Virtual Butler: What Can We Learn from Adaptive User Interfaces?

Cristina Conati

University of British Columbia, Department of Computer Science,
2366 main Mall, Vancouver, BC, Canada
conati@cs.ubc.ca

Abstract. In this paper, we discuss approaches and results from the field of User-Adaptive Interfaces that we believe can help advance the research on virtual butlers in general, and for the elderly in particular. We list principles underlying the design of effective mixed-initiative interactions that call for formal approaches to dealing both with the uncertainty on modeling relevant cognitive states of the user (e.g., goals, beliefs, preferences), as well as with the tradeoff between costs and benefits of the agent's actions under uncertainty. We also discuss the need for virtual butlers to understand the affective states of their users, and to what extent they need to be transparent by providing means for their users to understand the rationale underlying their adaptive interventions.

1 Introduction

User-Adaptive interfaces (UAI) is an interdisciplinary field that integrates research in Artificial Intelligence (AI), Human Computer Interaction (HCI) and Cognitive Science to create user interfaces that *autonomously* and *intelligently* adapt to the needs of individual users. Providing meaningful adaptation involves building a model of user traits relevant to adequately tailoring the interaction, i.e., a *user model*. Depending on the nature of the task and the extent of the support that the UAI aims to provide, these traits may include simple performance measures (such as frequencies of interface actions), domain-dependent cognitive traits (such as knowledge and goals), meta-cognitive processes that cut across tasks and domains (such as reasoning and learning skills), and affective states (such as moods and emotions).

The field of UAI has much in common with research on devising intelligent home assistants, or virtual butlers. In this paper, we discuss how to apply ideas that have been the focus of UAI research to the development and deployment of virtual butlers. In particular, we introduce principles that were originally proposed as the basis for successful mixed-initiative interactions with UAIs [1]. We then discuss a particular form of user modeling, i.e., modeling of the user's affective states from causes and effects, that can help virtual butlers create a long-term, comfortable relationship with their users. Finally, we address an important issue that we believe is key for the acceptance of technology designed to have a high impact on a user's everyday life, especially if the users are elderly people who may be not comfortable with high-tech

R. Trappl (Ed.): Your Virtual Butler, LNAI 7407, pp. 29–41, 2013.

solutions. This is the issue of transparency/trust, or the degree to which the user needs/wants to understand the rationale underlying the behaviours of an intelligent assistant in order to trust its services.

2 Principles of Mixed-Initiative Interfaces

In the late 1990s', Microsoft research had already put forward the idea of a Virtual Butler for desktop assistance, as a result of lessons learned from the deployment of the Office Assistant for Microsoft Office. The principles underlying the design of this new form of desktop assistance where spelled out in a seminal paper [1]. The key point of the proposed paradigm was that adaptive interactions should be *mixed-initiative*. That is, the user and the adaptive component collaborate to achieve optimal personalization of service. We discuss here a subset of the twelve principles reported in that paper, that we believe are fundamental for the development of successful virtual butlers for the home.

(1) Developing significant value-added automation. Because automated services come with overheads due to their potential lack of transparency and reliability, they should be used only to support tasks that cannot be suitably aided through simpler solutions (see, for instance, the self-cleaning glass vs. glass-cleaning robot example in the chapter by Helmut Stesse).

(2) Considering uncertainty about a user's needs. Being able to provide automated services proactively often requires understanding user goals, beliefs, and preferences. There is bound to be uncertainty in assessing these elements, especially in a rather unconstrained environment like the home, where the user can engage in many different and possibly unrelated activities. This uncertainty should be explicitly taken into account with formal probabilistic techniques rather than by using ad hoc heuristics with less principled theoretical underpinnings.

(3) Employing dialog to resolve key uncertainties. One possible way for a virtual butler to reduce the uncertainty in assessing its user's needs is to ask the user directly. While this is an option that the agent should always consider, the decision of whether to engage in a dialogue with the user should be mediated by an awareness of the potential cost of needlessly bothering the user.

(4) Considering both the costs and benefits of each possible action when deciding what to do. Principle #3 above is an instance of the more general principle of weighing the potential costs and benefits of each possible course of action when deciding how the agent should act, given the uncertainty that permeates the agent's assessment of the user. For instance, if the agent is uncertain about how much a user needs a specific service now, it should evaluate both *costs* and *benefits* of interrupting the user to provide the service, vs. deferring the action to a time when it will be less intrusive, vs. asking the user for more information to help the agent with its decision. Similarly, considerations of potential costs and benefits of actions should help the agent scope the precision of service to match the uncertainty over user goals. That is, doing less but doing it correctly under uncertainty can be more useful that try to provide more specific help that may be unwanted. (related to notion of "useful incompetent helper" in Iba's chapter).

(5) Allowing efficient direct invocation and termination. It is crucial to provide efficient means for users to directly invoke or terminate automated services, to make up for poor decisions by the agents.

(6) Employing socially appropriate behaviours for agent−user interaction. An agent should be endowed with default behaviours and courtesies that match the social expectations for a benevolent assistant. In the context of this book, the expectations are those of the elderly users of virtual butlers.

(7) Continuing to improve over time. This capability can be achieved via two different strategies. The first entails giving the agent the capability to learn by observation, without requiring additional explicit information from the user. The second strategy entails providing mechanisms that allow users to provide explicit feedback to aid the system's learning (e.g., mechanisms to complete or refine inferences and doings of agent).

3 How Far Have UAI Gone?

Adaptive techniques have been investigated for many types of applications, including recommender systems, intelligent tutoring systems, e-games and e-tourism [2] for an overview. There has also been encouraging progress with mixed-initiative approaches, but most of the existing systems tend to be task and/or domain specific. Examples include the TRIPS system for mixed-initiative problem solving [3]; the DiamondHelp mixed-initiative system for task guidance [4]; the MapGen system for mixed-initiative planning, deployed to help with ground operations in the Mars Exploration Rover mission; the Support the Customer (STC) system to provide GE customers with support to diagnose faults in their appliances.

The initiative that so far has gotten closest to the idea of a more general-purpose mixed-initiative intelligent assistant is the CALO project (Cognitive Assistants that Learns and Organizes), sponsored by DARPA. The project involved over 30 institutions in the USA, with the goal of creating cognitive software systems that relieve the workload of knowledge workers by (i) engaging in and leading routine tasks (such as scheduling, task execution, meeting management, information management) and (ii) assisting when the unexpected happens. As stated in the project's official website, *"A CALO should be able to reason, learn from experience, be told what to do, explain what it is doing, reflect on its experience, and respond robustly to surprise."* These are capabilities that we may want in a house assistant, and the CALO project made substantial progress on the necessary machine learning aspects of this research [30, 31]. Less progress has been made on successfully applying the proposed technologies in practice, in a user-friendly, mixed-initiative fashion.

Why, then, despite the many years of research and the many resources devoted to adaptive and mixed-initiative systems, are we far from general-purpose intelligent assistants? The reason is that any form of virtual butler is "AI-complete"; that is, it requires dealing with all of the traditional AI challenges (knowledge representation, reasoning, planning, problem solving, natural language processing), with the additional complexity that these challenges need to be tackled for many different task domains and for interacting properly with the user. In the rest of the paper, we will focus on the latter problem (interaction with the user), specifically on two issues that we believe are especially relevant for supporting smooth interaction between an elderly user and a virtual butler: (i) enabling the virtual butler to model and respond to the user's affective states; (ii) enabling the user to understand the reasons behind the virtual butler's interventions.

4 Modeling User Affect from Causes and Effects

What should a virtual butler understand about its users? Certainly their goals, preferences and beliefs (which is challenging enough), but also their *moods* and *emotions* (e.g., *affect*), if we want an agent that can create a long-lasting, balanced and comfortable relationship with the user [5].

Recent years have seen a flourishing of research directed towards adding affective components to human–computer interactions, with the assumption that "affect-sensitive interfaces" can better meet users' needs by creating a more natural dialogue between humans and computers. One key element of this endeavour is the computer's capability to recognize the user's emotional states during the interaction, i.e., to have a model of the user's affect (or *affective user model*). Possible sources of information to assess user emotions include causal information on both context and the person's relevant traits, as well as diagnostic information on visible bodily reactions. However, the information provided by these sources is often ambiguous and even contradictory, making emotion assessment a task riddled with uncertainty, especially in situations that can give rise to multiple emotions, possibly overlapping and rapidly changing, as for instance during the course of an emotionally charged conversation.

Consistently with the second principle listed in Section 2 (*Considering uncertainty about a user's needs*), to handle this uncertainty we have proposed a probabilistic framework for modeling user affect that uses Dynamic Decision Network (DDN) [6] to leverage information on both the possible causes and the observable effects of the user's affective reaction [7, 33]. Most existing research in modeling user affect has focused on devising models that can capture *which* affective states a user is experience during a given interaction. [8, 10, 9, 11, 12, 32]. Our approach is designed to also provide insights on *why* a user is in a particular affective state, thus better enabling an interacting agent to adequately respond to the user's emotional reactions.

4.1 The Affect Modeling Framework

A DDN is a graph where nodes represent either stochastic variables of interest or points where an agent needs to make deliberate decisions. Arcs in the graph capture the direct probabilistic relationships between the nodes. Each node has an associated probability distribution representing the conditional probability of each of its possible values, given the values of its parent nodes. As evidence on one or more network variables becomes available, *ad hoc* algorithms update the posterior probabilities of all the other variables, given the observed values.

Fig. 1 shows a high-level representation of two time-slices in our DDN-based framework for affective modeling. Each time slice represents the model's variables at a particular point in time and, as the figure shows, the network can combine evidence on both the causes and effects of emotional reactions to assess the user's emotional state. The sub-network above the nodes *Emotional States* is the predictive component of the framework. It represents the relations between possible causes and emotional states as described in the OCC cognitive theory of emotions [15]. According to this theory, emotions derive from cognitive appraisal of the current situation, which consists of events, agents, and objects. The outcome of the appraisal depends on how the situation fits with the individual's goals and preferences. For instance, depending on whether the current event (e.g., the outcome of an interface agent's action) does or does not fit with the individual's goals, the person will feel either *joy* or *distress* toward the event. Correspondingly, if the current event is caused by a third-party agent, the person will feel *admiration* or *reproach* toward the agent; if the agent is oneself, the person will feel either *pride* or *shame*. Based on this structure, the OCC theory defines 22 different emotions, described in terms of their valence and the entity they relate to.

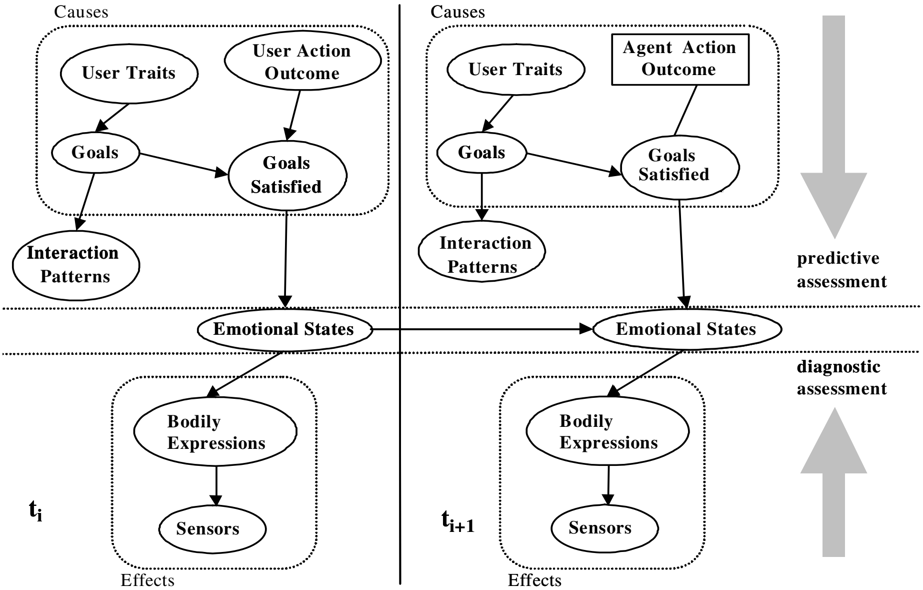

Fig. 1. High-level representation of the DDN for affective user modeling

We adopted this particular theory of emotion for our affective modeling framework, rather than alternative models that define emotions in terms of their level of valence and arousal [13], because its clear and intuitive representation of the causal nature of emotional reactions lends itself well to devising computational models that can assess *why* a user feels given emotions, as well as *what* these emotions are. This more fine-grained information enhances the capability of an interactive agent to adequately respond to user affect. For instance, if the agent can recognize that the user feels a negative emotion because of something wrong she has done (*shame* by OCC definition), it can decide to provide hints aimed at making the user feel better about herself. If the agent recognizes that the negative feelings are caused by its own behaviour (*reproach* by OCC definition) it may decide to take actions to make amends with the user. These specific interventions are not possible with approaches that detect emotions with no explicit knowledge of their causes [32].

Our OCC-based DDN for affective user modeling includes variables for goals that a user may have during the interaction with an autonomous agent, (nodes *Goals* in Fig. 1). The events subject to the user's appraisal are any outcomes generated by the user's or the agent's action (nodes *User Action Outcome* and *Agent Action Outcome* in Fig. 1). Agent action outcomes are represented as decision variables in the framework, indicating points where the agent decides how to intervene in the interaction. The desirability of an event in relation to the user's goals is represented by the node class *Goals Satisfied,* which in turn influences the user's *Emotional States* (we will call this part of the model *appraisal-subnetwork* from now on)

The user's goals are a key element of the OCC model, but assessing these goals is not trivial, especially when obtaining them directly from the user would be too intrusive. Thus, our DDN also includes nodes to infer user goals from indirect evidence (*goal-assessment subnetwork*). User goals can depend on *User Traits* such as personality and can influence user *Interaction Patterns*, which in turn can be inferred by observing the outcomes of individual user actions. Thus, both the relevant user traits and action outcomes can be used in the DDN as evidence for assessing user goals. The subnetwork below the nodes *Emotional States* is the diagnostic part of the affective modeling framework, representing the interaction between emotional states and their observable effects. *Emotional States* directly influence user *Bodily Expressions*, which in turn affect the output of *Sensors* that can detect them. Because in many situations a single sensor cannot reliably identify a specific emotional state, our framework is designed to modularly combine any available sensor information, and gracefully degrade in the presence of partial or noisy information.

4.2 Using the Framework for Affective User Modeling

We have instantiated the modeling framework described in the previous section to model user emotions during the interaction with PrimeClimb, an educational game designed to help students practice number factorization. The game includes a pedagogical agent that can provide individualized support when the student does not seem to be learning from the game [17]. Therefore, the affective model is designed to capture both feelings generated by the player's performance in the game (labeled as *pride/shame* in the OCC theory) as well as feelings generated by the agent's interventions (labeled as *admiration/reproach* in the OCC theory). The model also captures user's emotions towards game states (labeled as *joy/regret* in the OCC model).

As part of the iterative design and evaluation approach we adopted to build the affective model, we started by instantiating and evaluating the predictive part of our modeling framework. Creating the predictive part of the model required several user studies to identify common user goals during game playing, the probabilistic relationships between student personality traits, goals and interaction patterns to define the goal assessment network [17], and the probabilities that the outcomes of various student and agent actions satisfy each of the possible goals [18] for the appraisal network. We then experimented with adding to the model a diagnostic component that uses Electromigraphy (EMG) sensors placed on the user's forehead to detect frowns as signs of negative affect. Preliminary results show that with this model can achieve up to 73% accuracy in modeling emotions towards the agent [33].

4.3 Using the Framework for a Socially Intelligent Virtual Butler

The main advantage of the affective modeling approach described above is that, by explicitly modeling causes of affect, it gives an agent fine-grained information on how to respond to the affect. The second advantage is that it is flexible in taking advantage of the available sources of affective information, leveraging data on both causes and effects when available, but still being able to degrade gracefully when any of the potential information sources becomes unavailable or unreliable. This flexibility is enhanced if one adds to the model goals that explicitly represent lower level dimensions of affective reactions, i.e., valence and arousal. These dimensions are generally easier to assess than specific emotions, so the system has the chance to "*do less but do it more accurately*" in the presence of high level of uncertainty over a user's specific emotions, as suggested by the mixed-initiative principle #3 in Section 2. A third advantage is that the framework lends itself well to be used by an agent that takes both costs and benefits of its actions into account when deciding how to act. Dynamic decision networks are set up specifically to support a decision-theoretic approach to agent behaviour. In a decision-theoretic model [24], costs and benefits of agent behaviours are expressed as preferences over world states *S* (e.g., the possible affective states of a user for an affect-sensitive agent). In turn, these preferences are encoded via a utility function *U(S)*, which assigns a single number to express the desirability of each state S. Furthermore, for each action *a* available to the agent, and for each possible outcome state S' of that action, P(S'|E, *a*) represents the agent's belief that action *a* will result in state S', when the action is performed in a state identified by evidence E. The expected utility of an action *a* is then computed as

$$EU(A) = \Sigma_{S'} P(S'|E, a)U(S')$$

A decision-theoretic agent selects the action that maximizes this value when deciding how to act. DDNs allow modeling decision-theoretic behaviour via the inclusion of nodes that represent an agent's utilities, in addition to nodes representing probabilistic events in the world and the agent's decision points. By relying on propagation algorithms for Bayesian networks, DDNs allow computing the agent's action (or sequence of actions) with maximum expected utility given the available evidence on the current state of the world.

One potential drawback of our proposed approach is that it requires the detailed modeling of user goals and preferences, as well as how a user appraises different circumstances in the surrounding environment based on these goals and preferences. This modeling is bound to require a substantial amount of empirical data for each new user in order to be done accurately, but its cost is lessened by the fact that a virtual butler needs to understand user goals and preferences regardless of whether or not it includes an affective component. What is left then is the cost of modeling a user's appraisal criteria, and understanding how this cost compares with the gain in quality of the system's affective responses.

There are also two approach-independent issues that need to be investigated in order to devise emotionally intelligent virtual butlers. The first is deciding which emotions the virtual butler should be able to capture. The second is what should the virtual butler do to respond to those emotions.

The first issue requires, again, evaluating the cost of modeling each additional emotion against the value that can be gained in terms of usefulness/acceptance/impact of the virtual butler. The OCC model, for instance, accounts for twenty-two different emotions, including emotions related to feelings towards aspects of an entity (*liking, disliking, love, hatred*) and emotions related to appraising events in terms of their usefulness for others (*happy-for, resentment, gloating, pity*) or in terms of expected consequences for self *(hope, fear).* While it is very possible that most users may encounter each of these emotions at one point or another, it is necessary to evaluate which ones are prominent enough to warrant attention, and among these which ones can/should be a concern for a virtual butler. For instance, the only reason to model a morally negative but positively valenced emotion such as *gloating* would be to try and discourage it, but this should hardly be a mission for the virtual butler for an elderly user.

Once the affective states that the virtual butler should recognize have been determined, the question becomes how they should affect the butler's behaviour. There are at least two levels at which affective information can be included in the agent's operation. One level, which we will define as *affect oriented*, involves having the agent respond to the user's affective state to directly influence it; for instance, one could envision a virtual butler that can detect its user's negative affective states and act specifically to help the user overcome them. A second level, which we will call *task oriented*, sees affective information as one of the factors that the agent must take into account to decide how to best accomplish a given task. Suppose, for instance, that a virtual butler needs to communicate to its user that her friend is cancelling a plan to go and see a movie the following evening, and that the agent has the choice to do it right away or wait until later in the day. While giving the news right away would give the user more time to make alternative arrangements if desired, the agent may decide to delay the action if it detects that the user is already in a negative state, especially if it thinks that the negative state is caused by feelings of being lonely.

Both these levels will require extensive investigations to define the impact of the agent's actions on an elderly user's affect. These investigations may be able to leverage existing theories on affective interventions from emotional psychology and existing knowledge on the effects of emotionally responsive artificial agents in domains other than domotics for the elderly [19, 21, 22, 20, 23, 34]. However, because there is very limited work on the dynamics of affective interactions between artificial agents and elderly users, ad-hoc empirical studies will need to be conducted to fill the theoretical gaps, especially given the focus on users who may be less familiar with artificial agents and perhaps less open to the idea of having empathic relations with one.

5 Enabling the User to Understand Its Butler

User modeling allows an adaptive agent to understand its user, but shouldn't the user also understand the system? The term "understand" here relates to comprehending the *rationale* underlying the agent's interventions and suggestions, and is connected to one of the main principles of good design in HCI: interface *transparency*, or the extent to which a user can understand system actions and/or has a clear picture of how the system works [2]. There are at least two reasons to believe that allowing an agent to expose the rationale underlying its behaviours to the user may improve its effectiveness. The first is that this capability would greatly improve the mixed-initiative aspect of the agent-user interaction, because the agent and the user can *discuss* the agent's decisions based on how well the agent can justify them, as opposed to having a one-shot mixed-initiative exchange where all the user can do is either accept or reject the agent's service. The second is that understanding the rationale underlying an agent's behaviour may increase the user's *trust* in the agent: the user may not be as put off by an agent's less-then-ideal action if the agent can show that it was suggested based on reasonable assumptions and sound reasoning.

There are numerous examples of adaptive or mixed-initiative systems that provide access to all or part of their rationale. For example, there are adaptive systems for education that include *inspectable student models*. These systems allow users to view and sometimes edit their student model, with the assumption that these operations give users a sense of what causes the particular adaptive behaviour to occur [25, 26]. Provision of rationale has also been explored in recommender systems [27], in expert systems [28] and in mixed-initiative approaches to interface customization [29] Evaluations provide encouraging evidence that the rationale can increase system transparency [26, 29], promote reflection [26], and improve users' reactions to system recommendations [27]. If not properly designed, however, rationale can be difficult to use [25, 29] and can even lead to less favourable responses towards the system [27]. [29] also showed that interest and willingness to look at the system rationale are strongly user dependent. In their work, rationale relates to describing to the user how MICA, an adaptive system that supports user customization of MS Word menus/toolbars, generates its customization suggestions.

MICA tries to identify the user's optimal personalized interface (PI from now on) by evaluating which menu and toolbar items should be included in the PI and which should reside solely in the full interface (FI), accessible via a button click in the PI. MICA then generates corresponding customization suggestions. To do so, MICA relies on a user model that assesses the user's *time performance* given a particular PI.

The performance assessment relies primarily on three factors: (1) *Usage Frequencies:* how often the user is expected to access each menu item; (2) *Expertise:* user's familiarity with the existing menus, to account for the fact that users with lower expertise are likely to be more negatively impacted by excess functionality; (3) *Interface Size:* detailed layout information on the FI and the PI under consideration, including the number of items present and where they are located.

MICA's rationale component describes why the system is making recommendations and the relevant user- and interface-specific factors influencing its decision-making process. Presenting this rationale has the potential to provide valuable insight into how the system works; however, effectively communicating the information to the average user is a challenging design task, particularly since MICA's algorithm is relatively complex. [29] dealt with the challenge via an iterative design and evaluation process that led to the rationale-explanation component shown in Figure 2.

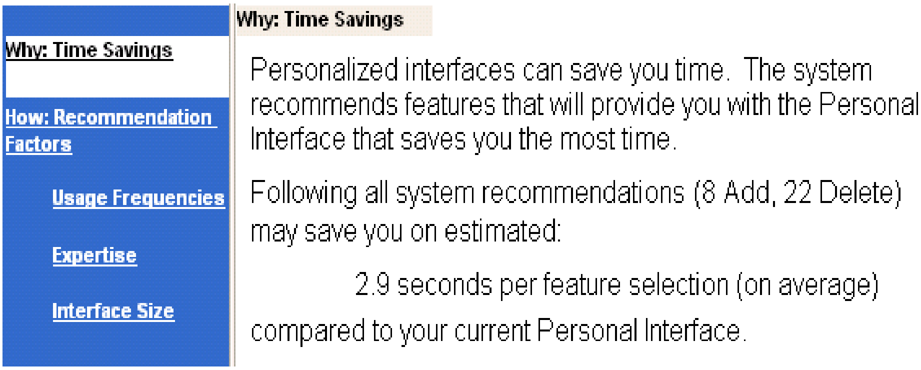

Fig. 2. "Why" component of MICA's rationale explanation

In this component, the user can access the rationale as soon as MICA generates a recommendation for customization. Once invoked, the dialogue box in Fig. 2 appears, including information on *why* and *how* the system makes recommendations. The "Why" component, displayed on the right of Fig. 2, indicates that the recommendations are based on time savings and provides an estimated savings per feature invocation (based on the user model's performance assessment) should the user choose to accept all recommendations.

The "How" component is a simplified explanation of MICA's decision-making. The first screen, "How: Recommendations Factors," explains that MICA balances the three factors described above (Usage Frequencies, Expertise, Interface Size). Next, three screens describe each factor in greater detail (two are shown in Fig. 3).

Findings from a formal qualitative study on the acceptance and impact of MICA's rationale functionality [29] indicate that the majority of users prefer to have the rationale present, but that a not-insignificant group of users do not need or want the information. For some users, the rationale led to increased trust, understanding, predictability, and motivation to accept recommendations. Others, however, felt that the rationale was just common sense, or was unnecessary in a mixed-initiative system or productivity application. Some users said that they did not need to see the rationale because they had an inherent trust in the system.

These findings suggest that, contrary to previously stated guidelines [2], system transparency may not, in fact, be important to all users in all contexts. But we should bear in mind that the tasks studied in [29] related to a productivity application, and were obtained in a laboratory setting with no serious consequences for having suboptimal task performance. These circumstances are likely to reduce the user's need to make sure that a system's suggestions are actually worth following, especially when weighed against the time and effort required to parse the system's explanations. Thus, there is a substantial amount of groundwork that needs to be done to assess how important system transparency is in the context of the interaction of elderly users with their virtual butler. If transparency turns out to be important, then researchers will need to focus their effort on understanding

- which level of understanding a virtual butler needs to promote (e.g., visualization of the system's user model vs. more explicit explanations of the inferences that generated the current model's predictions)
- How to best promote the chosen level(s) from an HCI point of view.

6 Discussion and Conclusions

In this paper, we have discussed approaches and results from the field of User-Adaptive Interfaces that we believe can help advance the research on virtual butlers in general, and for the elderly in particular. We have listed principles underlying the design of effective mixed-initiative interactions that call for formal approaches to dealing both with the uncertainty on modeling relevant cognitive states of the user (e.g., goals, beliefs, preferences), as well as with the tradeoff between costs and benefits of the agent's actions under uncertainty. We have also discussed the need for virtual butlers to understand the affective states of their users, and we have introduced a framework for affective user modeling that leverages both causes and effects of emotional reactions to assess the user's specific emotions, and why they arise. Finally, we have addressed the issue of system transparency, i.e., whether it is important/feasible that elderly users understand the rationale underlying the interventions of their butlers in order to make the most effective use of them.

What should emerge from this chapter, and from this book overall, is that there are still many more open questions than solutions along the road to devising Virtual Butlers for elderly citizens. Finding answers to these open questions should be a multidisciplinary endeavor, where psychologists and sociologists study the general principles underlying the interaction of elderly users with these kinds of advanced technologies, and IT experts use these principles to shape the technologies so that they can best suit this specific user population.

References

1. Horvitz, E.: Principles of Mixed-Initiative User Interfaces. In: Proceedings of CHI 1999, ACM SIGCHI Conference on Human Factors in Computing Systems, Pittsburgh, PA (1999)
2. Jameson, A.: Adaptive Interfaces and Agents. In: Sears, A., Jacko, J. (eds.) The Human-Computer Interaction Handbook: Fundamentals, Evolving Technologies and Emerging Applications, 2nd edn., pp. 433–458. CRC Press, Boca Raton (2008)
3. Ferguson, G., Allen, J.: Mixed-Initiative Dialogue Systems for Collaborative Problem-Solving. AI Magazine: Special Issue on Mixed-Initiative Assistants 28(2), 2–32 (2007)
4. Rich, C., Sidner, C.: DiamondHelp: A generic collaborative task guidance system. AI Magazine: Special Issue on Mixed-Initiative Assistants (2007)
5. Bickmore, T., Picard, R.: Establishing and maintaining long-term human-computer relationships. ACM Trans. Comput.-Hum. Interact. 12(2), 293–327 (2005)
6. Dean, T., Kanazawa, K.: A Model for Reasoning about Persistence and Causation. Computational Intelligence 5(3) (1989)
7. Conati, C.: Probabilistic Assessment of User's Emotions in Educational Games. Journal of Applied Artificial Intelligence, Special Issue on Merging Cognition and Affect in HCI 16(7-8), 555–575 (2002)
8. Healey, J.A., Picard, R.W.: Detecting Stress During Real-World Driving Tasks Using Physiological Sensors. IEEE Transactions on Intelligent Transportation Systems 6(2) (2005)
9. Prendinger, H., Mori, J., Ishizuka, M.: Recognizing, Modeling, and Responding to Users' Affective States. In: Ardissono, L., Brna, P., Mitrović, A. (eds.) UM 2005. LNCS (LNAI), vol. 3538, pp. 60–69. Springer, Heidelberg (2005)
10. Litman, D.J., Forbes-Riley, K.: Predicting Student Emotions in Computer-Human Tutoring Dialogues. In: 42nd Annual Meeting of the Association for Computational Linguistics, ACL (2004)
11. Mandryk, R.L., Inkpen, K.M., Calvert, T.W.: Using Psychophysiological Techniques to Measure User Experience with Entertainment Technologies. Journal of Behavior and Information Technology 25 (2006)
12. D'Mello, S.K., et al.: Automatic Detection of Learner's Affect from Conversational Cues. User Modeling and User-Adapted Interaction 18(1-2) (2008)
13. Lang, P.J.: The emotion probe: Studies of motivation and attention. American Psychologist 50 (1995)
14. D'Mello, S.K., Chipman, P., Graesser, A.C.: Posture as a predictor of learner's affective engagement. In: Proceedings of the 29th Annual Cognitive Science Society (2006)
15. Ortony, A., Clore, G.L., Collins, A.: The Cognitive Structure of Emotions. Cambridge University Press (1988)
16. Lang, P.J., et al.: Look at Pictures: Affective, Facial, Visceral, and Behavioral Reactions. Psychophysiology 30 (1993)
17. Zhou, X., Conati, C.: Inferring User Goals from Personality and Behavior in a Causal Model of User Affect. In: Proceedings of IUI 2003, International Conference on Intelligent User Interfaces, Miami, FL, U.S.A, pp. 211–218 (2003)
18. Conati, C., Maclaren, H.: Data-Driven Refinement of a Probabilistic Model of User Affect. In: Ardissono, L., Brna, P., Mitrović, A. (eds.) UM 2005. LNCS (LNAI), vol. 3538, pp. 40–49. Springer, Heidelberg (2005)
19. Klein, J., Moon, Y., Picard, R.: This Computer Responds to User Frustration: Theory, Design, Results, and Implications. Interacting with Computers 14(2), 119–140 (2002)

20. Burleson, W., Picard, R.W.: Gender-Specific Approaches to Developing Emotionally Intelligent Learning Companions. IEEE Intelligent Systems 22(4), 62–69 (2007)
21. Prendinger, H., Ishizuka, M.: The Empathic Companion: A Character-Based Interface That Addresses Users' Affective States. Applied Artificial Intelligence 19(3-4), 267–285 (2005)
22. Paiva, A., Dias, J., Sobral, D., Aylett, R., Woods, S., Hall, L.E., Zoll, C.: Learning By Feeling: Evoking Empathy With Synthetic Characters. Applied Artificial Intelligence 19(3-4), 235–266 (2005)
23. McQuiggan, S.W., Rowe, J.P., Lester, J.C.: The effects of empathetic virtual characters on presence in narrative-centered learning environments. In: CHI, pp. 1511–1520 (2008)
24. Howard, R.A., Matheson, J.E. (eds.): Readings in Decision Analysis. Decision Analysis Group, SRI International, Menlo Park, California (1977)
25. Czarkowski, M., Kay, J.: How to Give the User a Sense of Control over the Personalization of Adaptive Hypertext? In: Proc. of Adaptive Hypermedia and Adaptive Web-Based Systems (in Conjunction with UM 2003), pp. 121–131 (2003)
26. Zapata-Rivera, J.D., Greer, J.E.: Interacting with Inspectable Bayesian Student Models. International Journal of AI in Education 14 (2003)
27. Herlocker, J., Konstan, J.A., Riedl, J.: Explaining Collaborative Filtering Recommendations. In: Proc. of CSCW, pp. 241–250 (2000)
28. Horvitz, E., Breese, J., Henrion, M.: Decision Theory in Expert Systems and Artificial Intelligence. Journal of Approximate Reasoning 2, 247–302 (1988)
29. Bunt, A., McGrenere, J., Conati, C.: Understanding the Utility of Rationale in a Mixed-Initiative System for GUI Customization. In: Conati, C., McCoy, K., Paliouras, G. (eds.) UM 2007. LNCS (LNAI), vol. 4511, pp. 147–156. Springer, Heidelberg (2007)
30. Gervasio, M., Murdock, J.: What Were You Thinking? Filling in Missing Dataflow Through Inference in Learning from Demonstration. In: Proc. IUI 2009. ACM Press (2009)
31. Shen, J., Fitzhenry, E., Dietterich, T.G.: Discovering frequent work procedures from resource connections. In: IUI, pp. 277–286 (2009)
32. Arroyo, A., Cooper, D., Burleson, W., et al.: Emotion Sensors Go To School. In: AIED, pp. 17–24 (2009)
33. Conati, C.: Combining cognitive appraisal and sensors for affect detection in a framework for modeling user affect. In: Calvo, R.A., D'Mello, S.K. (eds.) New Perspectives on Affect and Learning Technologies. Springer, New York (in press)
34. D'Mello, S., Lehman, B., Sullins, J., Daigle, R., Combs, R., Vogt, K., Perkins, L., Graesser, A.: A Time for Emoting: When Affect-Sensitivity Is and Isn't Effective at Promoting Deep Learning. In: Aleven, V., Kay, J., Mostow, J. (eds.) ITS 2010, Part I. LNCS, vol. 6094, pp. 245–254. Springer, Heidelberg (2010)