

Our work is motivated by work on introspective reasoning, where the agent reasons about the decisions that lead to actions taken and how these actions affect the environment. Meta-AQUA [7] reasons about the processes that lead to a decision which resulted in an anomaly and considers three types of anomalies: novel situations, incorrect background knowledge and mix-indexed knowledge structure; the difference between the last two is that in the latter the agent has the knowledge but it is not retrieved in the appropriate circumstances. Fox and Leake [13] present a mechanism to fix these retrieval mechanisms using introspective reasoning techniques. In our work, we are focusing on novel situations when there is an expectation failure.

## 7 Conclusion and Future Research

We have described a formalism for agents that enables them to distinguish between those anomalies that they must deal with from those that they do not. Real world scenarios often deal with deadlines and it is practically impossible for an agent to worry about all the anomalies it comes across, reason, react and achieve its prime goals within the given deadline. Although our experiment setting is simulated, adding a deadline to our experiment clearly shows that the performance of the GDA agent is better than the eager and the baseline agents.

For future work, we would like to work on several different enhancements that can improve the performance and reasoning capabilities of the GDA agent in our future research. First, adding an importance factor to the problem formalism would help the agent to prioritize anomalies that are classified as a problem with the goals it possesses. Moreover, given that an agent only has finite resources, prioritizing the anomalies could also assist an agent to delegate goals to other agents. Second, adding goal monitors [10] after formulating goals could help an agent to adapt as the world changes. For example, during mine clearance, if the establishment of a Q-route changed from one location to

another then it is highly likely that it would not need to continue to clear mines along the originally proposed Q-route, or if the budget can be created through profits or donations. Finally, if the number of anomalies flagged were excessive given what might be anticipated in a particular context, then this could serve as a cue for the agent to generate a goal with a broader scope than the current goal. For example, our experimental setting has around ten mines within the proposed Q-route. Instead of clearing just the mines on the agent's path from GA1 to GA2, if the number of mines encountered were too great, then the agent could generate, or delegate to another agent, a goal to survey the entire region between GA1 and GA2. Finally, in the labor relations domain, we would like to explore adding other factors such as customer satisfaction, profits, or time and possibly go a step further and create a multiagent scenario.

## Acknowledgements

This research was supported by NSF under grant 1849131 and by ONR under grant number N00014-18-1-2009. We thank the reviewers for the comments and suggestions.

## References

1. Aamodt, A.: Explanation-driven case-based reasoning. *European Workshop on Case-Based Reasoning* (pp. 274–288). Heidelberg, Berlin: Springer (1993).
2. Aleskerov, E., Freisleben, B., Rao, B.: Cardwatch: A neural network based database mining system for credit card fraud detection. *Proceedings of the IEEE/IAFE 1997 computational intelligence for financial engineering* (pp. 220–226). IEEE (1997).
3. Bengfort, B., Cox, M. T.: Interactive Knowledge-Goal Reasoning. 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 (2015).
4. Chandola, V., Banerjee, A., Kumar, V.: Anomaly detection: A survey. *ACM computing surveys*, 41(3), 15 (2009).
5. Cox, M. T.: Perpetual self-aware cognitive agents. *AI magazine*, 28(1), 32 (2007).
6. Cox, M. T., Dannenhauer, D., Kondrakunta, S.: Goal operations for cognitive systems. *Proceedings of the Thirty-First AAAI Conference on Artificial Intelligence* (2017).
7. Cox, M. T., Ram, A.: Introspective multistrategy learning: On the construction of learning strategies. *Artificial Intelligence*, 112(1–2), 1–55 (1999).
8. Cox, M. T.: The Problem with problems. In proceedings of eighth Annual conference on Advances in Cognitive Systems (2020).
9. Dannenhauer, D., Munoz-Avila, H.: Raising expectations in GDA agents acting in dynamic environments. *Proceedings of the Twenty-Fourth International Joint Conference on Artificial Intelligence* (2015).
10. Dannenhauer, Z. A., Cox, M. T.: Rationale-based perceptual monitors. *AI Communications*, 1–16 (2018).
11. Eyorokon, V. B., Yalamanchili, P., Cox, M. T.: Tangent Recognition and Anomaly Pruning to TRAP Off-Topic Questions in Conversational Case-Based Dialogues. *International Conference on Case-Based Reasoning* (pp. 95–109). Springer (2018).

12. Floyd, M. W., Aha, D. W.: Incorporating transparency during trust-guided behavior adaptation. *International Conference on Case-Based Reasoning* (pp. 124–138). Springer (2016).
13. Fox, S., Leake, D.: Using introspective reasoning to refine indexing. *Proceedings of the Thirteenth International Joint Conference on Artificial Intelligence*, 391397 (1995).
14. Ghallab, M., Nau, D., Traverso, P.: *Automated Planning: theory and practice*. Elsevier (2004).
15. Gogineni, V. R., Kondrakunta, S., Brown, D., Molineaux, M., Cox, M. T.: Probabilistic Selection of Case-Based Explanations in an Underwater Mine Clearance Domain. In *International Conference on Case-Based Reasoning* (pp. 110-124). Springer, Cham (2019).
16. Gogineni, V. R., Kondrakunta, S., Molineaux, M., Cox, M. T.: Case-Based Explanations and Goal Specific Resource Estimations. In *The Thirty-Third International FLAIRS Conference* (2020).
17. Jaidee, U., Muñoz-Avila, H., Aha, D. W.: Integrated learning for goal-driven autonomy. *Proceedings of the Twenty-Second International Joint Conference on Artificial Intelligence* (2011).
18. Kumar, V.: Parallel and distributed computing for cybersecurity. *IEEE Distributed Systems Online*, 10, 1 (2005).
19. Li, P. C.: *Planning the optimal transit for a ship through a mapped minefield*. Masters Thesis, Dept Of Operations Research, Naval Postgraduate School, Monterey, CA (2009).
20. Molineaux, M., Klenk, M., Aha, D.: Goal-driven autonomy in a Navy strategy simulation. *Proceedings of the Twenty-Fourth AAAI Conference on Artificial Intelligence* (2010).
21. Molineaux, M., Kuter, U., Klenk, M.: DiscoverHistory: Understanding the past in planning and execution. *Proceedings of the 11th International Conference on Autonomous Agents and Multiagent Systems Volume 2* (pp. 989–996). International Foundation for Autonomous Agents and Multiagent Systems (2012).
22. Munoz-Avila, H., Aha, D. W., Jaidee, U., Klenk, M., Molineaux, M.: Applying goal driven autonomy to a team shooter game. *Proceedings of the Twenty-Third International FLAIRS Conference* (2010).
23. Panjala, U. S., Eyorokon, V., Cox, M.: A Mixed-Initiative Approach to Knowledge Discovery. *Proceedings of the International Conference on Cyber Warfare and Security* (p. 533). Academic Conferences International Limited (2017).
24. Powell, J., Molineaux, M., Aha, D. W.: Active and interactive discovery of goal selection knowledge. *Proceedings of the Twenty-Fourth International FLAIRS Conference* (2011).
25. Ram, A.: Indexing, elaboration and refinement: Incremental learning of explanatory cases. *Case-Based Learning* (pp. 7–54). Boston, MA: Springer (1993).
26. Schank, R. P.: *Explanation patterns: Understanding mechanically and creatively*. Psychology Press (2013).
27. Spence, C., Parra, L., Sajda, P.: Detection, synthesis and compression in mammographic image analysis with a hierarchical image probability model. *Proceedings IEEE Workshop on Mathematical Methods in Biomedical Image Analysis* (pp. 3–10). IEEE (2001).
28. Wilson, M. A., Molineaux, M., Aha, D. W.: Domain-independent heuristics for goal formulation. *The Twenty-Sixth International FLAIRS Conference* (2013).