Lamboudis, D. & Economides, A. A. (2002). Managing time thresholds in mixed-initiative learning environments. In M. Driscoll and Th.C. Reeves (eds.) Proceedings E-Learn 2002 World Conference on E-Learning in Corporate, Government, Healthcare & Higher Education, pp. 1746-1749, Montreal, Canada, 15-19 October, AACE (Association for the Advancement of Computing in Education). ISBN: 1-880094-46-0.

Managing Time Thresholds in Mixed-Initiative Learning Environments

Dimitris Lamboudis, Anastasios Economides
University of Macedonia
Thessaloniki, Greece
dimlamb@uom.gr, economid@uom.gr

Abstract, The effectiveness of pedagogical agents, in terms of their believability and adaptivity, in mixed-initiative learning environments can be considerably affected by the "timing" of agents' actions. This paper describes an approach to manage time thresholds, which takes into account the response "style" of each individual learner. The proposed approach has provided very positive results in the context of an agent-based, mixed-initiative learning environment, which we are currently developing.

Introduction

Mixed-initiative problem solving lies at the heart of knowledge-based learning environments, aiming to provide an individualised learning experience. Learning environments, which are inhabited by animated pedagogical agents, constitute one of the most prominent paradigms of mixed-initiative systems (Johnson, 2000), (Johnson et al, 2000), (Andre et al, 1999). While learners are actively engaged in problem solving activities, agents monitor their progress and provide to them feedback which is according to their individual profile, aiming to increase learning effectiveness and efficiency. Typical scenarios include an introduction to the subject by the agent, followed by a test, or a task to be accomplished, to evaluate the learner's level of knowledge acquisition (Lester et al, 1999). While the learner tries to perform the task in hand, the agent monitors his/her actions, and continuously evaluates various factors in order to engage and assist. The pedagogical agents' potential to couple feedback functionalities with a strong visual presence, makes them an ideal example for studying mixed-initiative interactions (Shaw et al, 1999).

Mixed-initiative systems must consider a set of key decisions in their effort to support joint activity, including: *when* to engage learners with a service, *how* to best contribute to solving a problem, *when* to pass control back to users, and *when* to query users for additional information (Horvitz, 1999). In order to reach to *situated* decisions, agents make "guesses" about learners' needs. These "guesses" usually depend on the evidence obtained through the "keyhole" of the user interface, collaborative statistical data about the learner (Zuckerman and Albreciit, 2001) and explicitly asked information, most commonly in the form of a query the user has to go through in the beginning of a session.

Moreover, in the context of such systems, there are many additional parameters and principles that should be taken into account by the agent in order to reach a turn taking decision (Horvitz, 1999), (Lester et al, 1999), (Bates, 1994). For example, the personality and the emotional state of the agent, the advisory history, the idle time elapsed etc.

This paper focuses on a specific aspect of agent behaviour in this context, namely *timing*, which is directly connected (among others), to *time thresholds*. Time thresholds are defined as the amount of time the learner is allowed to spend in order to successfully complete a task without the agent's help.

Idle time is one of the main variables used by the agents to infer that the learner has difficulties in understanding, solving and, in general, successfully proceeding in the learning procedure. Idle time is usually defined over a predefined time threshold, i.e. agents compare the time elapsed until the learner response, with a predefined time threshold. The actual value of this threshold is most of the times derived from statistical data and corresponds to the "mean time" that learners spend for the completion of the particular task (Lester et al, 1999).

It can be argued, however, that this approach has two main limitations:

- The pattern of the agent's behaviour is soon revealed to the student; this fact may considerably compromise agent's believability, and therefore, agent's effectiveness (Lester et al, 1997a).
- It does not take into account the individual learner's characteristics (it is rather targeted to the "average learner"), which is the main objective in personalized learning environments (Sampson, Karagiannidis, Kinshuk, 2002)

This paper proposes an alternative approach for managing time thresholds in agent-based, mixed initiative learning systems. The proposed algorithm takes into account the response of each individual learner, and has provided very positive results in the context of an agent-based, mixed-initiative learning environment, which we are currently developing.

The Proposed Algorithm

Our approach attempts to overcome the limitations identified in the previous section. The algorithm, instead of spontaneously engaging the agent when idle time exceeds a predefined time threshold, allows the learner to have "a second chance" by extending the threshold. This second chance (the extension of the predefined time threshold) is not provided unconditionally, since this would be equivalent to just set another rigid threshold, although greater than the initial one. Instead, when the threshold is reached, the agent decides to extend it by some probability P_e and not to extend it by some probability $P_a = 1\text{-}P_e$. Thus, in the "worst case", the agent will behave conventionally, i.e. like in the existing systems. However, there is a possibility, which is partially defined by the designer, at least as far as the initial value of P_e is concerned, that the agent will give the learner a second chance. Yet, if this possibility is heavily depending on the initial value of P_e, it would be just another ad hoc intervention of the designer, lacking any adaptive characteristics.

Instead, the probability of extending the time threshold (i.e. the definition of P_e), is determined by the agent, through the algorithm which checks if this extension of time has any affects on the learning procedure, that is, if it helps the learner to achieve his/her goals. In case it does, it reinforces the value of P_e. In the long run, this means that independently of the initial values of P_e and P_a the system will favour the option that actually helps the learner.

In more detail, the algorithm is shown in (Fig. 1). The corresponding notation and assumptions are as follows:

- A learning procedure that can be represented by a set of n hierarchically ordered Tasks, $T=\{T_i, i=1... n\}$;
- A set of corresponding time thresholds, $t=\{t_i, i=1... n\}$;
- An initial value of P_e^0 (the corresponding $P_a^0 = 1\text{-} P_e^0$);
  Where $P_e^0 = P$ *(extend time threshold in* $T_1 \mid t > t_1$*)*, i.e. the conditional probability of extending the time in the first task, given that the time threshold has been reached;
- A constant ?p to represent the reinforcement of P_e^{i-1};
  Where $P_e^{i-1} = P$ *(extend time threshold in* $T_i \mid t > t_i$*)*, i.e. the conditional probability of extending the time in the i_{th} task, given that the time threshold has been reached;
- A constant ?t to represent the extension amount of the time threshold; (Both ?p and ?t can be either constant values or percentages);
- A Boolean e, to serve as a flag to declare if there has been a time extension $e=1$ or not $e=0$;

Given these assumptions the algorithm receives t_i and P_e^{i-1} as input and process them as follows:

While (there is no positive response) Do
    if (e=1) then
        act and proceed to Task $(t_i\text{-}?t, P_e^{i-1})$
    else
        If $(t > t_i)$ *then*
            $t_i = t_i + ?t$ *by* P_e^{i-1} *and* $e=1$;
            act by P_a^{i-1} *and proceed to Task* $(t_i\text{-}?t, P_e^{i-1})$
end while
if (e=1) then
    $P_e^i = P_e^{i-1} + ?p$
    $P_a^i = 1\text{-} P_e^i$
Next Task (t_{i+1}, P_e^i)

Note that in case of action-taking by the agent, the algorithm remains to the same task but with a decreased threshold. Since the learner has received help he/she should be able to complete the task in less time. The agent is assumed to provide exhaustive help if needed; thus the algorithm ends with a positive response and avoids stack overflow.

The proposed algorithm can overcome the shortcomings of existing approaches:

1. The use of probabilities, instead of a predefined time threshold, increases the possibility that the agent behaviour will not be revealed to the learner, thus the believability of our agents can be enhanced.
2. Moreover, the time thresholds are dynamically adapted to the individual learner response, thus contributing to a learning experience which is driven by the learner characteristics

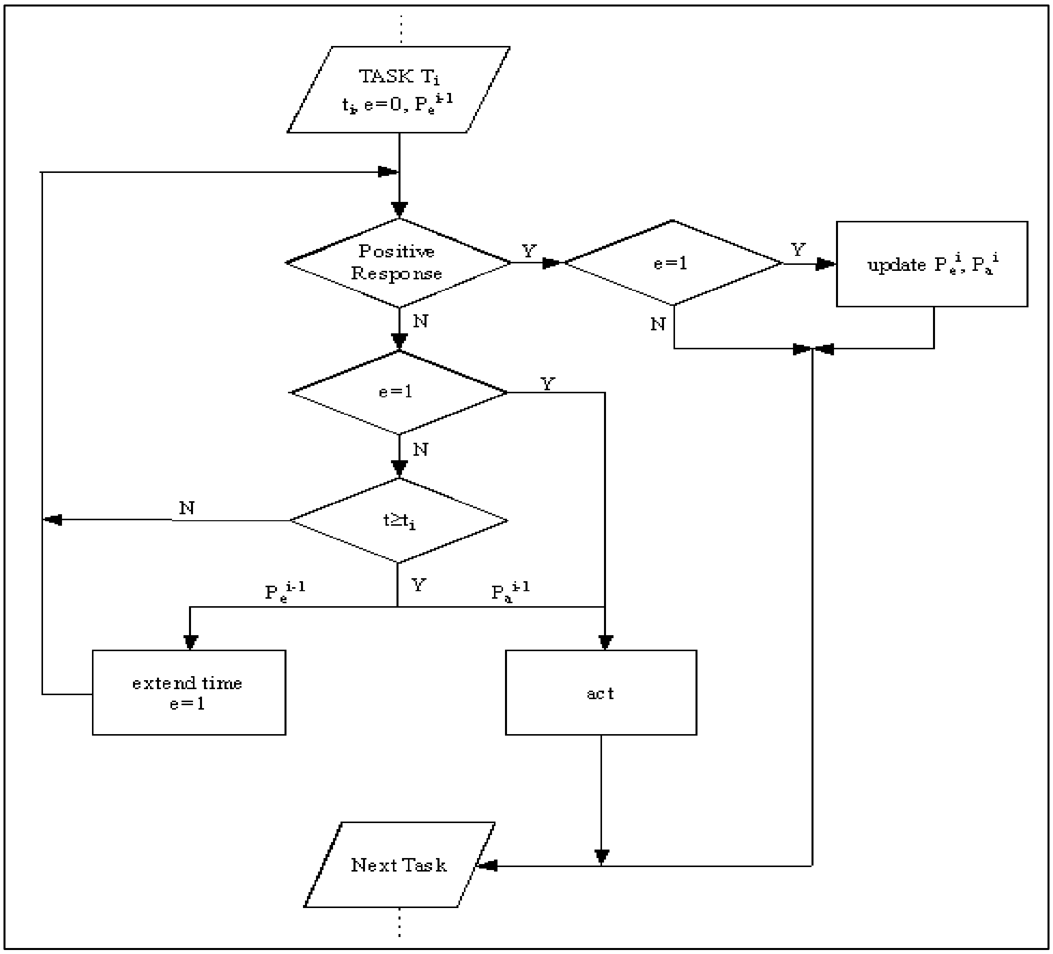

Figure 1: The proposed algorithm for managing time thresholds in agent-based, mixed-initiative learning systems

It is important to note that:

- The algorithm controls *when* an intervention will occur and not if the intervention will take place or not. The agent *will*, after all, engage if needed. Thus the algorithm maintains the pedagogical mainstream of decisions but introduces small variations to adapt to the particular user's style.
- The case of decreasing the value of P_e is not concerned. It seems rational to do so if the agent receives positive feedback before the threshold is reached but this would have the affect of balancing the value of P_e close to the initial value thus cancelling the adaptivity of the algorithm. After all, in the worse case, the agent will wait a bit longer before acting.
- Keeping the sequence of P_e ascending may lead to a value close or even equal to the unit ($P_e = 1$). This should not be avoided in general, since it actually means that the need of the user for extra time is very strong. However, in order to maintain the second characteristic of the algorithm (hiding the behaviour pattern), the designer could incorporate a control for the upper limit of the value of P_e.

Discussion

This paper has proposed an algorithm for managing the "timing" of agents actions in the context of mixed-initiative learning systems. The idea of the algorithm is to modify the timing of agents interventions, based on a probabilistic model, which takes into account the responses of each individual learner.

The proposed algorithm is being implemented in the context of an agent-based, mixed-initiative learning environment that we are currently developing. We have conducted some preliminary experiments with students of our department to evaluate the algorithm. In particular, we interviewed students using the system with pre-defined time thresholds, against students using the system where thresholds follow the proposed approach. This informal evaluation provided very positive feedback: the fact that agent's behaviour cannot be easily determined by the students, makes the agent believable, and enhance their learning effectiveness.

References

Andre E., Rist T., and Moeller J. (1999). Employing AI Methods to Control the Behaviour of Animated Interface Agents. *Applied Artificial Intelligence Journal, 13 4-5*.

Bates, J. (1994). The Role of Emotion in Believable Agents. *Communications of the ACM,* 37 (7), 122-125.

Horvitz, E. (1999). Principles of Mixed-Initiative User Interfaces. *Proceedings of CHI '99, ACM SIGCHI Conference on Human Factors in Computing Systems, Pittsburgh, PA, May 1999.*

Johnson W. L. (2000), *Pedagogical Agents. MIT Press.*

Johnson W. L., Rickel J. W., Lester J. C. (2000). Animated Pedagogical Agents: Face-to-Face Interaction in Interactive Learning Environments. *International Journal of Artificial Intelligence in Education 2000,* 11, 47-78.

Lester, J. C., Converse, S. A. Kahler, S. E. Barlow, S. T., Stone, B. A. and Bhogal, R. (1997a). The Persona Effect: Affective Impact of Animated Pedagogical Agents. *Proceedings of CHI '97 (Human Factors in Computing Systems,* pp. 359-366.

Lester J.C., Stone B.A., and Stelling G.D. (1999). Lifelike Pedagogical Agents for Mixed-Initiative Problem Solving in Constructivist Learning Environments. *User Modelling and User-Adapted Interaction,* 9, (1-2), pp. 1-44, 1999.

Sampson D., Karagiannidis, C., Kinshuk, (2002). Personalised Learning: Educational, Technological and Standardisation Perspective. *Interactive Educational Multimedia,* number 4 (April 2002), pp. 24-39

Shaw E., Ganeshan R., Johnson W. L., and Millar D. (1999). Building a Case for Agent-Assisted Learning as a Catalyst for Curriculum Reform in Medical Education. *International Conference on Artificial Intelligence in Education*.

Zuckerman I., and Albreciit D. W. (2001). Predictive Statistical Models for User Modeling. *User Modeling and User Adapted Interaction,* 11.

Acknowledgements

Part of the R&D work reported in this paper is carried out in the context of the "EP.E.N.D.Y.SH" project, partially funded by the Greek Ministry of Education.