A Framework for Dynamical Intention in Hybrid Navigating Agents

Eric Aaron and Henny Admoni

Department of Mathematics and Computer Science
Wesleyan University
Middletown, CT 06459

Abstract. As a foundation for goal-directed behavior, the reactive and deliberative systems of a hybrid agent can share a single, unifying representation of intention. In this paper, we present a framework for incorporating *dynamical intention* into hybrid agents, based on ideas from *spreading activation* models and *belief-desire-intention* (*BDI*) models. In this framework, intentions and other cognitive elements are represented as continuously varying quantities, employed by both sub-deliberative and deliberative processes: On the reactive level, representations support some real-time responsive task re-sequencing; on the deliberative level, representations support common logical reasoning. Because cognitive representations are shared across both levels, inter-level integration is straightforward. Furthermore, dynamical intention is demonstrably consistent with philosophical observations that inform conventional BDI models, so dynamical intentions function as conventional intentions. After describing our framework, we briefly summarize simple demonstrations of our approach, suggesting that dynamical intention-guided intelligence can potentially extend benefits of reactivity without compromising advantages of deliberation in a hybrid agent.

1 Introduction

Intention-guided inference is often based on propositional, deliberation-level representations, but some goal-directed intelligence can be based on dynamical, sub-deliberative representations of intention, as well. For example, consider an animated agent completing tasks, running errands in a grid world. Along with various desires and beliefs —such as, e.g., its belief that it does not have a letter to mail— the agent starts out with intentions in its cognitive system, one for each task it might perform; each intention has an associated cognitive *activation* value, representing the intensity of commitment to the corresponding task, the task's relative *priority*. As the agent begins, its plan of action is represented by its ordering of intention activations (i.e., priorities) from highest to lowest: For this example, its intention to *deposit a check* at the bank (I_{DC}) has the highest activation, so it is first in the planned task sequence; the agent's intention to *get its child* from school (I_{GC}) has slightly lower activation and is second in the planned sequence; and the remainder of the sequence follows from the agent's

E. Corchado et al. (Eds.): HAIS 2009, LNAI 5572, pp. 18–25, 2009.

other intentions, including (for illustrative purposes) its intention to mail a letter (I_{ML}), which is inconsistent with its belief that it does not have a letter.

As the agent progresses through its sequence of tasks, its cognitive state changes due to both deliberative and sub-deliberative processes. Its cognitive activations continuously evolve, for instance, unobtrusively causing I_{ML} to become negative, consistent with the agent's belief that it has no letter. Before it reaches the bank, continuous evolution also causes the activation of I_{GC} to slightly exceed that of I_{DC}; the agent then deliberates about what its highest priority task and overall task sequence should be, using geographic and task-specific knowledge to resolve uncertainty. Thus, the same representations of intentions are employed in two different contexts: Reactive processes, not conventional deliberation, resolve the mail-related inconsistency; and deliberative inference is invoked when the agent is called upon to select a new current task from two candidate tasks of essentially equivalent, maximal priority.

In this paper, we present a framework supporting such deliberative and reactive intelligence in *hybrid dynamical cognitive agents* (*HDCAs*, for short). The design of HDCAs' cognitive systems is influenced by the *belief-desire-intention* (or *BDI*) theory of intention [1]; the theory and its related implementations (e.g., [2,3] and successors) suggest that BDI elements (beliefs, desires, and intentions) are an effective foundation for goal-directed intelligence. HDCAs' sub-deliberative cognitive models are notably influenced by some *distinguishing properties* that differentiate intention from desire (noted in [1]). For examples, an intensely committed intention I diminishes impacts of other intentions on the intensity of I; the strongest intentions (i.e., intentions with the most intense commitment) need not correspond to the strongest desires; and intentions, not desires, govern HDCAs' task priorities.

In conventional deliberative agents, BDI-based intentions are represented and manipulated with mechanisms that do not emphasize *continuous-modeled* cognition, but HDCAs' cognitive models interconnect BDI elements in a continuously evolving system inspired by (though significantly different from) *spreading activation* frameworks of [4,5]. Each BDI element in an HDCA is represented by an activation value, indicating its salience and intensity "in mind" (e.g., how intensely held a belief or committed an intention), and cognitive evolution is governed by differential equations, so elements' activation values affect rates of change of other elements' activations. HDCAs employ these cognitive representations of dynamical intention on both reactive and deliberative levels, distributing the burden of goal-directed intelligence and enabling smooth hybrid integration.

2 Hybrid and Deliberative Structure

The reactive / deliberative structure of HDCAs is illustrated in Figure 1, showing sub-deliberative, continuous-modeled cognitive processes and dynamical navigation; deliberative task sequencing and path planning processes; and cognitive representations shared across levels. Each level employs cognitive representations in its own manner, but the representations fully support both levels, for

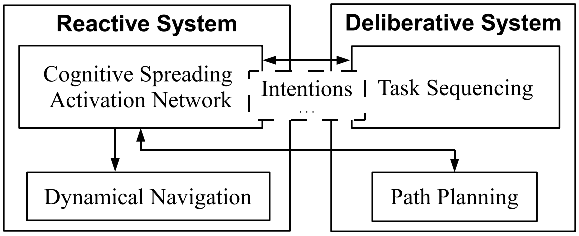

Fig. 1. System-level architecture of an HDCA, showing deliberative and reactive levels. Representations of intentions are shared by the sub-deliberative cognitive spreading activation network and the deliberative task sequencing process.

straightforward hybrid integration. The particular deliberative task sequencing and path planning processes of HDCAs in this paper are simple, although other, more complicated methods could be readily employed. Planners essentially derive "utility" values for each option, each task or path segment, based on geographic information, task-specific knowledge, and cognitive activations. Plans are then simply sequences (e.g., of tasks) in decreasing order of utility; higher commitment to an intention (i.e., higher intention activation) translates to higher utility for the associated task but is not the sole factor in determining task sequence.

Deliberation is designed to be invoked only in situations that are not well handled by fully reactive processes. In our illustrative examples, deliberation occurs only in two circumstances: if the current task is unexpectedly interrupted (e.g., by a blockaded street); or if the agent is called upon to change its current task —due to completing the previous task, evolutions of intention activations, or any other cause— and must select from multiple candidates with essentially equivalent intention activations (i.e., within a given threshold value of each other; see section 5 for an example). Unlike reactive task re-sequencing, which occurs when any intention activation values change relative ordering in their sequence and depends only on those values, HDCAs' deliberative processes also include domain-specific facts and global world knowledge. In section 5, for instance, deliberative task sequencing encodes that borrowing a book precludes buying a book, and it considers locations at which tasks are completed, making geographic distance critical to task sequencing. Deliberation also re-evaluates an agent's entire task sequence, adjusting activations of cognitive elements so that, e.g., tasks earlier in the sequence have higher activations on corresponding intentions, and precluded tasks have highly negative intentions. After deliberation, an agent simply continues with its new cognitive activation values in the reactive behavior of its new highest priority task.

In addition to being a hybrid reactive / deliberative system, an HDCA is a *hybrid dynamical system* (*HDS*, for short), a combination of continuous and discrete dynamics, modeled by a *hybrid automaton* [6,7]. A hybrid automaton is a finite state machine in which each discrete state (or *mode*) can be viewed as a continuous-modeled behavior, containing differential equations that govern system evolution in that mode. Transitions between modes (including those from

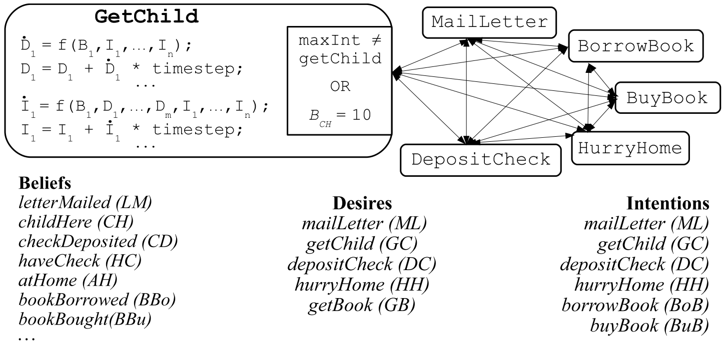

Fig. 2. Hybrid dynamical system modes and BDI elements, including abbreviations for BDI element names, for HDCAs in this paper

a mode to itself) are instantaneous, occurring when *guard* conditions are met, and may have discontinuous *side effects*, encoding discrete system dynamics. Hybrid dynamical systems are often apt models for navigating robots or animated agents (e.g., [8,9]), and HDCAs' reactive and deliberative structures naturally correspond to HDS elements: Each task of an HDCA is a reactive behavior, implemented as an HDS mode, and agents switch among these individually continuous tasks; deliberation in HDCAs only occurs during such transitions.

3 Reactive Structure

An HDCA's physical state (heading angle ϕ, position (x, y)) continuously varies as it navigates. For HDCAs in this paper, agent steering is based on [10], and simple intersection-to-intersection navigation in a grid world is similar to the method in [6], but other dynamical approaches could have been equally effective.

The continuous physical state smoothly integrates with an HDCA's cognitive system, which is based on continuously evolving activations of BDI elements (beliefs, desires, intentions); differential equations govern continuous evolutions of all elements, physical and cognitive. (Element values can also be changed discretely, as effects of mode transitions; after completing a task, for example, HDCAs' mode transitions discretely set the corresponding intention activation to the minimum possible value.) Figure 2 shows BDI elements (and abbreviations for their names) and the mode transition model for HDCAs in this paper, which is simplified to a one-to-one correspondence between intentions and actions.

In HDCAs, BDI elements are represented by *activation* values, restricted to the range $[-10, 10]$. Near-zero values indicate low salience, and greater magnitudes indicate greater salience and intensity of associated concepts, so, e.g., more active intentions represent more commitment to and urgency of the related actions. Negative values indicate salience of the opposing concept, e.g., for intentions, intention not to perform the related action. These cognitive activations are interconnected in differential equations; a partial cognitive system (with many terms and equations omitted) is in equation 1, where beliefs, desires, and intentions are represented by variables beginning with B, D, and I, and time-derivative variables are on the left in each equation:

$$\begin{aligned} \dot{D_{DC}} &= a_1 B_{HC} + a_3 I_{DC} - a_5 I_{GC} + \cdots \\ \dot{I_{DC}} &= b_1 B_{HC} + b_3 D_{DC} - b_6 D_{HH} + b_8 I_{DC} - b_{10} I_{GC} + \cdots . \end{aligned} \tag{1}$$

This illustrates interconnectedness: Elements exert *excitatory* or *inhibitory* influence by increasing or decreasing derivatives. Variables stand for activations of cognitive elements (e.g., desire to deposit a check, D_{DC}). Coefficients encode impacts of connections; most are constants, but intention coefficients contain functions that also encode distinguishing properties of intention (see section 4).

There is also a mechanism for *perception* in HDCAs, by which proximity to an item (e.g., a UPS deposit box, a street blockade) affects agents' cognitive systems. Current HDCAs have only limited perceptual structure; potential for substantial extensions exists but is not discussed in this paper.

4 Properties of Intention

In HDCAs, there are mechanisms that ensure consistency with distinguishing properties of intention (noted in [1]) that apply to our dynamical account of intention—intentions are *conduct controlling* elements that, when salient, *resist reconsideration* and *resist conflict* with other intentions.[1]

For *reconsideration resistance*, we encode two criteria: any *high-active* intention I_a (i.e., having high activation magnitude) tends to minimize impacts on I_a from other intentions; and the magnitude of this effect grows as the activation (magnitude) of I_a grows. To enable this, for intentions I_a and I_b, for every $a \neq b$, the differential equation for $\dot{I_a}$ includes the following structure:

$$\dot{I_a} = \ldots \;\; - k_n \cdot PF(I_a) \cdot I_b \;\; \ldots . \tag{2}$$

For example, in equation 1, the coefficient of I_{GC} in the equation for $\dot{I_{DC}}$ has the form $b_{10} = k_{10} \cdot PF(I_{DC})$, with *persistence factor* PF defined as

$$PF(I_a) = 1 - \frac{|I_a|}{\sum_i |I_i| + \epsilon}, \tag{3}$$

where noise term $\epsilon > 0$ prevents division by 0, and i ranges over all intentions. For $b \neq a$, $PF(I_a)$ multiplies every intention I_b in the equation for $\dot{I_a}$, so as $PF(I_a)$ nears 0 (i.e., as I_a grows in magnitude relative to other intentions), contributions of every such I_b are diminished, and when $PF(I_a) = 1$ (i.e., $I_a = 0$), such contributions are unaffected. The denominator encodes that I_a is less reconsideration-resistant when other intentions are highly active.

[1] These are not the only properties of intention that are emphasized in [1]; they are, however, properties that can apply to reactive-level intention, not requiring, e.g., future-directedness incompatible with reactive implementations.

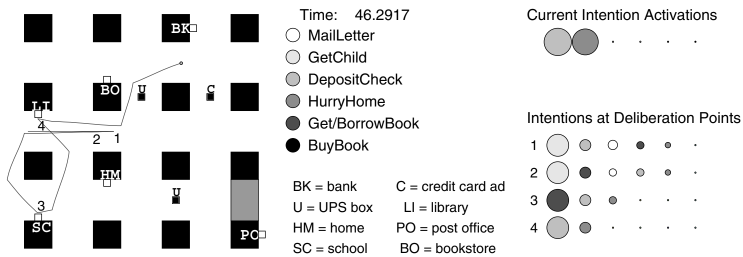

Fig. 3. Screen display of a simulation in progress. A map of the four-by-four grid world, left, shows buildings and obstacles (black boxes), targets (white squares abutting buildings), the current position of a moving blockade (gray box), and an agent's path. Numbers identify locations at which the agent invoked deliberation. On the right of the display, task sequences (including results of deliberations) are represented by sequences of color-coded circles, where radii indicate the magnitudes of activations of intentions.

Due to this and other activation-oriented mechanisms for *conduct control* and *conflict resistance* (see the supplementary website for this paper [11]), dynamical HDCA intentions are consistent with distinguishing properties of intention in [1]. Because an HDCA's deliberative level can be straightforwardly implemented to be consistent with [1], and because cross-level interconnections do not violate relevant properties, the dynamical intentions in HDCAs function as conventional intentions, rather than as some other cognitive elements inconsistent with [1].

5 Demonstrations and Experiments

Simulations of HDCAs navigating in a grid world show some features of our dynamical cognitive representations: intention-guided intelligence on both reactive and deliberative levels, extending reactive capabilities without sacrificing deliberation; and consistency of intention representations with philosophical foundations [1] of conventional BDI models. The simulations' code shows another feature, the straightforward integration of reactive and deliberative levels, united by shared cognitive representations. We briefly summarize our simulations in this section; the supplementary website for this paper [11] has animations from our simulations and more information about our code and demonstrations.

Our featured simulation of HDCA intelligence is similar to the errand-running scenario in section 1 of this paper. The HDCA navigates from intersection to intersection, eventually reaching targets in the grid world of Figure 3. That display also lists the tasks the agent might perform (*MailLetter*, ..., *BuyBook*) and indicates relative strengths of intention activations, thus showing the agent's current planned task sequence. The dynamical cognitive activations continuously evolve over time, sometimes re-sequencing tasks. The agent's simple deliberative level (see section 2) is engaged in only two circumstances: unexpected interruptions, due to an unpredictable, moving *blockade* that can block entry to any empty street; and *uncertainty* when the agent is called upon to change its current task, i.e., when there are two or more candidates for maximal priority with intention activation values differing by less than a threshold value T (see website [11] for implementation details). When deliberation is employed, the agent replans its full task sequence, incorporating task-domain and geographic knowledge.

This featured simulation contains several instances of both deliberative and reactive task sequencing. The agent in the simulation begins with maximal intention I_{DC}, but before reaching its target, activation on I_{GC} evolves to narrowly exceed that of I_{DC}. This prompts the agent to select a new current task when candidate tasks are nearly equivalent, thus invoking deliberation (at point 1 in Figure 3). The agent then generates a plan using simple methods, based on intention activations, geographic distances, and the knowledge that the agent must perform only one of *BorrowBook* and *BuyBook*. The plan is encoded by changing intention activations so that the intention corresponding to the first action in the plan gets the highest activation, etc. The agent then continues completing tasks, later deliberating when stopped by the blockade (point 2) or after completing a task when there is uncertainty about its new current task (points 3 and 4).

In addition, the agent's continuously evolving intention activations result in reactive-level task re-ordering throughout its task sequence. Most notably, as in the scenario in section 1, the agent intends to mail a letter but believes it has no letter, and it reconciles the inconsistency without deliberation: The activation of I_{ML} eventually becomes negative, consistent with its belief. Overall, the agent successfully completes tasks without superfluous activity (e.g., *BuyBook*) or action based on inconsistent intentions (I_{ML}), demonstrating the deliberative- and reactive-level intelligence enabled by its shared cognitive representations.

Other simulations directly demonstrated the consistency of HDCA intentions with distinguishing properties from [1] discussed in section 4. Some experiments simply tested HDCAs' reactive cognition, without navigation or task completion, demonstrating that terms in cognitive systems' differential equations (e.g., PF from section 4) successfully encode properties of intention. For example, when PF is present, activation values of high-active intentions are not significantly diminished by conflicting intentions, but when PF is absent, conflicting intentions can decrease high-active intention activations; similar experiments verified conflict resistance. Additionally, other experiments that did incorporate navigation demonstrated that agents' strongest desires and strongest intentions do not necessarily correspond, and intentions, not desires, serve to control HDCAs' conduct. (For more information, see this paper's supplementary website [11].)

6 Discussion and Conclusions

Hybrid dynamical cognitive agents are based on dynamical, continuous-modeled cognitive representations that extend reactive intelligence without sacrificing deliberation. For the simple implementations and explanations in this paper, HDCA deliberation is part of instantaneous HDS mode transitions, but other implementations could instead model time during deliberation—additional HDS modes could straightforwardly support this without requiring alterations to the fundamental agent model. Modeling agents as hybrid automata also supports straightforward interconnections between deliberative and reactive levels, and it potentially enables HDS analysis methods to verify some aspects of agent behavior (e.g., [12]; see [7], however, for theoretical restrictions on HDS analysis).

Because HDCAs' cognitive representations are shared by deliberative and reactive structures, intention-guided intelligence can be distributed over both levels. Moreover, HDCAs' unconventionally represented intentions are demonstrably consistent with distinguishing properties of intention noted in [1] that inform conventional BDI models, so HDCAs' intentions function as conventional intentions. Although our simple demonstrations were about a single HDCA, with limited perceptual mechanism and world interaction, they suggest more general utility of the underlying ideas: With more powerful methods for deliberative inference and perception, this dynamical intention framework could lead to intention-guided agents that rely on reactive intelligence, employing deliberation only when needed, making hybrid agents even more robust, efficient performers in dynamic multi-agent scenarios and other applications.

References

1. Bratman, M.: Intentions, Plans, and Practical Reason. Harvard University Press, Cambridge (1987)
2. Georgeff, M., Lansky, A.: Reactive reasoning and planning. In: AAAI 1987, pp. 677–682 (1987)
3. Rao, A.S., Georgeff, M.P.: Modeling rational agents within a BDI-architecture. In: Proc. of Principles of Knowledge Representation and Reasoning, pp. 473–484 (1991)
4. Collins, A.M., Loftus, E.F.: A spreading activation theory of semantic priming. Psychological Review 82, 407–428 (1975)
5. Maes, P.: The dynamics of action selection. In: IJCAI 1989, pp. 991–997 (1989)
6. Aaron, E., Ivančić, F., Metaxas, D.: Hybrid system models of navigation strategies for games and animations. In: Hybrid Systems: Computation and Control, pp. 7–20 (2002)
7. Alur, R., Henzinger, T., Lafferriere, G., Pappas, G.: Discrete abstractions of hybrid systems. Proc. of the IEEE 88(7), 971–984 (2000)
8. Axelsson, H., Egerstedt, M., Wardi, Y.: Reactive robot navigation using optimal timing control. In: American Control Conference (2005)
9. Aaron, E., Sun, H., Ivančić, F., Metaxas, D.: A hybrid dynamical systems approach to intelligent low-level navigation. In: Proceedings of Computer Animation, pp. 154–163 (2002)
10. Goldenstein, S., Karavelas, M., Metaxas, D., Guibas, L., Aaron, E., Goswami, A.: Scalable nonlinear dynamical systems for agent steering and crowd simulation. Computers And Graphics 25(6), 983–998 (2001)
11. Aaron, E., Admoni, H.: Supplementary HAIS 2009 material (2009), `http://eaaron.web.wesleyan.edu/hais09_supplement.html`
12. Asarin, E., Dang, T., Girard, A.: Hybridization methods for the analysis of nonlinear systems. Acta Informatica 43(7), 451–476 (2007)