This paper appears in the Proceedings of the Ninth International and Interdisciplinary Conference on Modeling and Using Context (CONTEXT'15), November 2-6, 2015, Larnaca, Cyprus, and in Lecture Notes in Artificial Intelligence (Springer). Copyright Springer and the authors. The final publication is available at http://dx.doi.org.

Modeling Erroneous Human Behavior: A Context-Driven Approach

Chris Wilson and Roy M. Turner

School of Computing and Information Science
University of Maine
Orono, ME 04469 USA
`chris.wilson@maine.edu, rturner@maine.edu`

Abstract. Artificial agents that model aspects of human behavior often model behaviors that an observer would regard as normal. In recent years, agents that exhibit observable erroneous behaviors have become common in a variety of applications, including simulated impaired agents to test assistive technologies and realistic agents in video games. In this paper, we present a context-driven approach to modeling plausible human behavior and a framework for modeling erroneous behavior which focuses on impairing an agent's ability to recognize and deal effectively with anticipated contextual changes.

Keywords: simulation, cognitive impairment, context-mediated behavior, assistive technologies

1 Introduction

Artificial agents that model aspects of human behavior can be found in a variety of applications ranging from scientific and military simulation to commercially available video games. In many cases, these agents are designed to model aspects of normative human behavior [4]. In recent years however, agents that model erroneous human behavior have become prevalent in both scientific simulation and commercially available video games.

One important use for simulated humans is in developing and testing cognitive orthotics, also known as assistive technologies for cognition (ATCs), meant to support patients with cognitive deficiencies (e.g., from dementia) [6]. Agents that can simulate erroneous human behavior save time and money and avoid the possible ethical issues with using actual patients. Tests of the Autominder system [7], for example, were done using an agent that would forget to perform activities on its daily agenda [8]. Similarly, Serna et al. [9] developed an agent that would perform a task's steps out of order, and in some cases, incorrectly.

It is also useful for non-player characters (NPCs) in some video games (e.g., first-person shooters, FPS) to exhibit erroneous behaviors to give the human player(s) a competitive advantage in order to avoid frustrating game play. For example, an NPC may often stand in the open for a period of time, exposing their position, before returning fire [5].

Most current simulations of erroneous behavior focus on forgetting or poor task performance. However, errors can also occur when a task is remembered and performed correctly, but out of context. For example, dementia patients often exhibit contextually-inappropriate behaviors such as wandering from the house in the middle of the night or leaving food unattended on the stove to engage in another activity.

We are developing an approach to modeling erroneous behavior that is based on errors in contextual reasoning. The approach starts with an agent that is able to behave appropriately for its current context and that is able to understand how that context will evolve as a result of pursuing a goal, then taking action to modify any problematic features of its context to be appropriate for the accomplishment of the goal. Unlike our previous work on context-mediated behavior (CMB) [11], which used knowledge of known contexts to decide how to behave while in them, the current approach uses contextual knowledge to drive behavior, including how to change the context to allow goal accomplishment. In addition to being a promising reasoning approach in its own right, this provides an elegant way to model compromised behavior by impairing the agent's ability to recognize and deal effectively with anticipated contextual changes.

In the remainder of the paper, we briefly discuss related work. We then discuss the problem and our overall approach. Next, we look at the kinds of contextual knowledge needed. We then turn to a discussion of how in our approach an agent formulates a contextually-appropriate plan. Finally, we look at how such an agent can be compromised to produce realistic, contextually-inappropriate plans.

2 Related Work

All agents use some form of contextual knowledge, implicit or explicit. In rule-based agents, this is encoded in rule antecedents describing when a rule is applicable. In planners, context is usually contained in preconditions or filter conditions that are part of the operator schema description.

A problem with contextual knowledge being local to rules or operators, however, is that unless care is taken, it can cause an agent to exhibit the same behavior regardless of the current context. For example, in the Fallout 3 video game, a village populated with friendly townsfolk is attacked by mutants, whom the player must repel. After the attack, surviving villagers still greet the player in the same friendly way, even though they are surrounded by the corpses of their neighbors [10]. While this could be remedied by adding additional constraints to when particular greetings are appropriate, this would tend to cause an explosion of such qualifiers. Worse, from the perspective of our current problem, the erroneous behavior does not necessarily reflect the kind that an impaired person would exhibit.

Some researchers have proposed the use of smart objects and situations to make an agent's behavior more context-appropriate [10]. For example, a smart object might, depending on the current context, inform an agent how to hold or gaze at it. These approaches do not, however, aid an agent in planning to achieve a specific goal. (In fact, since the agent can't know a priori what the object will tell it, this may actually hinder an agent's ability to plan.) For our current focus, one could imagine altering objects' actions to cause contextual errors; however, this would move the focus from the agent being modeled to the environment and could result in all agents present behaving inappropriately for the context.

In earlier work, we argued for the benefits of explicitly representing contexts and contextual knowledge with respect to acquiring, learning, reasoning about, and using such knowledge [11]. In our context-mediated behavior (CMB) approach, known contexts are represented as contextual schemas (c-schemas), which both describe the contexts as well as prescribe how to behave while in them. This can address some of the limitations of related approaches, and it provides a way to inject errors into an agent's behavior by impairing its ability to reason about its context.

In the current project, we develop this idea to create a believable impaired agent. However, we are not only interested in using contextual knowledge to mediate aspects of an agent's current behavior, but also aspects of its future behavior. To this end, we treat an agent's context as a dynamic object that continuously evolves as an agent works to achieve its goals. Similar to other approaches, we view an agent's context as a collection of smaller contextual objects that evolve at different rates.

A related approach is taken by Brézillon and colleagues [3]. They propose three types of contextual knowledge, namely *contextualized, contextual* and *external* knowledge, that can be applied to a problem-solving step. Contextualized knowledge describes any knowledge that is used by an agent during a problem-solving step, whereas contextual knowledge is any knowledge that is not explicitly used during a step, but that constrains it. External knowledge is all other knowledge that has nothing to do with the problem-solving step.

This approach allows an agent's contextual knowledge to evolve during problem solving. For example, while pursuing a goal, a piece of contextualized knowledge might become either contextual or external knowledge. In our approach, the use of c-schemas allows us to explicitly represent both contextualized and contextual knowledge and also provides, in the current work, a way to reason about future contexts in order to help create plans.

3 Overview of our Approach

Similar to the approaches of Pollack [8] and Serna [9], we will model erroneous human behavior by implementing an artificial agent that exhibits plausible aspects of normative human behavior as it works to achieve its goals, then compromise aspects of this agent's cognitive function in order to induce the effects of a cognitive impairment.

Where our approach differs from others is in how we view, and therefore model, normative and erroneous behavior. In our approach, we regard normative agent behavior as exhibiting contextually-appropriate behavior while working to achieve its goals. We will model this type of behavior with a context-aware agent. Erroneous behavior will be achieved by compromising our agent's ability to perform basic contextual reasoning when formulating a plan for achieving a goal.

4 Required Contextual Knowledge

Our model of normative reasoning assumes that an agent uses contextual knowledge to mediate its planning process to ensure that it commits to contextually-appropriate goals. It is critical to represent context in the normative model so that it also supports impairment to give rise to erroneous behavior.

We use the term *context* to mean any identifiable configuration of environmental, goal-related, and agent-related features that has predictive power for an agent's behavior. Some features of an agent's `current context` exist as a result of the goal(s) currently being pursued. These features, which we refer to as the current *problem-solving context* (PSC), are ephemeral and are usually removed from the current context by the actions which achieved the goal (e.g., the grill cheese sandwich being prepared is not yet cooked) or by "cleanup" actions that either part of the plan or specified by the context representation (e.g., remove pan from stove, turn off stove).

Others features, which we refer to as the agent's *persistent context*, are longer-lived and persist across successive context changes. For example, once a person is dressed, being dressed will persist across successive context changes (driving to work, being at work, etc.). In our approach, we consider the agent's PSC and PC as sub-contexts of the agent's current context.

In addition to its current context, we assume that an agent has general knowledge about the context that results from pursuing a goal. For example, when considering going out to eat, most people know they will be out in public, there will be other patrons and wait staff present, and they will be expected to be appropriately dressed. This general contextual knowledge, which we refer to as an *implied context*, influences how the person plans to achieve his or her goals by ensuring that plans result in context-appropriate behaviors. In our approach, implied contexts are used during the agent's planning process.

Features of the agent's persistent context can be used to impose constraints on its future contexts, that is, on the future contexts it can be in. For example, when preparing a hot meal, the agent should focus its attention on the food in the pan to prevent it from burning or causing a kitchen fire. As long as the stove remains on and the pan remains on the stove, the agent should not enter any context in which its position is not the same as the current context (i.e., in the kitchen). Contextual constraints help the agent determine if committing to a goal will cause contextually-inappropriate behavior. In the cooking example, the agent knows that it should not commit to any goal (e.g., checking the mail) that causes it to leave the kitchen. We use contextual constraints to help the agent formulate contextually-appropriate plans.

5 Modeling Normative and Erroneous Behavior

In our approach, an agent avoids contextually-inappropriate goals and actions by using contextual knowledge to formulate what we will refer to as a *contextually appropriate plan* (CAP). A CAP is any plan that ensures the context induced as a result of pursuing the goal in consideration does not violate any contextual constraints that may be imposed by persistent features of the agent's current context.

A CAP is constructed by first considering the evolution of the agent's current context as a result of executing default procedural knowledge pertaining to a goal. This process allows the agent to identify persistent contextual features which are then merged into the implied context associated with the goal in consideration. The resulting context is a representation of the agent's future context surrounding the pursuit of the goal. Features of this context are then compared against any contextual constraints that may be imposed. This not only prevents the agent from committing to a goal that violates these constraints, but allows it to identify features of the current context that affect the appropriateness of the implied context. The latter information can be used to modify the plan for achieving a goal in a way that remedies any problematic features.

Our method of impairment induces contextually-inappropriate behavior by impairing the agent's ability to recall contextual knowledge when formulating a plan for achieving a goal. Some of the agent's contextual knowledge at any given time is in its working memory (i.e., the current context), while the rest is in its long-term memory (i.e., its c-schemas). To this end, our model of cognitive impairment borrows ideas from ACT theory [1] about how memories are stored and recalled.

Each feature in the agent's current context in working memory is assigned a strength S which is represented using the ACT equation for "memory trace" strength [2]. By increasing the decay rate of S for a contextual feature or by preventing it being committed to memory we can cause an agent to formulate a plan based on inaccurate contextual information, thus increasing the likelihood of committing to a goal in a contextually-inappropriate manner.

We can compare c-schemas in long-term memory to *chunks* in ACT-R's declarative module [1]. Borrowing from ACT-R, a c-schema c is given an *activation weight* A_c, which is represented using the ACT equation for chunk activation. We can impair the ability of an agent to retrieve appropriate c-schemas by manipulating the parameters of this equation to reduce A_c. Doing so will result in incomplete or wrong contextual knowledge being returned from long-term memory, which increases the likelihood of committing to achieve a goal in a contextually-inappropriate manner.

6 Conclusions and Future Work

Agents that simulate human behavior play an ever increasing role in a variety of scientific and commercial applications. For some applications, modeling plausible impaired human behavior is as important as modeling normative human behavior. In this paper, we presented a context-mediated approach to modeling plausible human behavior and a framework for impairing the cognitive function of a context-aware agent in order to simulate plausible erroneous behavior.

Currently, our work is in the very early stages. At the time of this writing, we have implemented and conducted basic preliminary tests of our context-aware agent and we are in the process of implementing our impairment framework. In the near future, we will analyze the plausibility of the resulting erroneous behavior. We will also examine how our approach can be used to model other types of erroneous behaviors beyond those associated with a cognitive impairment like dementia.

References

1. Anderson, J.R.: The Architecture of Cognition. Harvard University Press, Cambridge, MA, USA (1983)
2. Anderson, J.R.: A spreading activation theory of memory. Journal of Verbal Learning and Verbal Behavior 22(3), 261–295 (1983)
3. Brézillon, P., Pomerol, J., Saker, I.: Contextual and contextualized knowledge: An application in subway control. International Journal of Human-Computer Studies 48(3), 357–373 (1998)
4. Kormányos, B., Pataki, B.: Multilevel simulation of daily activities: Why and how? In: IEEE International Conference on Computational Intelligence and Virtual Environments for Measurement Systems and Applications (CIVEMSA). pp. 1–6. IEEE (2013)
5. Lidén, L.: Artificial stupidity: The art of intentional mistakes. AI Game Programming Wisdom 2, 41–48 (2003)
6. O'Brien, A., Ruairi, R.M.: Survey of Assistive Technology Devices and Applications for Aging in Place. Proceedings of the Second International Conference on Advances in Human-Oriented and Personalized Mechanisms, Technologies, and Services pp. 7–12 (2009)
7. Pollack, M.E., Brown, L., Colbry, D., McCarthy, C.E., Orosz, C., Peintner, B., Ramakrishnan, S., Tsamardinos, I.: Autominder: An intelligent cognitive orthotic system for people with memory impairment. Robotics and Autonomous Systems 44(3), 273–282 (2003)
8. Rudary, M., Singh, S., Pollack, M.E.: Adaptive cognitive orthotics: combining reinforcement learning and constraint-based temporal reasoning. In: Proceedings of the Twenty-First International Conference on Machine Learning. p. 91. ACM (2004)
9. Serna, A., Pigot, H., Rialle, V.: Modeling the progression of alzheimers disease for cognitive assistance in smart homes. User Modeling and User-Adapted Interaction 17(4), 415–438 (2007)
10. Sloan, C., Kelleher, J.D., Namee, B.M.: Feeling the ambiance: Using smart ambiance to increase contextual awareness in game agents. In: Proceedings of the 6th International Conference on Foundations of Digital Games. pp. 298–300. ACM (2011)
11. Turner, R.M.: Adaptive Reasoning For Real-World Problems: A Schema-Based Approach. Psychology Press (1994)