Goal Reasoning: Foundations, Emerging Applications, and Prospects

David W. Aha

2018. *AI Magazine*, vol. 29, no. 2, pp. 3-24.

Presented by Leen-Kiat Soh

Abstract

Goal reasoning has a bright future as a foundation for the research and development of intelligent agents

- the study of agents that can deliberate on and self-select their objectives, which is a desirable capability for some applications of deliberative autonomy
  - This capability is of interest to several AI subcommunities and applications
- Increasing attention on the importance of how agents reason about goals (e.g., AI safety)

Aha's group has focused on how goal reasoning can assist with controlling autonomous systems

Outline

Introduction on Goal Reasoning

Related Work

Summary of Work from Aha’s group

- Foundations
- Emerging Applications

Current and future research directions

Introduction

Goal reasoning (GR) as the process by which intelligent agents continually reason about the goals they are pursuing, which may lead to goal change

- This flexibility may allow them to behave *competently* when they are *not preencoded* with a model that dictates what goals they should pursue in *all* encounterable situations

Introduction: GR Agents

GR agents can

- deliberate on a space of goals,
- dynamically adjust goal priorities, and
- perform goal-management functions (e.g., formulation, commitment, and suspense)

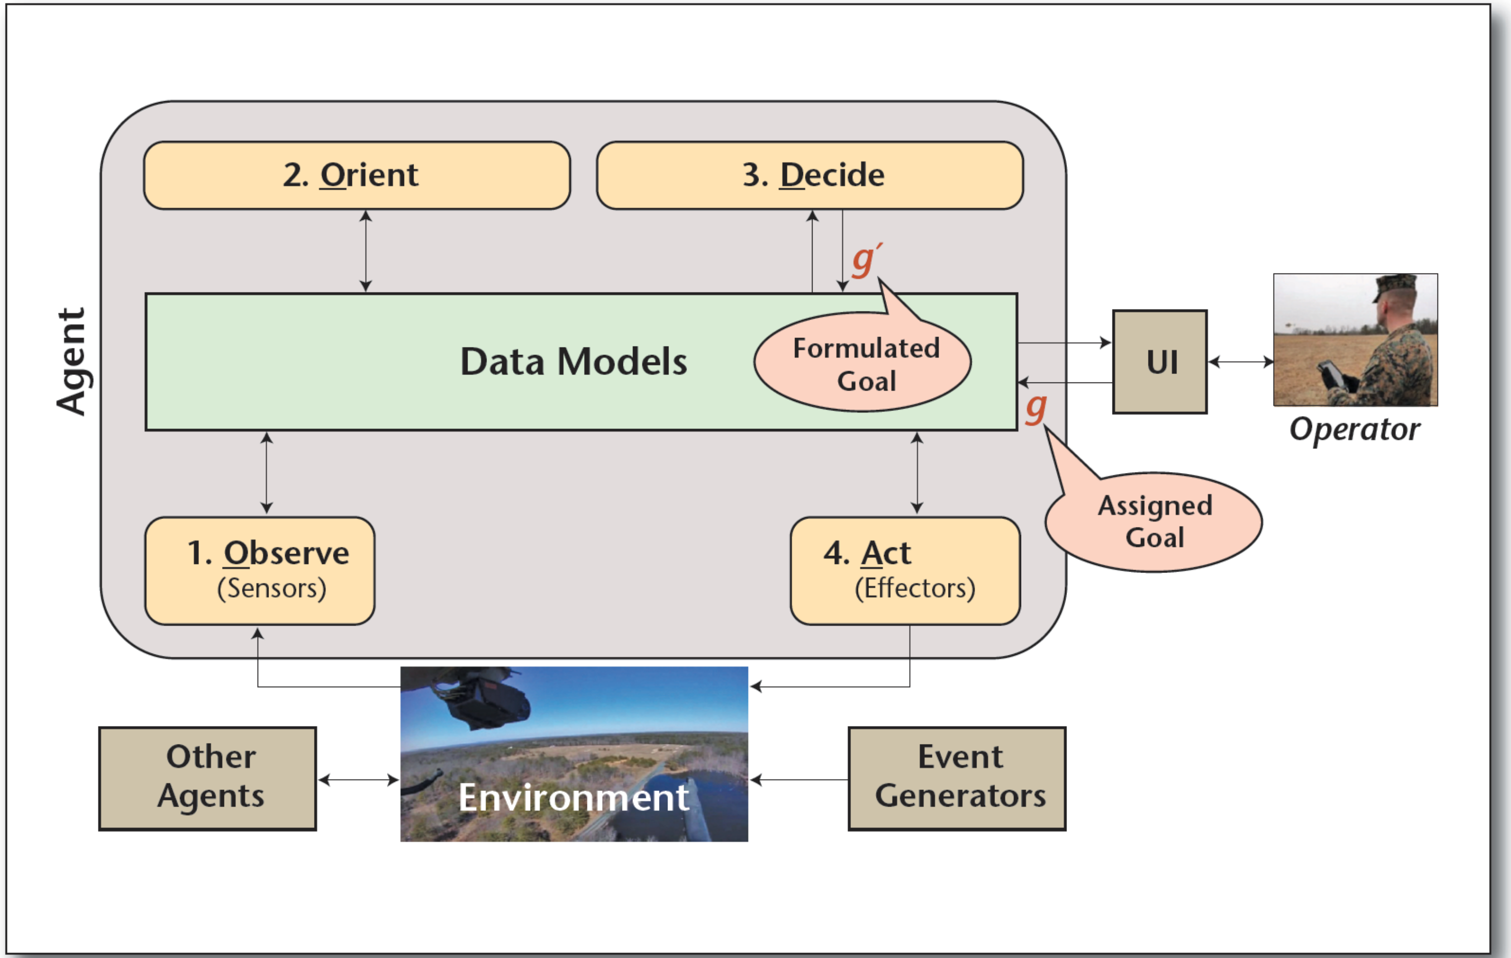

Figure 1. Goal Reasoning Agents Can Formulate Their Own Goals.

Introduction: Dimensions for Goals

Type: Goals can be declarative (referring to belief states) or procedural (referring to actions)

Specificity: Goals may refer to a concrete instance or an abstraction (for example, region of belief states, sequence of actions)

Duration: Goals may refer to a static time point or be durative

Purpose: Some goals are designed to learn world knowledge (i.e., query or knowledge goals), while others are attainment goals (i.e., they *exploit* such knowledge)

Condition: Goals can be unconditional, or conditioned on beliefs or other goals

Persistence: Goals may, or may *not*, be interruptible

Introduction: Aha's Focus

Focus on goals that are declarative, that are specific points in belief space *not* involving knowledge acquisition, that are unconditional and that can be interrupted

Introduction: Complex Environments

GR agents are intended for complex environments

Environment Dimension	Simple	Complex
Operator Availability	Constant	Intermittent or Inaccessible
Goal Model	Complete	Partial
Accessibility	Full	Partial
Updates	Static	Dynamic
Action Effects	Discrete	Continuous
Action Outcomes	Deterministic	Stochastic
Agents	Single	Multiple

Table 1. Goal Reasoning Agents Are Most Appropriate for Complex Environments.

Introduction: GR Agents?

GR agents are ***not*** intended for all environments and scenarios

- If the agent's human operator is always available, then they could potentially provide continuous control, alleviating any need for agent self-control: *GR is not relevant*
- If the agent is given a complete function that determines what goal should be pursued for all encounterable situations, then there is no need to perform dynamic inference in support of goal reasoning: *GR becomes a retrieval task*
- Their ability to perform goal reprioritization is *not* useful unless they encounter situations that warrant goal reprioritization: e.g., impasses or affordances requiring goal deliberation

Related Work

Symbolic Planning

Cognitive Architectures

Intelligent Agents

Related Work: Symbolic Planning

Most research on symbolic task planning pertains to the following problem: given initial and goal states from a set of states, and a model of actions that can be applied to traverse among these states, generate a plan that can be applied in the initial state to traverse into the goal state

- **no monitoring of the plan's execution takes place**, and **the agent cannot change the goal**

Newer formulations:

- Continual planning: human operators can provide an agent with additional goals during run time
- Oversubscription planning: planner must reason about which among conflicting goal(s) it should attempt to achieve
- Conditional goals by reasoning about trade-offs among sensing costs and goal rewards

One perspective is that GR is a methodology for plan monitoring in the context of planning and acting

Related Work: Cognitive Architectures

Solutions for GR have frequently been included in cognitive architectures

- **SOAR**: universal subgoaling provides a process for responding to impasses during problem solving by posting a new subgoal to solve
  - E.g., TacAir-Soar was provided with a top-down goal hierarchy that encodes doctrine, missions, and tactics for its simulated air vehicles to perform, along with a bottom-up hierarchy of rules to guide interrupt processing
  - E.g., use of appraisal theories in SOAR to support GR processes
- **ICARUS** extended: Nominate top-level goals (from a long-term memory of general goals) and continuously manage them through a prioritization function
- **Act-R** + model: Replace its architectural goal stack for managing goals, and show that goal-directed behavior can be explained using mechanisms of goal activation and associative priming
- **MIDCA**: Model a metacognitive process for goal change and a cognitive process for goal generation to manage unexpected events in dynamic environments

Related Work: Intelligent Agents

Many GR contributions have been proposed in the context of intelligent agents

- A perpetual self-aware cognitive agent that was designed for continuous autonomous operation in complex environments: integration of planning, execution, and goal generation components
- Use an agent's context (i.e., its environment situation) to constrain goal selection and prioritization
- Motivated agents based on **BDI** architectures have addressed the representation of goal types, their properties, and reasoning lifecycles
  - MADbot: an agent that dynamically generates goals in response to its internal motivation model
- Many agent programming languages support automated planning in the context of BDI architectures

Foundations

Their GR research began with investigations of **goal driven autonomy** (GDA), which is a simple anomaly-driven agent model

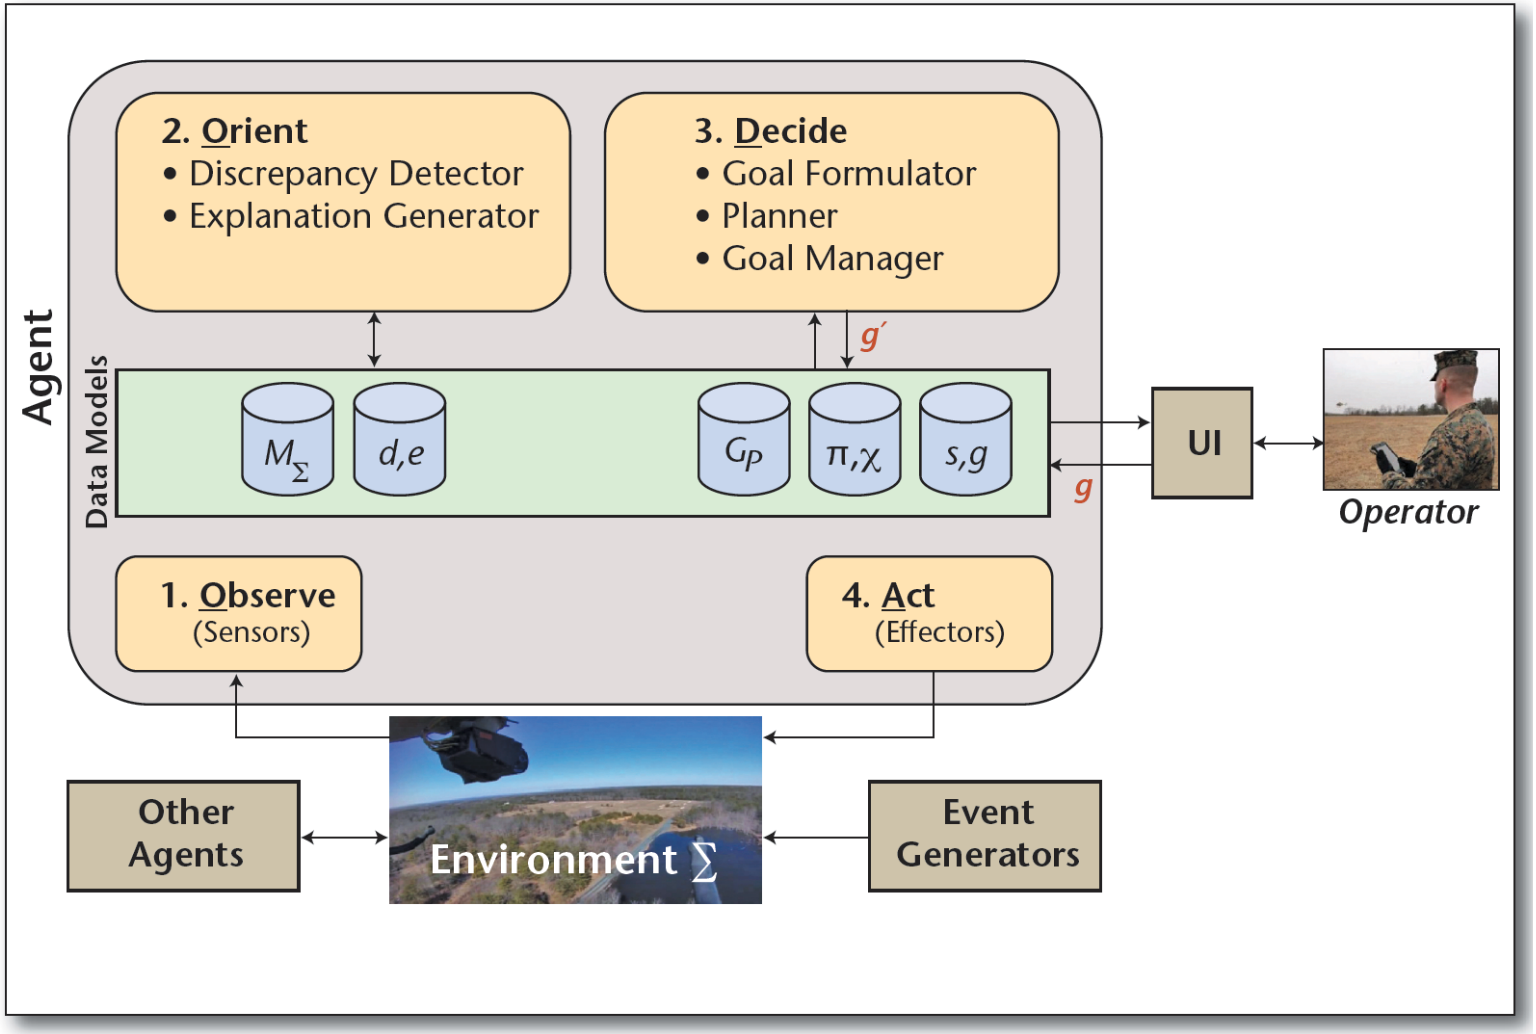

Figure 2. A Depiction of the Goal-Driven Autonomy (GDA) Model of Goal Reasoning.

Foundations: Key Steps

Step 1, *discrepancy detection*, compares the observations obtained from executing action a_i in belief state s_i with expectation χ_i (that is, this tests whether any constraints are violated, corresponding to unexpected observations). If a discrepancy d is found, then it is given to step 2.

Step 3, *goal formulation*. Given d, e, and $s_{(i+1)}$, this process generates a goal $g' \in G$ (not shown).

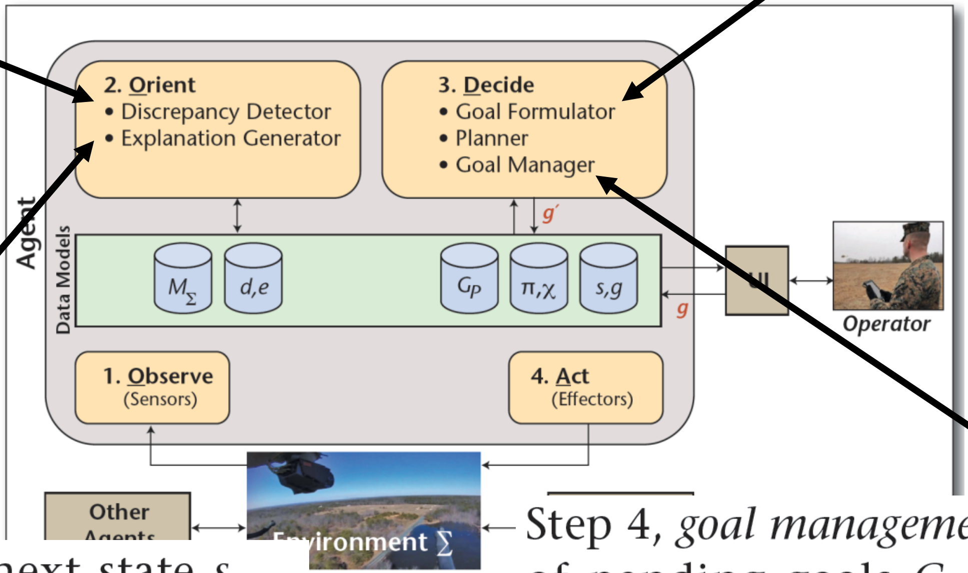

Step 2, *explanation generation*. Given next state $s_{(i+1)}$ (provided by γ), s_i, and d, this process hypothesizes an explanation $e \in E$ (not shown in figure 2) of its cause.

Step 4, *goal management*. Given a set (initially empty) of pending goals $G_P \subseteq G$ and g', this process may update G_P and will select the next goal g to feed to the planner.

Foundations: Limitations & Next?

While GDA can model simple GR processes, it does *not* explicitly model

- goal constraints,
- the relation of goals to tasks for achieving them, or
- processes for suspending or revising goals whose plans are not executable

This limitation motivated the development of a more comprehensive process model for GR

Foundations: Goal Refinement

Goal refinement, an extension of plan refinement that models the progressive refinement of goals through the *addition of constraints*

- Goal refinement can represent the context in which a goal is pursued by a GR agent

Aha's goal refinement process model is the **Goal Lifecycle**

Foundations: Goal Lifecycle

This model transitions a goal node (that is, a pairing $GN = (g, C)$ of goal g with constraint set C) through increasingly detailed modes (e.g., formulated, selected) by applying constraint-refinement strategies that progress goal nodes toward completion

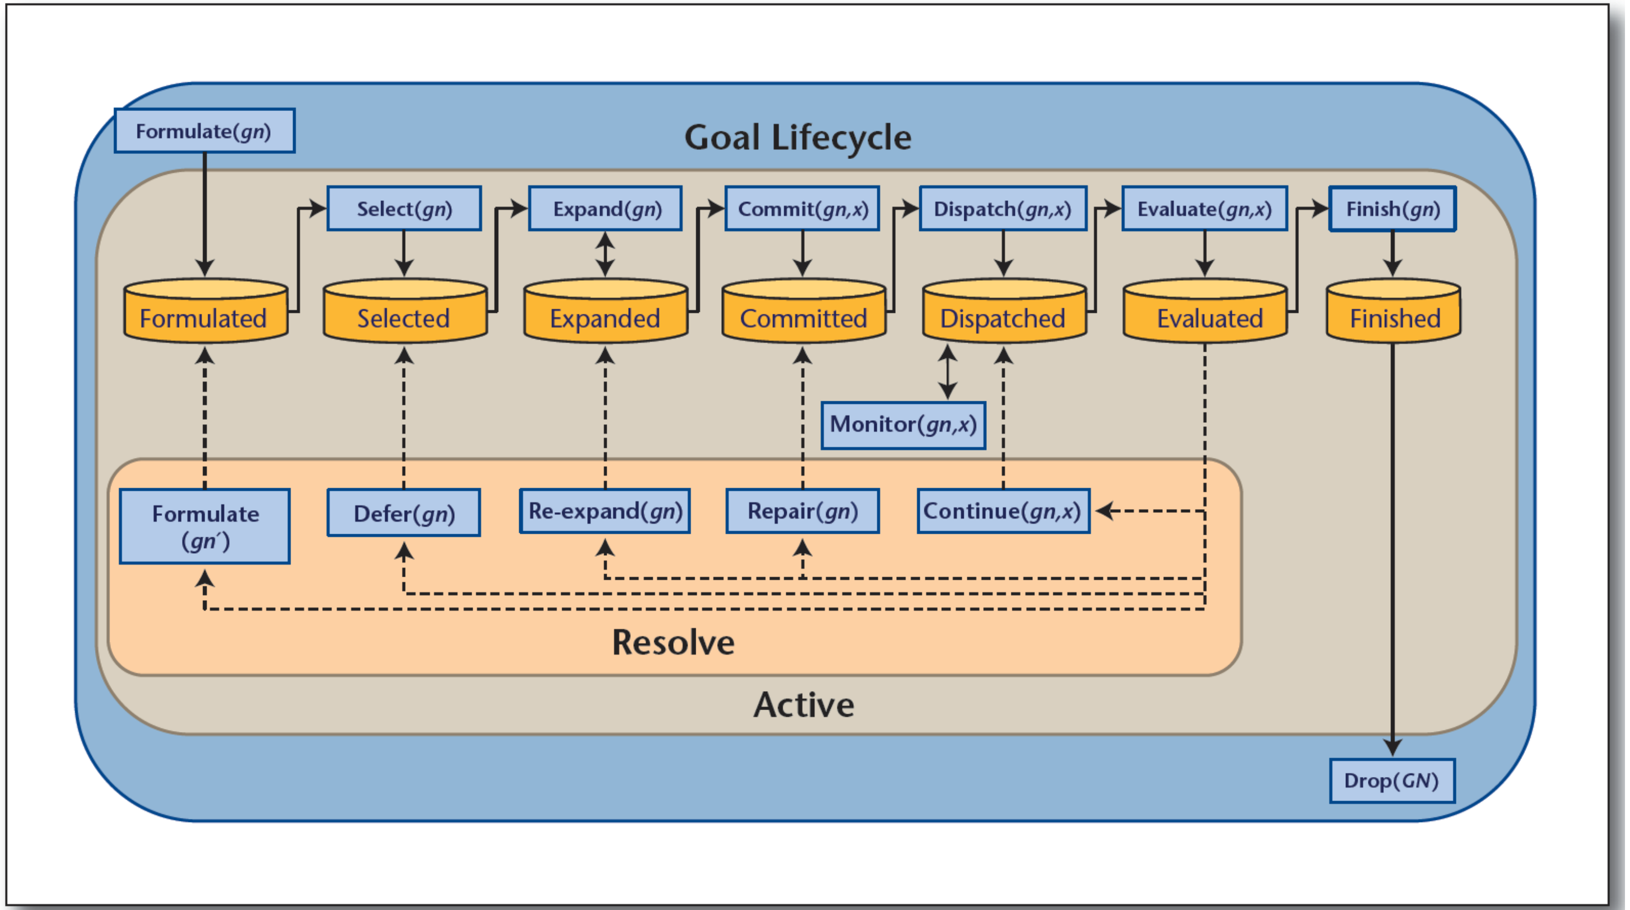

Figure 3. The Goal Lifecycle — A Goal Refinement Model of Goal Reasoning.

Foundations: Goal Lifecycle

Constraint Refinement Strategies

Formulate creates a new goal node and enters it into the Goal Lifecycle by defining its initial constraints, criteria, and prerequisites

Select chooses which goal(s) to actively pursue; it ensures that the goals' prerequisites are met and that the agent has the resources to pursue them

Expand generates a set of expansions *X* (e.g., plans, decompositions of nonprimitive goals, or trajectories of primitive goals) to achieve a goal *g* in goal node GN, and a set of expectations for each

Commit picks an expansion $x \in X$ to pursue from those generated by expand

Dispatch executes the committed expansion and defines the criteria by which *g* can be evaluated during execution

Foundations: Goal Lifecycle Inside GR Agent

Captures decision points during a goal's activation,

- as a set of ***decide*** subprocesses (figure 4) subsuming those in figure 2

A new data structure (GN) to record substantial information associated with each goal node:

- Goal
- Associated constraints, mode
- Selected expansion/plan
- Plan expectation
- Associated discrepancies
- Etc.

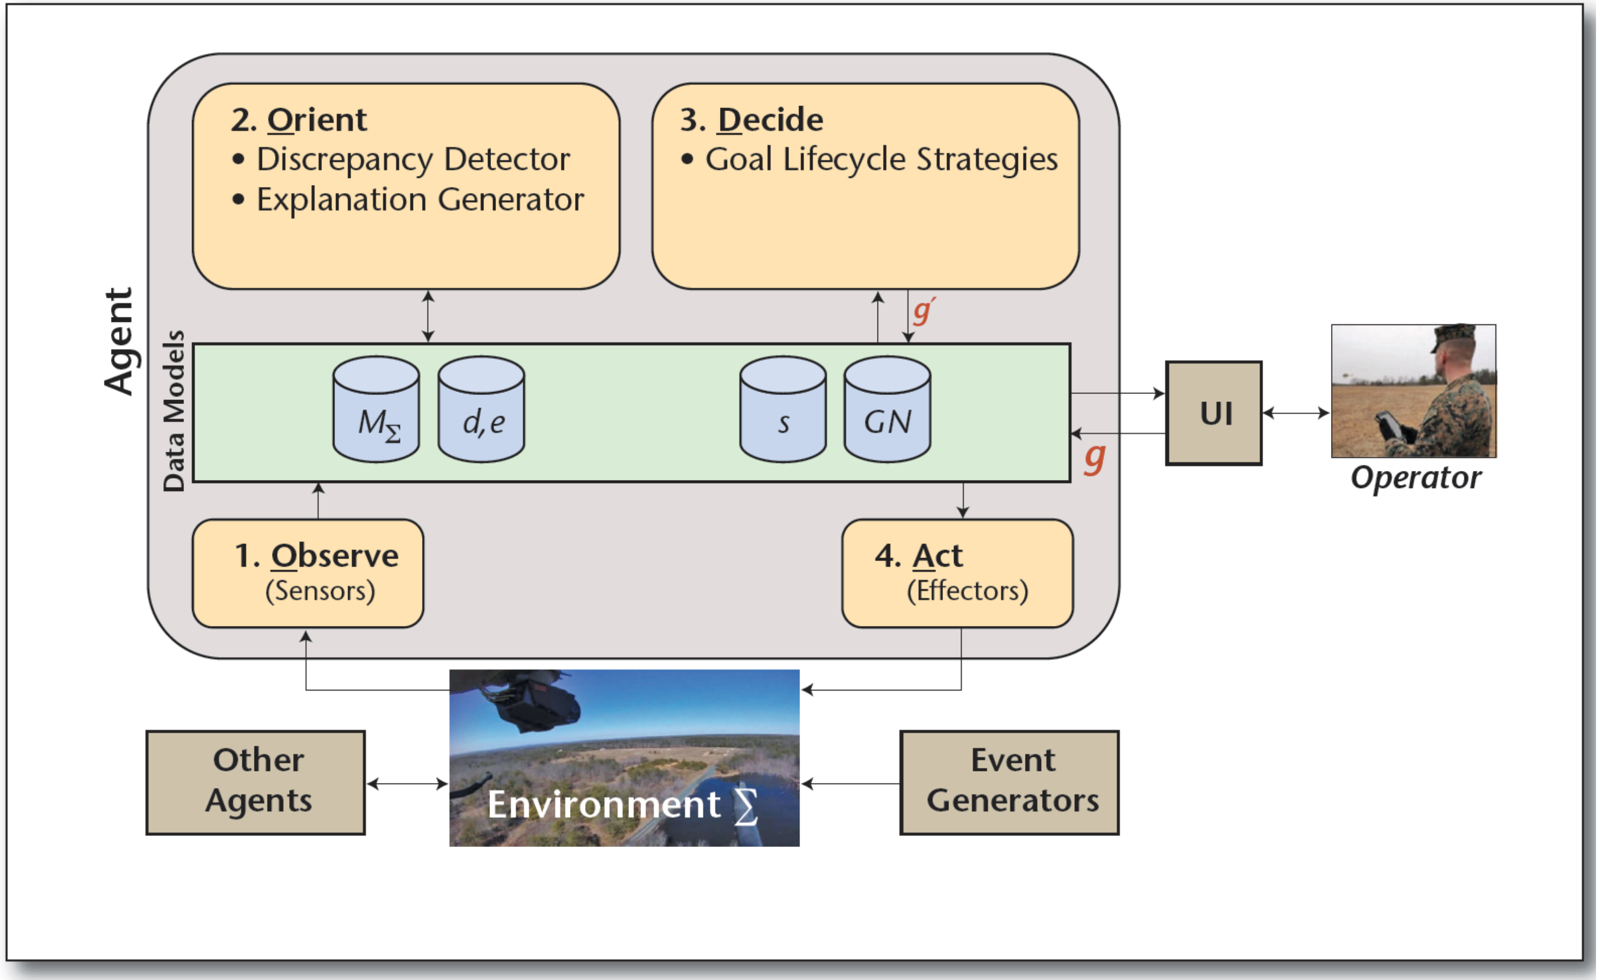

Figure 4. A Depiction of a Goal Lifecycle Model of Goal Reasoning.

Foundations: Summary

The Goal Lifecycle provides a formal structure for goal refinement, such that the GR agent can deliberate on and adapt its goals in response to dynamic and unpredictable events

In Aha’s application of GR, their GR agents employ variants of GDA or more comprehensive Goal Lifecycle models.

Emerging Applications

Underwater Vehicles

Beyond-Visual Range Air Combat

Foreign Disaster Response

Emerging Applications Unmanned Underwater Vehicles

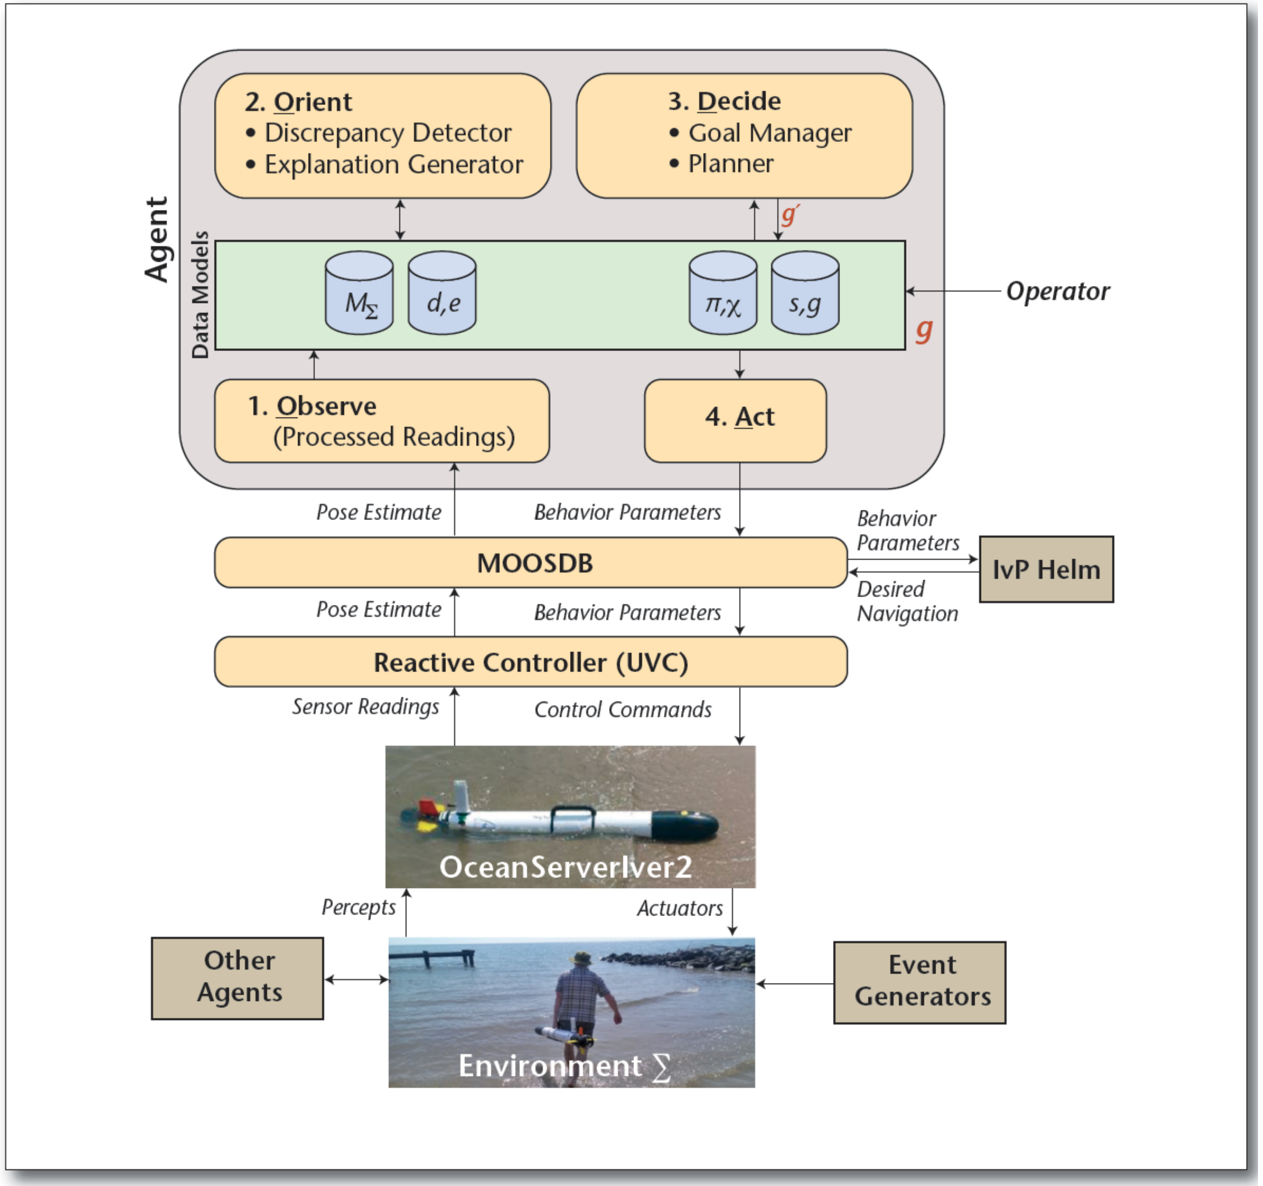

Figure 5. A GDA Model of Goal Reasoning for Unmanned Underwater Vehicle (UUV) Control.

Emerging Applications
Unmanned Underwater Vehicles

Initial Test

In simulated neutral and hostile mode scenarios, where only the latter included active sonar pings

The simulated USV starts in the center of the UUV's target survey region, picks a heading to one of the two endpoint regions, loiters for a specified time (to ensure the UUV encounters it), and then begins traversing

The USV emits an engine noise with a detectable radius

Meanwhile, the UUV departs its launch point toward its survey region

It detects the unexpected engine noise and interprets it as a discrepancy that triggers the explanation generator to identify that a contact is within range, but without detecting a (hostile) ping, no new goal will be formulated

In contrast, when a ping is encountered, the explanation generator concludes that there is a hostile vehicle within range, and goal formulation recommends (with high priority) a goal to retreat to the safe point

When the pinging is no longer detected, the explanation generator will conclude that there is no hostile vehicle in that region

At this time, goal generation directs the UUV to resume its prior goal, and it completes its mission

In 25 trials, in which the mode was randomly varied along with other independent variables (for example, the simulated USV's route), the UUV responded appropriately each time

Emerging Applications
Unmanned Underwater Vehicles

Actual Test

From six trials they collected data with the UUV traversing at the sea surface or maintaining a depth of 0.75 meters

Equal numbers of trials were used for hostile and neutral USVs

In each trial, the UUV correctly detected the USV and its active pinging (for the hostile condition), and reacted by explaining the discrepancy and formulating the correct goal in response

Due to the calm marine environment, they did not encounter significant positional sensor drift during the relatively short mission

Generously modeling their expectations using PHOBOS's ranged values provided sufficient tolerance for the noise that we did observe in the surface and underwater trials

Emerging Applications
Beyond-Visual Range Air Combat

Beyond-Visual Range (BVR) air combat (or, simply, BVR) is a modern form of air-to-air warfighting

In BVR, opposing teams of aircraft engage over large distances (that is, over 100 kilometers), where each team attempts to destroy their enemy (using active radar homing missiles with ranges of approximately 50 kilometers) or to force them to retreat

- Similar to close-range dogfighting, BVR engagements can involve multiple aircraft (teammates and adversaries) operating in a contested airspace.

BVR is less reactive and involves more deliberation, with positioning and timing being more important than motion planning

- continuous, partially observable (due to limited sensor ranges), and noisy (due to sensor errors)
- aircraft behaviors must satisfy tight real-time constraints to evade opponent attacks and avoid dangerous maneuvers (e.g., flying too low, colliding with teammates)
- Substantial uncertainty, as the adversary’s assets (air and ground), configurations, and preferred tactics are not always known a priori
- Highly dynamic (e.g., the battle situation can change rapidly leading to goal or mission changes)

Thus, controlling wingmen UAVs in future mixed human-UAV BVR teams motivates the development of GR-controlled agents.

Emerging Applications
Beyond-Visual Range Air Combat

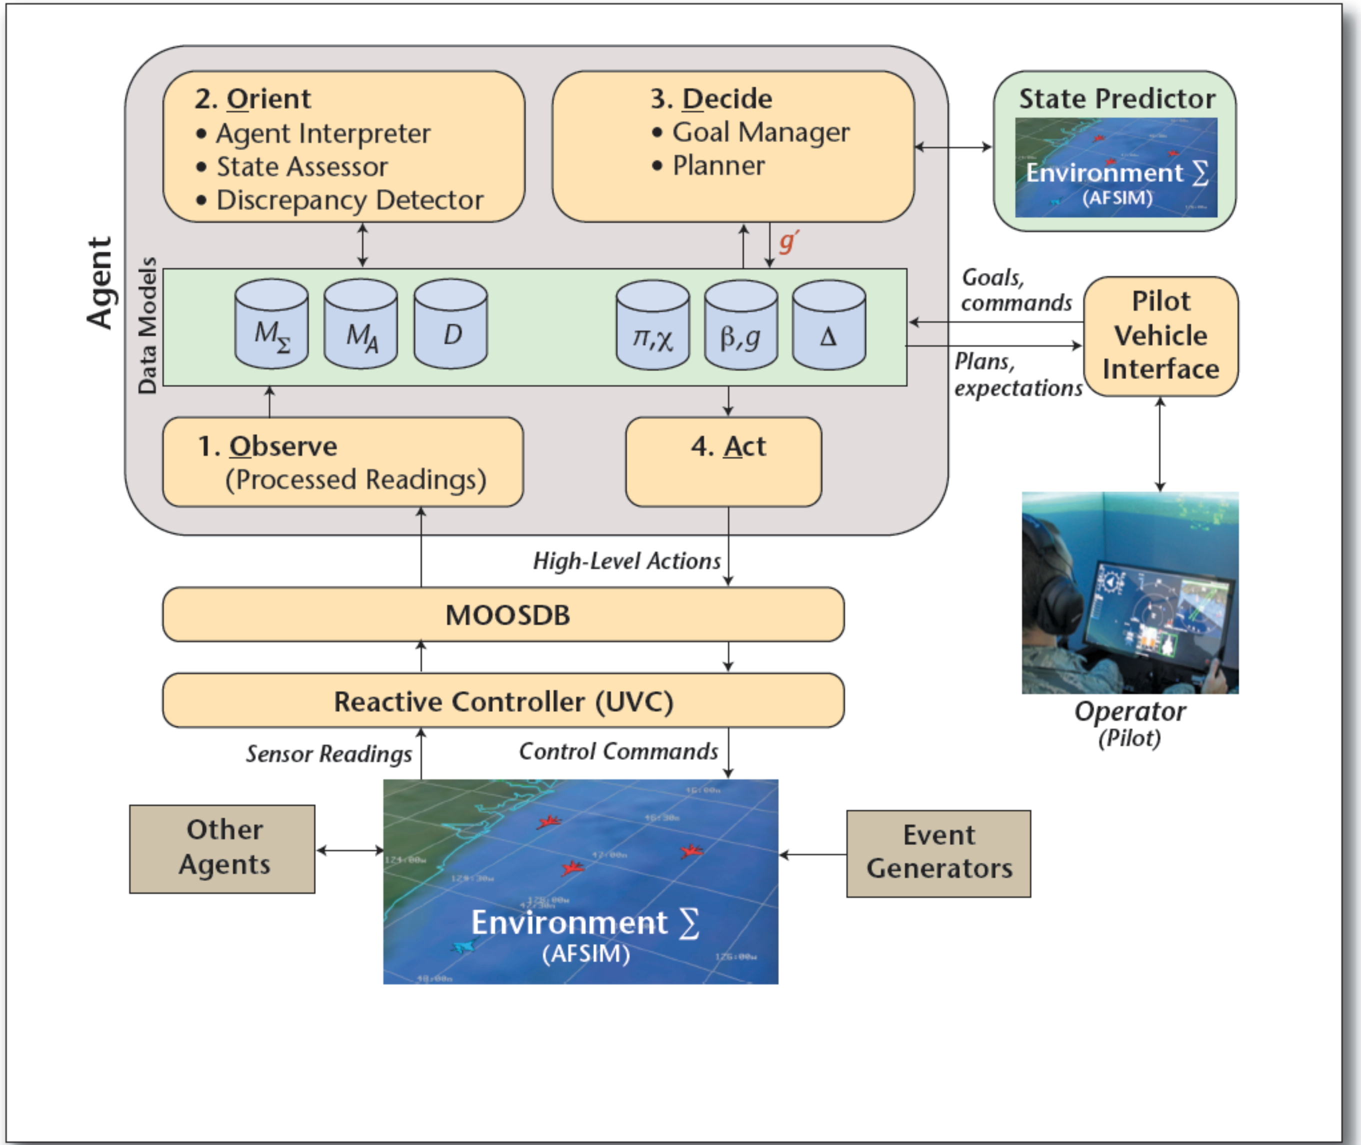

Figure 8. A Modified GDA Model of Goal Reasoning for Controlling a UAV Wingman in a Mixed Human / UAV Beyond-Visual-Range (BVR) Air Combat Team.

Emerging Applications
Beyond-Visual Range Air Combat

Discrepancies

- Incoming Missile: unexpected hostile missiles (which allows the system to dynamically respond to an attack and attempt to evade the missile)
- Model Changed
- Flanking Hostile
- Expectations Violated: violations of any of the current plan's expectations, generated by the state predictor
- Out of Ammo
- Low on Fuel: whether that resource is running low
- Opportunistic Target

Emerging Applications
Foreign Disaster Response (FDR)

An FDR mission's objective is to provide, across the globe, humanitarian aid after a natural disaster strike, when many lives can be in peril and first responders must react quickly

These missions can benefit from a heterogeneous team of UAVs an unmanned ground vehicles (UGVs):

- that rapidly survey the area
- that identify key locations (e.g., of survivors, damaged infrastructure) and traversable routes for ingress and egress,
- that locate VIPs, and
- that serve as mobile communication relays

Human coordination can be challenging

- as commands must be translated to actions by appropriate team members, the team must keep the human operator informed, and the robots must react intelligently to changes in their environment or their internal state

These dynamic conditions may cause them to change their tasks or even their objectives (e.g., if the current one is unachievable)

Emerging Applications
Foreign Disaster Response

To address this problem, they designed the **situated decision process** (SDP):

- under operator guidance, uses a GR approach to control and coordinate a robot team (i.e., managing and executing their goals)
- uses a **centralized control** approach that provides commands to independent vehicles
- to be used by a forward operating base

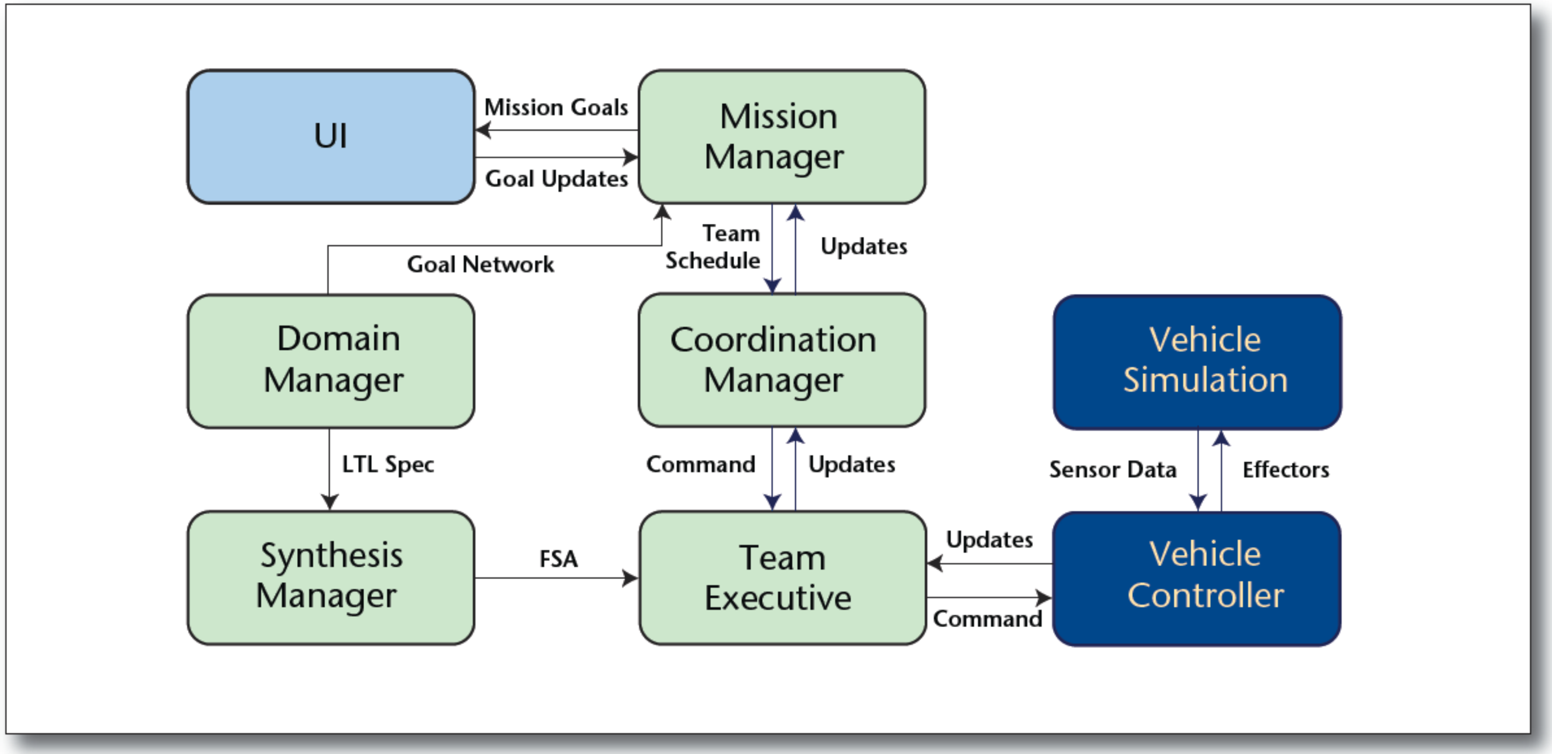

Figure 11. Conceptual Design of the Situated Decision Process (SDP).

The Mission Manager performs goal reasoning, creating a schedule of actions for a team of vehicles, each of which executes a synthesized FSA.

Prospects

GR can be a foundation for (unembodied) proactive decision aids in collaborative decision making contexts

- Observing how military command staff interact during a mission and supporting them

A more sophisticated GR agent would need to perform automated scene understanding and situation assessment

- Proposed: *Integrated deep learning (for image recognition) and natural language understanding techniques with GR agents*

The GR agent would also need to detect and reason about observations not anticipated by its action, event, or agent models

- Proposed: *DiscoverHistory, which performs continuous explanation generation by using heuristic constraints to search the space of plausible explanations*

Prospects

The GR agent would also need extensions for more robust decision making

For example, current GR agents that instantiate the Goal Lifecycle assume there exists *only* a single goal node, and only one algorithm for each of its strategies (for example, selection, expansion)

This approach is limiting: Why not compare the utilities of different goal nodes before selecting which one to process, and not necessarily discard ones that are not immediately selected?

Also, different algorithms for a lifecycle strategy (e.g., planners for the expansion strategy) may be appropriate for different problem solving contexts and could be made available to a GR agent for selection

Thus, they are investigating a metareasoning method for selecting a goal node (or nodes) and a strategy algorithm (or algorithms) to apply

They are also devising methods that reason about how to proceed if a goal cannot be processed in its current form (e.g., can more specific versions of it be processed that would satisfy a human operator's intent?).

Prospects

A GR agent that serves as a proactive intelligent decision aid must be able to **explain its models**, **its reasoning for a recommendation** (and other decisions), and **the expected outcomes from applying a recommendation to its human teammates**

That is, a GR agent should be transparent so that its teammates can calibrate their trust in it and, in doing so, make appropriate decisions

- a user interface designed to expose a GR agent's models and reasoning can increase an operator's **situation awareness** and **trust** in the agent's decision making
- **Explainable AI:** develop and assess the utility of explainable GR agents in human-agent teaming contexts