Towards Quantifying the Completeness of BDI Goals

(Extended Abstract)

John Thangarajah, James Harland
Computer Science & Info Tech.
RMIT University, Australia
{johnt,james.harland}@rmit.edu.au

David N. Morley
Artificial Intelligence Center
SRI International, USA
morley@AI.SRI.COM

Neil Yorke-Smith
American University of Beirut
and University of Cambridge
nysmith@aub.edu.lb

ABSTRACT

Often, such as in the presence of conflicts, an agent must choose between multiple intentions. The level of completeness of the intentions can be a factor in this deliberation. We sketch a pragmatic but principled mechanism for quantifying the level of completeness of goals in a Belief-Desire-Intention–like agent. Our approach leverages previous work on resource and effects summarization but we go beyond by accommodating both dynamic resource summaries and goal effects, while also allowing a non-binary quantification of goal completeness.

Categories and Subject Descriptors

I.2.11 [**Artificial Intelligence**]: Distributed Artificial Intelligence—*Intelligent agents.*

Keywords

partial completeness; goal reasoning; resource summaries

1. INTRODUCTION AND BACKGROUND

In agent systems in the Belief-Desire-Intention (BDI) tradition, the most common conceptualization of goal accomplishment is discrete: a goal is either complete (usually, a plan for it has succeeded), or it is incomplete (whether execution of a plan or plans for it has begun or not) [1, 7].

For the deliberation that an intelligent agent undertakes about its goals – such as the decision about which intention to focus on next – the agent is thus limited to a coarse binary approximation of goal completeness. If the agent were able to compute a finer-grained approximation of the level of completeness of its goals, it could make more nuanced and potentially more suitable decisions. For example, when resolving goal conflicts [4], the agent may choose to continue with the goal that is more complete than the other.

While the notion of partially-complete goals has been defined in the literature [8, 7], reasoning frameworks to date have largely left unanswered *how* to compute the level of completeness of a goal in a realistic and principled manner.

The closest work to our own is that of van Riemsdijk and Yorke-Smith, who formalize the concept of a partially-complete goal for a BDI-like agent [7]. They capture partial satisfaction of a goal using a progress metric A, and a minimum value $a_{min} \in A$ that the goal must attain for the agent to consider it completely satisfied. These authors describe how an agent can reason based on such a representation, but do not provide any detailed computational mechanisms.

Our focus in this line of work is *to specify a principled and general approach that can be used computationally to quantify a measure of completeness for a given goal.* It is not our aim here to specify how an agent subsequently uses this information, i.e., its deliberation mechanisms.

There are several factors that may contribute towards assessing the completeness of a goal [3]: resources, deadlines, number of actions/plans complete, time elapsed, effects realized, etc. In this work, we propose the use of two factors to determine a quantifiable measure of completeness: *resource consumption* and the *effects* of achieving the goal.

First, we use resource consumption to provide a measure of the level of *effort* the agent has dedicated towards satisfying a goal. There has been previous work on representing resource requirements and continuously refining them as the agent executes its goals [6, 2]. We build on this existing work to provide a quantifiable measure of completeness with respect to the agent's effort.

Second, the *effects* of a goal capture its desired outcome, generally in terms of conditions that should be true when the goal execution is complete [4]. For example, the effect of a goal of a Mars rover robot to scan an area for targets of interest is that the area is scanned. We use the effects of the goal to provide a measure of the level of goal *accomplishment*, since the purpose of the goal is indeed to bring about its intended effects. As with resources, we build on and extend existing work on representing and reasoning about the effects of goals and plans [5]. In that prior work, effects are represented as boolean predicates, such as *area-scanned* in the rover example. Because there can be situations where the conditions may be satisfied to a certain degree, such as 80% of the area is scanned, we extend the prior work to allow a non-boolean representation of effects.

2. QUANTIFYING COMPLETENESS

We consider resources of two types – consumable and reusable – and treat both types to be of equal importance, since their relative importance depends on the application domain. Reusable resources are available for use again once released from their current use.

In order to determine the level of completion of a goal g at the current time t, with respect to resources or effects, it is necessary to determine (1) the resources consumed and

Appears in: *Alessio Lomuscio, Paul Scerri, Ana Bazzan, and Michael Huhns (eds.), Proceedings of the 13th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2014), May 5-9, 2014, Paris, France.*

effects attained thus far in executing g, and (2) the resources required and effects that should be attained in order for the goal to complete from t.

While the former step can be computed accurately by monitoring the resource consumption and checking the current state of the world for effects achieved, the latter step is more complex. The nature of BDI agent systems are such that there are different ways (plans) of accomplishing a particular goal, and these may use different resources and bring about different effects. Moreover, plans may fail and unexpected events may occur. The deliberation on which way to achieve the goal (i.e., plan selection) is made dynamically during execution depending on the context the agent is in, and hence is not known in advance. Consequently we cannot always say *a priori* precisely what resources will be needed to accomplish a given goal. We therefore adopt and extend the look-ahead mechanism of [4, 5] which uses summary information to compute a lower- and upper-bound of future resource usage and of effects attained.

Resources as a Measure of Completeness.

The aim of our resource analysis is to provide an agent with a quantified measure of effort with respect to the amount of resources consumed thus far in executing a goal, in the context of the total resource requirements for achieving the goal. Hence we require the agent to keep track of the total resources consumed in executing each goal.

We use the necessary and possible resource summaries to provide a lower- and upper-bound resource consumption analysis, respectively. The lower-bound analysis calculates, for every resource that has been used by the current time or is necessary in the future, the percentage of the value of that resource that has been consumed at time t. The upper-bound computation makes the same computation but with the possible resource summary. In both cases, we aggregate the percentage values to attain a single value.

Effects as a Measure of Completeness.

In contrast to resources, a measure of effort, effects are a measure of accomplishment. As stated, the effects of a goal can be thought of as the state of the world that the agent wants to achieve in order to accomplish the goal. The percentage of these effects currently achieved gives a quantifiable measure of accomplishment. It may well be the case that some effects have a greater impact on the achievement of the goal, but for a domain-independent approach we treat them of equal importance.

We propose two computational approaches: the first based on the success condition of the goal and the second on the effect summaries of the goal.

The first approach computes the level of completeness of a goal g with respect to its success condition $S(g)$ by calculating the percentage of the value of each effect in $S(g)$ currently achieved by the agent relative to the initial value when the goal execution began.

Although relatively simple, this approach ignores effects other than those in the success condition of the goal – even for those goals where some (side-)effects are necessary in order to achieve the goal's effects.

The second approach therefore includes these effects as part of the quantification of completeness. We use effect summaries to compute, respectively, a lower-bound using the definite effect summary, and an upper-bound using the combined definite and potential effect summaries (since they are exclusive). We adopt the techniques developed in [5] for deriving and updating the effect summaries, but generalize their formulae to operate on a set of effects that are composed of key-value pairs and not simple predicates.

3. SUMMARY AND FUTURE WORK

This line of work is motivated by how an agent can obtain information to make the most suitable decisions about its courses of action. We have sketched a principled mechanism for computing completeness of goals of a BDI-style agent. To our knowledge, this work is the first to study such computation with an emphasis on tractable, pragmatic reasoning. Our technical approach, not fully described here, leverages and extends earlier works on efficient resource and effect summarization. Our implementation inherits their low computational overhead. An agent can use the estimations of goal completeness in order to inform its deliberation in important reasoning areas such as goal prioritization and conflict resolution [7, 4].

Future work includes evaluating our implementation on additional scenarios, considering more closely the potentially non-monotonic nature of effects, and investigating domain-dependent weighting of resource and effect types.

Acknowledgements. We thank the reviewers for their comments. JT acknowledges the ARC Discovery grant number DP1094627. NYS thanks the Operations group at the Judge Business School and the fellowship at St Edmund's College, Cambridge.

4. REFERENCES

[1] M. Georgeff and A. Rao. Rational software agents: From theory to practice. In *Agent Technology: Foundations, Applications, and Markets*, chapter 8, pages 139–160. Springer, New York, 1998.

[2] D. Morley, K. L. Myers, and N. Yorke-Smith. Continuous refinement of agent resource estimates. In *Proc. of AAMAS'06*, pages 858–865, 2006.

[3] J. Thangarajah, J. Harland, and N. Yorke-Smith. A soft COP model for goal deliberation in a BDI agent. In *Proc. of CP'07 Workshop on Constraint Modelling and Reformulation (ModRef'07)*, pages 61–75, 2007.

[4] J. Thangarajah and L. Padgham. Computationally effective reasoning about goal interactions. *J. Automated Reasoning*, 47(1):17–56, 2011.

[5] J. Thangarajah, L. Padgham, and M. Winikoff. Detecting and avoiding interference between goals in intelligent agents. In *Proc. of IJCAI'03*, pages 721–726, 2003.

[6] J. Thangarajah, M. Winikoff, L. Padgham, and K. Fischer. Avoiding resource conflicts in intelligent agents. In *Proc. of ECAI'02*, pages 18–22, 2002.

[7] M. B. van Riemsdijk and N. Yorke-Smith. Towards reasoning with partial goal satisfaction in intelligent agents. In *Programming Multi-Agent Systems VIII*, pages 41–59. Springer, New York, 2012.

[8] Y. Zhou, L. van der Torre, and Y. Zhang. Partial goal satisfaction and goal change: Weak and strong partial implication, logical properties, complexity. In *Proc. of AAMAS'08*, pages 413–420, 2008.