Mixed-Initiative Negotiation: Facilitating Useful Interaction Between Agent/Owner Pairs

Pauline M. Berry, Thierry Donneau-Golencer, Melinda Gervasio, Bart Peintner, and Neil Yorke-Smith

Artificial Intelligence Center, SRI International, Menlo Park, CA 94025 USA
{berry,donneau,gervasio,peintner,nysmith}@ai.sri.com

Abstract. A mixed-initiative agent for personal time management interacts not only with its human owner but also with other agents and humans that share or depend on the same time commitments. The assistive capabilities of such an agent include the ability to provide information and context for its owner, negotiate on behalf of its owner, and understand when autonomous action is possible, preferred, or expedient. It must operate without losing the trust of its owner or negotiating partners. Since time management is intensely personal, each such human-agent pair will evolve its own characteristics and working practices. This position paper considers how an existing mixed-initiative adaptive time management agent, PTIME, can be extended to the multiagent negotiation setting. We discuss opportunities for facilitating personalized mixed-initiative negotiation protocols and adjustable autonomy through demonstration, instruction, and advice. In addition, we explore user interaction considerations including the provision of explanation to build trust by enabling the owner to understand and correct agent decisions and suggestions.

Key words: mixed-initiative, time management, negotiation, explanation, suggestion, trust

1 Introduction

Scheduling meetings in an office environment commonly depends on email-based negotiation. Typically, a series of messages between a meeting organizer and potential participants broadcast the intent to have a meeting, possible time windows for the meeting, restrictions on the windows, counter-proposals, and eventually agreement on a time, duration, location, and participants. While this method can produce the desired result, it comes at great cost to potential participants: in addition to the time spent on the viewing their calendars and responding to email, productivity is lost through the context switch required for each round of the negotiations. These costs motivate the need for automated assistance in scheduling.

A number of fully or semi-automated scheduling systems have been developed [1–3] but they have largely aimed to bypass the human negotiation aspects of meeting scheduling and to aim for full automation. The systems therefore do not fit within the organization's socially accepted (though commonly non-formalized) workflows regarding meeting scheduling, and so suffer from low adoption rates.

We have seen this lack of adoption with our own semi-automated scheduling agent, PTIME [4]. Each user has his own PTIME agent that assists in scheduling his own calendar and coordinating his calendar with others. The PTIME agent adapts to its user becoming personalized to his needs and preferences over time. Our aim is for the agent to work within users' typical scheduling workflows, rather than trying to convince users to switch to a different paradigm of meeting scheduling. In particular, we aim to automate aspects of an organization's current scheduling workflows and predict users' responses to decisions or requests for information in order to make suggestions to them or anticipate their responses. Eventually, the user may gain enough trust in his agent to allow it to make some decisions on his behalf. The larger goal is to reduce the time and number of context switches needed to organize or become an attendee in a meeting.

This time management domain leads us to address several interesting issues at the intersection of adjustable autonomy and multiagent negotiation. The combination of these two problems is less well studied in the literature, and we believe that addressing them together may produce synergistic opportunities and more intuitive agent behavior. This position paper discusses our initial thoughts on how a user and software agent can interact in the context of meeting scheduling and multiagent negotiation, and emphasizes the key contribution of user-agent dialogue to adjustable autonomy.

2 Mixed Initiative Negotiation and Personal Time Management

Artificial intelligence techniques can solve the multi-participant, distributed or centralized meeting scheduling problem, supposing sufficient information, However, existing fully or semi-automated scheduling systems fail to address the personal nature of the domain [1, 2]. The process of negotiating meeting times, the tools employed, and the preferences over events all exhibit considerable variation between individuals. For example, the best solution for an over-constrained meeting request (i.e., where not all criteria can be fulfilled simultaneously) depends on the individual: one person prefers to reschedule an existing meeting, another prefers times outside the specified window, while still another prefers to omit a participant. Hence, for a system that goes beyond standard calendaring functionality and addresses the scheduling problem to be of value, it must facilitate negotiation and embrace personalization.

Mixed-initiative systems integrate human and automated reasoning to take advantage of their complementary reasoning styles and computational strengths. Taking a mixed-initiative methodology is a powerful paradigm for engaging humans in a software process. Previous work in this area explores mixed-initiative interactions between the agent and the user in terms of a balance between agent autonomy and user control [5]. Strategies are developed for allowing the user to personalize the default behavior of an agent, to subdivide tasks between the user and the agent according to the criticality of decisions, and to allow the user to inspect the agent's behavior and correct it. More recent work has argued that mixed-initiative systems must exhibit both *adaptable* strategies to allow the user to personalize the agent behavior and *adaptive* strategies that use AI techniques to personalize agent behavior automatically [6].

This development leads to an ongoing dialogue between the agent and its owner. The agent, by using this dialogue to refine its knowledge, can over time become a progressively more capable and trustworthy, scheduling agent [7]. Eventually this may lead to an agent with a self-adaptive ability. Myers and Yorke-Smith [8] discusses theories of proactivity in single agent behaviors. Work on Electric Elves [9] explores the problem of user-agent pairs in an organizational setting. Human agents were assisted in a variety of organizational tasks including meeting planning. Scerri et al. [10] explores practical progress toward adjustable autonomy in multiagent environments composed of human/agent pairs.

2.1 The PTIME Agent

PTIME (*Personalized Time Management*) is an intelligent, personalized calendar management agent that helps users handle email meeting requests, reserve venues, and schedule events. The agent is designed to unobtrusively learn scheduling preferences, adapting to its user over time. The ability to learn is a key aspect in the evolving interaction between human and agent. Details of the PTIME scheduling process and preference learners can be found in [4].

The PTIME multiagent environment consists of agent/human pairs interacting collaboratively to solve complex event scheduling problems. Figure 1 represents a set of PTIME/owner pairs within the context of a meeting negotiation initiated by human owner A and assisted by his PTIME agent Ea. Participants (B,C,D) and their PTIME agents (Eb, Ec, Ed) can share availability information and solution preferences, and can rank options. Note that not all participants need have a PTIME agent for the negotiation to progress. At present, the PTIME user organizing the meeting decides which meeting option to select, taking into consideration other participants' generic scheduling preferences. The selected meeting option is presented to invitees for inclusion or otherwise in their calendars. In the simplest form of negotiation supported, the other participants besides the organizer may simply accept the meeting request or not. PTIME is being developed to support more elaborate forms of negotiation while keeping the user in the loop. The question remains of why, when, and how such an agent should interact with the user and how that interaction should evolve over time.

2.2 An Example Multiagent Negotiation Use Case

In the context of our time management domain, there are two main types of interaction within a negotiation.

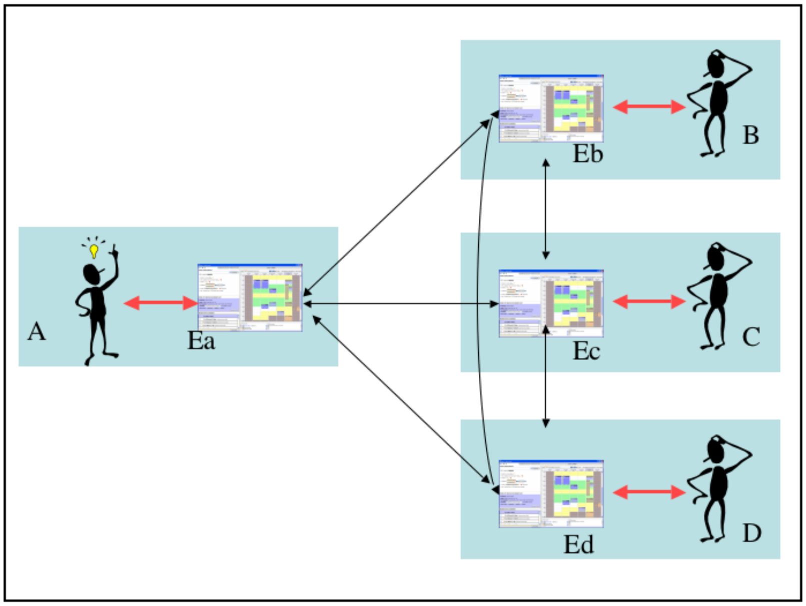

Fig. 1. Mixed Initiative Negotiation Use Case.

- A request for *information*: together, one agent/human pair can compose the request and the recieving agent/human pairs can answer. In each case the agent/human pair must evaluate the request and understand the context. The receiving agent/human pair can assess alternative responses and either answer or not. This form of request takes part during the negotiation process. An example use case is described in Figure 2.
- A request for a *decision*: in the case of meeting scheduling this typically takes the form of a request to accept or decline a suggested meeting time. Again, the agent/human pairs involved must evaluate the contents of the request, understand its context, assess alternatives, and respond. In this case the response may conclude the negotiation process or prolong it. A typical use case is described in Figure 3.

In a fully autonomous system the agent itself would manage these responses, but in a mixed initiative system it may require interaction between each agent and its human owner.

3 Facilitating Mixed-Initiative Negotiation in Agent Systems

We have argued that time management is a personal and sensitive aspect of people's day-to-day lives, and that an agent or assistant for this activity must be

1. ***Organizer A(Ea): Asks several participants to Rate meeting options s1, s2, s3***

   What are the *Interaction/Explanation/Suggest possibilities for receiving agents Eb, Ec, Ed and their owners?*

   a) The agent could pass the request directly to the user to rate options
   b) The agent could rate (some) options and present to the user to confirm or adjust
   c) The agent could additionally generate some alternative suggestions for its owner
   d) The agent could provide explanation on the ratings and on its confidence about them
   e) The agent could provide explanation on why it could not rate some options
   f) The agent could ask questions to lift the uncertainties that prevented it from rating these options
   g) The agent could rate (some) options on its own and return answers autonomously

2. ***Participant B(Eb): Replies with good rating for s1 but poor ratings on s1, s2***
3. ***Participant C(Ec): Replies with poor ratings on s1, s2, and s3 and adds suggestion s4***
4. ***Participant D (Ed): Fails to reply after some time period***

   How does the Organizer/Emma A(Ea) pair deal with varied responses from participants?

   a) Agent Ea could simply present replies and let owner pick
   b) Agent Ea could display each option with combined rating and let owner pick
   c) Agent Ea could evaluate all options (s1, s2, s3, s4), suggest highest ranking and let its owner, A, pick
   d) In each case Agent Ea could autonomously select the option on its owner's behalf

   How does the Organizer/Emma A(Ea) pair deal with no response from D(Ed)?

   a) Agent Ea could decide to wait X more minutes
   b) Agent Ea could send reminder to Ed
   c) Agent Ea could ask its owner
   d) Agent Ea could schedule without waiting for Ed's response (based in Ed's computed importance in meeting)
   e) As d) Agent Ea could act autonomously and provide its owner with a useful explanation on its handling of the situation

Fig. 2. Information-based Mixed Initiative Interaction in Negotiation Use Case.

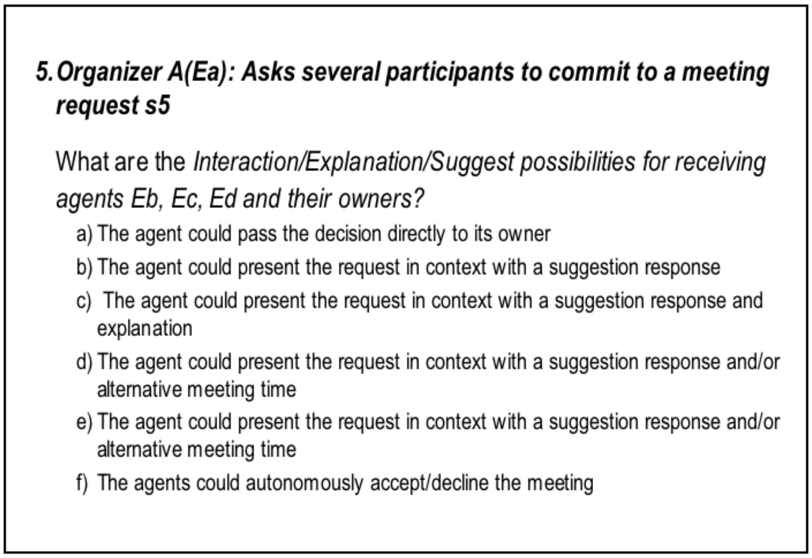

Fig. 3. Decision-based Mixed Initiative Interaction in Negotiation Use Case.

adaptable to its owner's needs. We have further argued that the agent must engage its owner in this process to both build trust and increase its ability to act for its owner over time. The thrust of our discussion is to explore useful interaction between the agent and its owner during the process of multiagent negotiation. We will discuss how the owner might inform the agent about acceptable behavior that is either desired by or helpful to the owner, and how in return the agent can assist its owner in understanding the context of the negotiation and possible actions available to, or autonomous actions taken by, the agent. The context is meeting scheduling and the interaction takes place when an agent/owner pair needs to share some information to satisfy the progression of negotiation, or the agent/owner pair must make a decision regarding the scheduling of a meeting. The mechanisms for interaction we will explore include observation, advice, suggestion, action, and explanation. These concepts are familiar to the field of intelligent user assistants [11, 12] and single agent-human adjustable autonomy [13] but can be applied to the problem of human-agent interaction in multiagent systems.

Interaction and the sharing of control between human/agent pairs are complicated by the multiagent environment. For example, what should be done if a participant in the negotiation fails to respond, or responds slowly? What options are available to the agent? Should it involve its owner and if so when? What happens when participants respond with unexpected or unhelpful answers?

Observation When a user makes a decision about what information to share or what action to take in a specific context, this is a demonstration to the agent about how that agent should behave under those exact conditions. Learning from these interactions is an unobtrusive form of learning that can impact an agent's behavior. It has some similarity to case-based reasoning, which has also been applied to learn scheduling preferences [14]. In PTIME we apply learning techniques to such interactions with the user to learn a model of user preferences, including scheduling preferences, location preferences, the importance of a person to a meeting or an event to a person. Each time a user makes a scheduling choice the agent will update its learned model of user preferences. This model is used to produce suggestions to the organizing user in the form of ranked options [4]. Similarly, it can be used to suggest a participant's response to a request for information or an accept/decline decision.

Our work in learning a model of user preferences has improved the ranking of meeting alternatives [4]. However, it is a more difficult task to derive rules for autonomous behavior from such models. Observations are noisy, the environment complex, and choices are often context dependent. Many research challenges remain in learning adjustable autonomy rules through observation. How can an agent break down its observations into chunks that are meaningful for a learner (for adjustable autonomy) and understandable to a user (for feedback). More likely, the observation must be combined with another type of interaction to be useful. For example,

- When an agent makes a mistake, the user can help it understand the problem and correct the learned rule (e.g., I declined this meeting not because of its time but because of its type).
- When an agent learns something, it can ask the user for confirmation that it is learning the correct rule.

Advice An agent should be able to accept its owner's advice and conform to organizational policies. Advice is defined as an enforceable, well-specified constraint on the performance or application of an action in a given situation. There can be hard advice, which can also be called *instruction*,and soft advice, which can be ignored by the agent and may define a preferred application of a constraint or action. [15] defines two types of policy: authorization and obligation. We extend this categorization to include preference:

- *Authorization* defines the actions that the agent is either permitted or forbidden to perform on a target.
- *Obligation* defines the actions that an agent must perform on a set of targets when an event occurs. Obligation actions are always triggered by events, since the agent must know when to perform the specified actions.
- *Preference* defines a ranking in the order, a rule of application, or a selection of an action under certain conditions.

Advice may be conflicting, it can be long-lived, and its relevance may decay over time. Advice can be used to influence the selection of procedures, negotiation protocols, or choices and also to influence adjustable autonomy. The application of advice is central to both PTIME for influencing preference learning and in the future for learning adjustable autonomy strategies. In the latter case, some examples of advice in PTIME are to

- Always accept meeting requests from personname="Bob Smith" or personrole=MyManager.
- Respond automatically to a meeting time rating request when confidence in ratings is "high".
- Never schedule a meeting time after 4pm without my confirmation.

Several research challenges arise from the use of advice in adjustable autonomy. What language can be used for specifying advice? How should an agent apply advice, especially soft advice? Should an agent ever override, or ignore, advice? Finally, how should an agent handle conflicting advice?

Suggestions Suggestion or recommendation systems are often based on the ability to rate or rank options. In the same way PTIME can rate options presented to it and rank them to display to the user or share with another human/agent pair. In addition, PTIME can suggest an alternative option that might suit its owner better than those presented as part of an ongoing negotiation process. An interesting addition to simply rating options according to a static evaluation function is to use an evolving model of the owner's preferences. In PTIME this model is a reflection of previous interactions between the agent and its owner. Some sample suggestions may include

- In response to a *rate meeting times* request: PTIME suggests
  - 9am Tues (medium = 3 stars)
  - 11am Tues (low = 2 stars)
  - 3pm Tues (not suitable = 0 stars)
  - An alternative option 10am Tues (high = 4 stars)
- In response to a *schedule meeting* request: PTIME suggests $ACCEPT$ with confidence level high

Particular challenges are created by the multiagent nature of the domain. Human/agent pairs are not necessarily reliable communicators. One agent may reply to a request immediately, having assumed responsibility for that action. Another may defer to its human owner and wait some undetermined time for a response. The initiating agent can use the mechanism of suggestions to present the user with alternative strategies to handle these particular cases.

Actions In a multiagent environment the agent may have some base level of autonomy and be capable of taking some actions, replying to another agent's request for information, updating a user's calendar, and so on. In a mixed initiative environment the human should be aware of these actions and understand their purpose and context. The research challenge is to enable the agent to move from a base level of autonomy to the ability to act autonomously in a context-sensitive way. Even more challenging would be the ability to acquire this capability through natural interaction with the user. For example, if an agent repeatedly observes the user declining back-to-back meetings, the agent may begin to suggest this response based on its learned model. If the user accepts the suggestion repeatedly, the agent may in the future automatically decline such meetings and notify the user.

Explanation Explanation provides an agent with mechanisms for justifying suggestions or recommendations it may make, or actions it has taken, to a user [16]. A human when deciding to accept or reject a suggestion made by another human will consider the quality of previous recommendations from that person and how that person's interest aligns with his own. Where there is doubt, the human will ask *why*. In the same way, a human user will, over time, learn to trust (or not) the suggestions made by an agent and by extension the actions taken by an agent. While the relationship is developing, the user will want the agent to justify its suggestions and action. Thus, providing sensible and intelligible explanations is beneficial and necessary.

An important subcategory of explanation is the presentation of context to the user. This has several purposes: to help the user navigate the set of possible options, examine an option at multiple levels of abstraction, and annotate an option or change. Visual displays of options, *options at a glance*, are powerful methods of conveying context [17]. For example, the known status of a negotiation process can be displayed visually, or a set of options could be displayed within the known calendar information. An example of context in PTIME is shown in in Figure 4. Here options are displayed on a calendar also displaying shared calendar and preference information relevant to the current context.

In addition to presenting the user with a visual display of context and simple explanations of conflicts in terms of constraints violated, we would like to explore more specific forms of explanation in PTIME. For example,

– PTIME suggests 1pm on Tuesday as an alternative meeting time since everyone is available and PTIME has learned that the user prefers afternoon meetings
– PTIME suggests that the user decline this meeting request because it has learned he does not like back-to-back meetings with this host.

Explanations of a suggestion or action are influenced by the agent's reasoning process, the problem context, and the agent's knowledge of the user's preferences. In most scheduling systems it would be sufficient to explain a solution in terms of constraints violated and evaluation function used. However, if the agent is adaptive and concerned with learning how to be proactive then the challenge is to also explain a suggestion in terms of the learned model, i.e., "why do you think you are allowed to . . . ?" If an agent is going to reveal its learned preferences to the user this has implications. The user may wish to "correct" the agent. Thus, explanations must be dynamic to support user's follow-up questions.

Fig. 4. Context provided by options at a glance. The three options are presented in the context of the participants' schedules. Also, the system has information about two of the participants' preferences and the colors encode that information. The user can delve deeper into the context to look at the meeting details and view preferences individually.

4 Conclusions and Ongoing Work

This position paper sets out a discussion of the role of the dialogue between agent and human in the evolution of adjustable autonomy in the context of the time management domain. The discussion is based on our work on PTIME, a time management assistant that learns the user's preferences, and interacts with its owner to improve over time and earn the user's trust to act autonomously when authorized. As we expand the role of multiagent negotiation within PTIME we can benefit from existing work in mixed-initiative methodologies and distributed meeting negotiation to enhance the user's experience with PTIME. Our work can apply to any domain that has characteristics similar to our multiple human/agent pair environment and negotiation protocols.

Acknowledgments. We thank all the members of the CALO project, especially Khang Duong and Julie Weber. This material is based upon work supported by the Defense Advanced Research Projects Agency (DARPA) under Contract No. FA8750-07-D-0185/0004. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of DARPA or the Air Force Research Laboratory.

References

1. Higa, K., Shin, B., Sivakumar, V.: Meeting scheduling: An experimental investigation. In: Proc. of IEEE Intl. Conf. on Systems, Man, and Cybernetics. (1996) 2023–2028
2. Palen, L.: Social, individual and technological issues for groupware calendar systems. In: Proc. of CHI'99. (1999) 17–24
3. Tullio, J., Goecks, J., Mynatt, E., Nguyen, D.: Augmenting shared personal calendars. In: Proc. of UIST'02. (2002) 11–20
4. Berry, P.M., Gervasio, M., Peintner, B., Yorke-Smith, N.: Balancing the needs of personalization and reasoning in a user-centric scheduling assistant. Technical Note 561, Artificial Intelligence Center, SRI International (2007)
5. Cesta, A., D'Aloisi, D., Brancaleon, R.: Considering the user in mixed-initiative meeting management. In: Proc. of Second ERCIM Workshop on User Interfaces for All. (1996)
6. Bunt, A., Conati, C., McGrenere, J.: Insights from the design and evaluation of a mixed-initiative personalization facility. In: Working Notes of CHI'08 Workshop: Usable Artificial Intelligence. (2008)
7. Kozierok, R., Maes, P.: A learning interface agent for scheduling meetings. In: Proc. of IUI'93. (1993) 81–88
8. Myers, K., Yorke-Smith, N.: Proactive behavior of a personal assistive agent. In: AAMAS'07 Workshop on Metareasoning in Agent-Based Systems. (2007)
9. Chalupsky, H., Gil, Y., Knoblock, C., Lerman, K., Oh, J., Pynadath, D., Russ, T., Tambe, M.: Electric Elves: Applying agent technology to support human organizations. In: Proc. of IJCAI'01. (2001)
10. Scerri, P., Pynadath, D., Tambe, M.: Towards adjustable autonomy for the real world. J. AI Research **17** (2002) 171–228
11. Lamberti, P.M., Prager, J.M.: Advice-giving using reason: An intelligent assistant for interactive computing. In: Seventh IEEE Conference on Artificial Intelligence Applications. (1991) 428–434
12. Rich, E.: Users are individuals: Individualizing user models. Intl. J. Man-Machine Studies **18** (1983) 199–214
13. Isbell, C.L., Pierce, J.S.: An IP continuum for adaptive interface design. In: Proc. of HCI International. (2005)
14. Sycara, K., Zeng, D., Miyashita, K.: Using case-based reasoning to acquire user scheduling preferences that change over time. In: Proc. of IAAI'95. (1995)
15. Sloman, M.: Policy driven management for distributed systems. J. Network and Systems Managements **2**(4) (1994) 333–360
16. Glass, A., McGuinness, D.L., Wolverton, M.: Toward establishing trust in adaptive agents. In: Proc. of IUI'08. (2008) 227–236
17. Berry, P.M., Moffitt, M., Peintner, B., Yorke-Smith, N.: Design of a user-centric scheduling system for multifaceted real-world problems. In: Proc. of ICAPS'07. (2007)