

Plan-Based Intention Revision*

Adam Amos-Binks

Department of Computer Science
NC State University
Raleigh, NC, USA, 27695
aaamosbi@ncsu.edu

R. Michael Young

School of Computing and
The Entertainment Arts and Engineering Program
University of Utah
Salt Lake City, UT, USA, 84112
young@eae.utah.edu

Introduction

Plan-based story generation has operationalized concepts from the Belief-Desire-Intention (BDI) theory of mind to create goal-driven character agents with explainable behavior. Two classes of agents have emerged from this representation. A ‘conscious rationalizer’ agent who can justify *adopting* an intention with an air-tight causally-linked sequence of actions and a ‘fairweather’ agent who *drops* their intention the moment a causal complication is introduced.

While adopting and dropping intentions is essential for goal-oriented agents, restricting agents to these two classes does not capture the dynamic nature of intentions. Intention revision models explainable behavior changes, which is especially important for the complex interactions between character intentions in interactive narratives.

We define an intention revision model with two parts. We first prescribe when agents should *reconsider* existing intentions with logic from BDI agent design. Second, we define how agents decide to *revise* intentions using persistent goals (P-GOAL) as characterized by Cohen and Levesque (1990). Using the QUEST (Graesser, Lang, and Roberts 1991) cognitive model of question answering, we describe an evaluation assessing the explainability of intention revision.

Previous work

Story planners have demonstrated a range of character agent intention. IPOCL story planning (Riedl and Young 2010) operationalized intention by only generating solution plans containing actions on a causally connected action sequence to a character agent’s intentional goal. Character intentions are structured into *intention frames* and are the source of our ‘conscious rationalizer’ agent. Building on this approach was Glaive and the CPOCL representation (Ware et al. 2014). CPOCL modelled conflict between agents by allowing two agents to pursue conflicting plans. This led to our ‘fairweather’ agent, who drops their intentional goal once a

*This material is based upon work supported in whole or in part with funding from the Laboratory for Analytic Sciences (LAS). Any opinions, findings, conclusions, or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of the LAS and/or any agency or entity of the United States Government.

Copyright © 2018, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

Algorithm 1 BDI control loop excerpt (Rao and Georgeff 1998)

```
1: ...  $\mathcal{B}(\text{beliefs}), \mathcal{D}(\text{desires}), \mathcal{I}(\text{intentions}), \pi(\text{plan})$   
2: get next observation  $\omega$   
3: revise  $\mathcal{B}$  on the basis of  $\omega$   
4: if (reconsider( $\mathcal{B}, \mathcal{I}$ )) then  
5:    $\mathcal{D} = \text{options}(\mathcal{B}, \mathcal{I})$   
6:    $\mathcal{I} = \text{filter}(\mathcal{B}, \mathcal{D}, \mathcal{I})$   
7:   if not sound( $\pi, \mathcal{B}, \mathcal{I}$ ) then  
8:      $\pi = \text{plan}(\mathcal{B}, \mathcal{I})$   
9: ...
```

causal link threat is introduced to their plan. Underlying both of these models of intention is Cohen and Levesque’s (1990) intentional logic for rational agents. At the core of this logic is the definition of a P-GOAL. A goal that an agent is committed to achieving and will only drop upon achieving it or believing it is unachievable. These foundations influenced the designs of decision making for BDI agents (Rao and Georgeff 1998). Of interest to this work is how a BDI agent might decide to change intentions, summarized in Alg. 1.

Assessments of BDI agent behavior often rest in the behavior’s explainability. One extensively used explainability measure is the QUEST cognitive model of question-answering (Q-A) in the context of stories. The model uses a QUEST knowledge structure (QKS) to represent mental models of a story’s intentional structure. QUEST also defines how to structure Q-A and includes a QKS search method to predict responses. Comparing QUEST predictions and people’s responses to Q-A assesses explainability.

Plan-based Intention Revision

In our formalization, we use typical definitions for POCL planning (e.g. (Penberthy and Weld 1992)), extended to characterize intentional structures (Riedl and Young 2010).

A character agent’s goal-driven behavior is within the context of a story plan that solves a story planning problem.

Definition 1 (Story plan) A story plan π is $\langle S, B, O, L, I \rangle$ where the set of steps is S , B the set of binding constraints on the variables of S , O the partial ordering of the steps in S , L the set of causal links joining steps from S , and finally I , the set of intention frames that define character subplans.

Intention frames structure story plan elements into goal-oriented behavior of individual character agents.

Definition 2 (Intention Frame) An intention frame is a tuple $\mathcal{I} = \langle c, g, m, \sigma, T \rangle$ where c is a character agent, g is c 's intentional goal, motivating step $m \in S$ with the effect $\neg g$, the satisfying step $\sigma \in S$ with g as an effect. A subplan for c to achieve g is a set of steps $T \subseteq S$ that c consents to, each step shares at least one causal link to another step in T , and achieves goal g . Steps in T occur after m and before σ .

In the BDI control loop, agents observe their environment, update their beliefs, and reconsider intentions based on them (lines 2-4, Alg. 1). For our plan-based model, observations are effects of actions that modify character beliefs represented as sets of consistent, non-modal, ground literals.

A character will reconsider a goal when an effect introduces a causal link threat to the intention frame's subplan.

Definition 3 (Reconsidered Intention Frame) A reconsidered intention frame, \mathcal{I}_R , is $\langle \mathcal{I}, \epsilon \rangle$ where ϵ is a literal that introduces a causal link threat to two linked steps in $T(\mathcal{I})$.

After reconsidering an intention frame, an agent must deliberate if their goal is still worth pursuing (line 5-7, Alg. 1). Cohen and Levesque (1990) prescribe that an agent may only drop a goal after achieving it or when the agent believes the goal is unachievable. We are interested in the latter:

Definition 4 (Unachievable Goal) A goal is unachievable, g_u , if using a character's belief state as the initial state, no subplan to achieve $g(\mathcal{I}_R)$ exists.

If a character finds $g(\mathcal{I}_R)$ is unachievable, they will drop it and adopt an intention frame containing an achievable goal that supports solving the story planning problem:

Definition 5 (Intention Revision) An intention revision is $\langle \mathcal{I}_R, \mathcal{I}' \rangle$ where $g(\mathcal{I}_R)$ is unachievable and \mathcal{I}' is an intention frame where $g(\mathcal{I}') \neq g(\mathcal{I}_R)$, and $c(\mathcal{I}') = c(\mathcal{I}_R)$.

We refrain from discussing how to remove \mathcal{I}_R and add \mathcal{I}' to an existing story plan. However, new approaches that manage dependencies between character intention frames is promising area (Amos-Binks, Potts, and Young 2017).

Evaluation

The use of QUEST to evaluate intentional plan structure explainability is based on its extensive validation as a cognitive model of narrative text comprehension. QUEST is limited, however, because it has not characterized text with character intention revisions. We outline how QUEST instruments can evaluate the explainability of our intention revision model.

The QKS is a directed graph composed of nodes representing events, states, and goals. Arcs describe seven types of relationships that capture the causal and goal relationships between nodes. Based on our intention revision definitions and QKS composition rules, we hypothesize the QKS in Figure 1 is created from reading text with intention revisions.

QUEST's Q-A model takes any two nodes from the QKS and structures them in a Q-A pair (e.g cols 5,6 in Table 1). Subjects rate the goodness of answer (GOA) of the 'A' node to the 'Q' node using a four-point Likert scale. QUEST predicts GOA based on a QKS graph traversal. Agreement of QUEST predictions and subject responses validate the QKS.




Figure 1: QKS created by plan-based intention revision

#	Pair	Type	GOA	Question	Answer
Q1	E1,S1	Cons	Good	What is a cons. of E1?	That $c(\mathcal{I})$ believes g_u (S1)
Q2	G2,E1	Why	Good	Why did $c(\mathcal{I}')$ want G2?	Because of E1
Q3	G2,G1	Why	Bad	Why did $c(\mathcal{I}')$ want G2?	Because $c(\mathcal{I})$ wanted G1

Table 1: Questions to assess intention revision explainability

In Table 1, Q1 assesses if story event E1 consequentially leads to S1 where $c(\mathcal{I})$ believes $g(\mathcal{I}_R)$ is unachievable. In Q2, we are interested if E1 also initiated the adoption of the goal in G2. Finally, Q3 ensures that the new goal, G2, is perceived as different from G1 and not a sub-goal. Together, this group of questions validates the QKS in Figure 1 is being constructed and produces an explanation. Our next steps are to first experiment using hand-written text with intention revisions, then use text translated directly from story plans.

Summary

To extend the behavior range of character agents in story plans, we have defined plan-based intention revision with two key properties. The first is when an agent should reconsider an intention frame. We use a causal-link threat to operationalize the *reconsider* function from the BDI control loop. The second property describes how an agent should *revise* their intentions. Based on when a P-GOAL should be dropped, we allow an intention revision when no subplan exists for a reconsidered intention. Finally, we discuss evaluating these two properties using the QUEST cognitive model.

References

- Amos-Binks, A. A.; Potts, C.; and Young, R. M. 2017. Planning Graphs for Efficient Generation of Desirable Narrative Trajectories. *Intelligent Narrative Technologies Workshop*.
- Cohen, P. R., and Levesque, H. J. 1990. Intention is choice with commitment. *Artificial Intelligence* 42(2-3):213–261.
- Graesser, A. C.; Lang, K. L.; and Roberts, R. M. 1991. Question answering in the context of stories. *Journal of Experimental Psychology: General* 120(3):254–277.
- Penberthy, J., and Weld, D. 1992. UCPOP: A Sound, Complete, Partial Order Planner for ADL. *International Conference on Knowledge Rep. and Reasoning* 103–114.
- Rao, A. S., and Georgeff, M. P. 1998. Decision procedures for bdi logics. *Journal of Logic and Comput.* 8(3):293–343.
- Riedl, M. O., and Young, R. M. 2010. Narrative Planning : Balancing Plot and Character. *JAIR* 39:217–267.
- Ware, S. G.; Young, R. M.; Harrison, B.; and Roberts, D. L. 2014. A Computational Model of Plan-based Narrative Conflict at the Fabula Level. *IEEE TCIAIG* 6(3):271–288.