

Guest Editorial

Special issue on goal reasoning

Mark Roberts ^{a,*}, Daniel Borrajo ^b, Michael Cox ^c and Neil Yorke-Smith ^{d,e}

^a *Naval Research Laboratory, USA*

^b *Universidad Carlos III de Madrid, Spain*

^c *Wright State University, USA*

^d *Delft University of Technology, The Netherlands*

^e *American University of Beirut, Lebanon*

Goals are a unifying structure across the variety of intelligent systems, and reasoning about goals takes many forms. In the most encompassing view, intelligent systems use goal structures (or goal rewards) to manage long-term behavior, anticipate the future, select among priorities, commit to action, generate expectations, assess tradeoffs, resolve the impact of notable events, and learn from experience. As a result, studies of goal reasoning appear in diverse subfields of AI such as motivated systems, cognitive science, automated planning, and agent-oriented programming to name but a few.

A community centered on this topic has conducted a series of workshops since 2010. The workshop series was first held at AAI 2010 (eleven submissions). It was held twice at the Advances in Cognitive Systems conference, in 2013 (eleven submissions) and 2015 (fourteen submissions). The workshop in 2016 moved to IJCAI (fourteen submissions) and continued in 2017 at IJCAI as well (fifteen submissions). This special issue collects extended versions of papers from the 4th Goal Reasoning Workshop held at IJCAI in 2016. Of the fourteen original submissions, seven have been extended for this special issue:

Anticipation of Goals in Automated Planning by Raquel Fuentetaja, Daniel Borrajo and Tomás de la Rosa examines how to anticipate the arrival of goals in on-line continual problems. The work described focuses on performing this task in a domain-independent manner, and the authors show that their approach can outperform reactive planning approaches in benchmark problems and a UAV scenario from prior work.

Learning-driven Goal Generation by Alberto Pozanco, Susana Fernández and Daniel Borrajo examines how to learn models that predict when goals will appear, which allows the planning process to consider current and future goals. The results demonstrate how the proposed approach works in a grid world domain inspired by unmanned aerial vehicles.

Goal Reasoning for Autonomous Underwater Vehicles: Responding to Unexpected Agents by Mark Wilson, James McMahon, Artur Wolek, David Aha, and Brian Houston examines how to embed Goal Driven Autonomy, a model of goal reasoning, in a small underwater vehicle. This paper shows that the system can help the vehicle respond to a dynamic environment. To our knowledge, this is the first example of a this approach running on an underwater vehicle.

Learning Task Hierarchies Using Statistical Semantics and Goal Reasoning by Sriram Gopalakrishnan, Héctor Muñoz-Avila, and Ugur Kuter examines how to learn hierarchical planning models from plan traces using semantic text analysis techniques. The authors demonstrate that their approach can effectively learn a planning model for a logistics domain.

Distributed Discrepancy Detection for a Goal Reasoning Agent in Beyond-Visual-Range Air Combat by Justin Karneeb, Michael Floyd, Philip Moore, and David Aha examines goal reasoning to assist pilots in beyond-visual range combat with a particular focus on detecting and responding to discrepancies. The proposed discrepancy management system is shown to improve mission success over a baseline.

Rationale-based Perceptual Monitors by Zohreh Dannenhauer and Michael Cox examines the relation-

* Corresponding author. E-mail: mark.roberts@nrl.navy.mil.

ship between planning, interpretation, and goals. This paper shows how an agent can track changes in its environment that impact executing plans and existing goals. The authors demonstrate that the proposed approach can effectively respond to changes in several planning benchmarks.

Investigating the Solution Space for Online Iterative Explanation in Goal Reasoning Agents by Christine Task, Mark Wilson, Matthew Molineaux, and David Aha examines how to characterize the ways in which an agent can generate explanations in partially observable environments. The paper proves that it is unimportant for an agent to know the ground truth explanation,

but that narrowing down the possible explanations can guide the agent towards more reasonable actions for goals.

These papers illustrate the importance to computational systems of the explicit capacity to reason about goals as well as behavior. Such capability reflects a particular human characteristic that distinguishes us from other species in terms of abstract reasoning about what is desirable in the future and the options we have to affect it. The objective of this special issue is to highlight some of the progress made in this area and to enlarge the scope of discussion concerning the potential benefits and costs that such research implies.