Proactivity in an Intentionally Helpful Personal Assistive Agent

Karen Myers and **Neil Yorke-Smith**
Artificial Intelligence Center, SRI International, Menlo Park, CA 94025, USA
{myers,nysmith}@ai.sri.com

Abstract

The increased scope and complexity of tasks that people perform as part of their routine work has led to growing interest in the development of intelligent personal assistive agents that can aid a human in managing and performing tasks. One desired capability for such agents is that they be able to act *proactively* to anticipate user needs, opportunities and problems, and then act on their own initiative to address them. This position paper outlines some initial thoughts on desired forms of proactive behavior, and identifies technical challenges in developing systems that embody such behaviors.

Introduction

We are interested in developing intelligent personal assistive agents that can aid a human in managing and performing complex tasks. Our overall goal is to reduce the amount of effort required by the human to complete the tasks she intends. Effort here encompasses both the activities necessary to perform the tasks, and the cognitive load in managing and monitoring them. Thus, a personalized assistive agent may aid its user *directly* by performing tasks on her behalf or in conjunction with her [7], and *indirectly* through actions such as providing context for her work, minimizing interruptions, and offering suggestions and reminders [3].

We are exploring these ideas within a system for intelligent personalized assistance called CALO [11]. The focus for CALO is to support a busy knowledge worker in dealing with the twin problems of information and task overload. CALO's current task-related capabilities are grounded in a *delegative BDI model* [12], in which the system adopts intentions only in response to being explicitly assigned them by the user. CALO can perform a variety of routine office tasks delegated by the user, such as arranging meetings and completing online forms, as well as more open-ended processes such as purchasing equipment or office supplies and arranging conference travel.

One limitation within the current CALO framework is the lack of a *proactive* capability that would enable CALO to anticipate needs, opportunities, and problems, and then act on its own initiative to address them. We are interested in developing proactive behaviors along these lines within CALO, to increase the overall effectiveness of the system as a personal assistant.

Our research objectives in this area are as follows. First, we want to understand the types of proactive behavior that would be helpful to incorporate into an assistive agent. Second, we want to characterize how an agent can best reflect over possible actions and current commitments (both user and system), as a guide to which intentions to adopt and how to pursue them. Finally, we would like to develop a theory of proactivity that characterizes both when an agent should take initiative to assist the user, and the nature of the assistance that should be given. In this position paper, we set out initial thoughts on these topics and identify some important questions for future research.

Characterizing Helpful Assistance

Ethnographic studies of human work habits and task management (e.g., [1, 3]) reveal that people usually achieve all their important tasks. We become adept at multi-tasking and remembering the things that really matter; however, we fail to achieve perfectly tasks with soft deadlines or forgettable details.

Let us distinguish tasks performed solely by the user (*user tasks*) from those performed solely by the agent (*agent tasks*), and those performed jointly in partnership (*shared tasks*). We assume that the user enjoys an instrumented work environment and that she employs electronic artifacts to keep track of her tasks, commitments, and calendar.

It is not our goal to address the general problem of inferring a user's intent [9, 5]. Although research on CALO encompasses recognizing from her actions what task a user is working on [11], our starting point (for the moment) is that the user has entered a description of her user tasks and tasks assigned to her CALO into an electronic todo list. In addition, we assume that CALO is told or can infer a mapping from these entries to formal models within a task ontology (i.e., the tasks have associated semantic descriptions).

Within this setting, we can envision a personal assistive agent aiding its user in many ways:

- achieve a goal/perform a delegated task
- collect information
- share information (with user or team)
- filter information/reduce interruptions
- remind and notify

- summarize across projects/time
- provide context for a task or message
- monitor task progress
- anticipate potential problems (e.g., lack of resources)
- explain status, reasons, and failures
- intervene in a user or shared task to help
- suggest relevant actions (e.g., link to resources for current task)
- provide team liaison (e.g., manage delegation requests)
- instruct the user (e.g., how best to use CALO)
- learn (e.g., set an agent learning goal of how to do a task)

Example Scenario CALO observes the items in your electronic todo list, what you are working on currently, what you have delegated to your CALO and to other people, and your commitments for the week ahead. CALO assesses that your workload is likely to be uncomfortably high at the end of the week. Via a chat message, CALO offers you a reminder of an important meeting early next week, with the suggestion that a paper review (on your todo list) could be transferred to a colleague (whom CALO identifies as having appropriate expertise and time in his schedule), to leave you time to focus on the meeting. In addition, CALO begins to prepare background material for the meeting without being explicitly asked. It attaches the relevant documents to the item in your todo list and the event in your calendar.

The above scenario illustrates two distinct types of proactive behavior for an agent. The first type, which we call *task-focused proactivity*, involves providing assistance for a task that the user either is already performing or is committed to performing; assistance takes the form of adopting or enabling some associated subtasks. Task-focused proactivity is exemplified in the above scenario by CALO collecting background information in support of a scheduled meeting.

The second type, which we call *utility-focused proactivity*, involves assistance related to helping the user generally, rather than contributing directly to a specific current task. An example of this type occurs in the scenario when CALO takes the initiative to recommend transferring a paper review task in response to the detection of high workload levels. This action is triggered not by a desire to assist with any individual task on the user's todo list, but rather in response to a higher-level desire (namely, workload balancing).

Principles We identify five principles to guide proactive behavior (compare the principles for intelligent mixed-initiative user interfaces in [8]):

- **unobtrusive**: not interfering with the user's own activities or attention, without warrant
- **valuable**: pertinent to advance the user's interests
- **capable**: within the scope of the agent's abilities
- **safe**: without negative consequences
- **user control**: exposed to the scrutiny and according to the mandate of the user

In the above example, CALO's actions are pertinent to the important upcoming meeting. CALO itself is not capable of reviewing the paper; identifying a colleague who potentially is able, CALO does not delegate the task from your todo list automatically, but leaves you in control to take the suggestion or not. This suggestion and the preparation of background materials are both *safe*, in that they result in no changes of state, other than a gain in information. Throughout, CALO's actions are unobtrusive: the communication is via a chat message with context, and the completed information gathering is again in context, attached to the relevant artifacts in your working environment.

Demands on a Theory of Proactivity

A theory of proactivity will likely have much in common with theories of collaboration, since both are rooted in the notion of an agent taking action to assist another. The leading candidates in collaboration frameworks are Joint Intention theory [2] and SharedPlans theory [7].

Joint Intention theory [2] formalizes the communication acts between agents to establish and maintain joint belief and intention: the obligations on what "message" to communicate and under what circumstances to do so. SharedPlans theory [7] specifies the collaborative refinement of a partial plan by multiple agents; it handles multiple levels of action decomposition and partial knowledge of belief and intention.

The characterization of *how* to provide proactive assistance could likely be modeled as a variant on these collaboration theories in which certain of the requirements for mutual belief and commitment are relaxed. Characterization of *when* to act proactively, however, is not considered within these theories. Here, we consider some of the factors that bear on this control issue.

A helpful assistive agent weighs the cost and benefits of potential intentions and the plans to achieve them [5]. We first need a *theory of action consequences* in order to define the concept of a *safe* action. For example, is a safe action one that maintains world state, other than adding to user or agent beliefs? Or is an action safe provided the state changes can be reversed (and at what cost)? Does safety also encompass not interfering with the user's actions? This theory is required for the agent to assess what actions are safe.

Second, we need a *theory of user desires* to describe what are the long- and short-term goals of the user. Such a theory provides a means of assessing the value of each agent action in terms of the user's objectives. The question for the agent is then: when are (unsafe) actions to be considered? If a task has many safe actions and high perceived benefit, should it be barred because one action is potentially unsafe, such as accepting a meeting request on the user's behalf?

Third, a theoretical basis to support the helpful behavior identified above must account for at least (1) user, shared, and agent tasks; (2) acting in support of another agent's goals; (3) restricting actions to those that are perceived safe. Finally, it must admit the timeliness of action and interaction, in order to support the agent's unobtrusive, pertinent, user-controllable mixed-initiative assistance.

Challenges in Ongoing Work

Acquiring Agent Understandable Tasks The scenario above hinges on the assistive agent's ability to infer associations among and reason over information such as todo items, calendar entries, projects, resources, plans, current task and location, and agent capabilities.

How can these various aspects of knowledge be populated? For example, semantic information about user and shared goals is critical. Three possible sources of information are (1) inference from user actions (i.e., intent and plan recognition) [9]; (2) explicit user declaration in a semantically grounded manner (as we assumed above); (3) inference from user non-semantic statement, possibly confirmed with an explicit disambiguation request. In the third case, for example, the agent might link the informally specified entry "book conference" on the user's todo list to the known task `conference-travel` by drawing on techniques such as those described in [6].

Since studies show that people decompose their work into projects and todos at differing abstract levels [3], a related challenge is the levels of abstraction at which to define a task ontology, and how to relate user-specified tasks into it.

While task-focused proactivity seeks to provide assistance to the user with immediate, tangible goals, in contrast utility-focused proactivity addresses more general objectives of the user. Utility-focused proactivity requires a representation of user desires that captures the user's unstated *interest goals* (in the terminology of the OCC cognitive model [13]) as well as explicitly stated *achieve* and *replenishment* goals. Since interest goals differ between people, a helpful assistive agent requires such a model of its user, in order to assess the value of agent actions.

Acquiring Task Parameters To act on an intention to perform a task, the agent must have instantiations for the task's input parameters. For example, it is not enough to identify a `conference-travel` task: the agent must know to which conference the user is referring. One approach is to acquire the parameters to instantiate the task fully, by asking the user to specify them. Since this risks disturbing the user with a perceived irrelevant request (unless the user has just asked her agent to perform the task, the request comes out of the blue), another approach is to guess the parameter values from learned history, or to perform information-gathering actions (e.g., look at the user's calendar).

A second approach is to act on a set of possibly matching, partially-instantiated tasks by performing safe, conditional actions, one for each possible task instantiation. For example, gather flight and hotel quotes for each conference. A strong notion of safety is needed here, since the agent must not reserve a flight to each venue! In conjunction or instead, the agent can perform conformant actions that support any of the possible tasks. For example, whatever the destination, the user will need to submit a travel authorization form, and many of the fields can be prepopulated.

Personalization and User Interaction By adapting to its user's preferred working and communication styles, and her capabilities and experience [4], an agent becomes a more helpful and thus trustworthy assistant over time. A part of personalization that is central to user experience is that of interaction [8]. When and by what modality does the agent communicate its beliefs and actions? When is it better not to act, to interrupt and ask, and to act [10]? What interfaces provide for efficient communication and collaboration? An effective assistive agent will deliberate not only about intentions to act, but also intentions to communicate.

Acknowledgments. The authors thank other members of the PExA project. This material is based upon work supported by the Defense Advanced Research Projects Agency (DARPA), through the Department of the Interior, NBC, Acquisition Services Division, under Contract No. NBCHD030010.

References

[1] V. Bellotti, B. Dalal, N. Good, P. Flynn, D. G. Bobrow, and N. Ducheneaut. What a To-Do: studies of task management towards the design of a personal task list manager. In *Proc. of CHI'04*, pages 735–742, Vienna, Austria, 2004.

[2] P. R. Cohen and H. J. Levesque. Intention is choice with commitment. *Artificial Intelligence*, 42:213–261, 1990.

[3] M. Czerwinski, E. Horvitz, and S. Wilhite. A diary study of task switching and interruptions. In *Proc. of CHI'04*, pages 175–182, Vienna, Austria, 2004.

[4] J. L. Edwards and G. Scott. LOCATE intelligent systems demonstration: Adapting help to the cognitive styles of users. In *Proc. of AAAI-06*, pages 1937–1938, Boston, MA, 2006.

[5] A. Fern, S. Natarajan, K. Judah, and P. Tadepalli. A decision-theoretic model of assistance. In *Proc. of AAAI'06 Workshop on Modeling Others from Observations*, Boston, MA, 2006.

[6] Y. Gil and T. Chklovski. Enhancing interaction with to-do lists: Using artificial assistants. In *Proc. of AAAI 2007 Spring Symposium on Interaction Challenges for Intelligent Assistants*, Stanford, CA, 2007.

[7] B. Grosz and S. Kraus. Collaborative plans for complex group action. *Artificial Intelligence*, 86(2):269–357, 1996.

[8] E. Horvitz. Principles of mixed-initiative user interfaces. In *Proc. of CHI-99*, pages 159–166, Pittsburgh, PA, 1999.

[9] E. Horvitz, J. Breese, D. Heckerman, D. Hovel, and K. Rommelse. The Lumiere project: Bayesian user modeling for inferring the goals and needs of software users. In *Proc. of UAI'98*, pages 256–266, Madison, WI, 1998.

[10] E. Horvitz, C. M. Kadie, T. Paek, and D. Hovel. Models of attention in computing and communications: From principles to applications. *Communications of the ACM*, 46(3):52–59, 2003.

[11] K. Myers, P. Berry, J. Blythe, K. Conley, M. Gervasio, D. McGuinness, D. Morley, A. Pfeffer, M. Pollack, and M. Tambe. An intelligent personal assistant for task and time management. *AI Magazine*, 28, 2007. To appear.

[12] K. L. Myers and N. Yorke-Smith. A cognitive framework for delegation to an assistive user agent. In *Proc. of AAAI 2005 Fall Symposium on Mixed-Initiative Problem-Solving Assistants*, pages 94–99, Arlington, VA, 2005.

[13] A. Ortony, G. L. Clore, and A. Collins. *The Cognitive Structure of Emotions*. Cambridge University Press, 1988.