Modelling Empathy in Social Robotic Companions

Iolanda Leite[1], André Pereira[1], Ginevra Castellano[2], Samuel Mascarenhas[1], Carlos Martinho[1], and Ana Paiva[1]

[1] INESC-ID and Instituto Superior Técnico, Technical University of Lisbon, Portugal
[2] HCI Centre, School of Electronic, Electrical and Computer Engineering
University of Birmingham, United Kingdom
iolanda.leite@ist.utl.pt

Abstract. Empathy can be broadly defined as the ability to understand and respond appropriately to the affective states of others. In this paper, we present a scenario where a social robot acts as a chess companion for children, and describe our current efforts towards endowing such robot with empathic capabilities. A multimodal framework for modeling some of the user's affective states that combines visual and task-related features is presented. Using this model of the user, we personalise the learning environment by adapting the robot's empathic responses to the particular preferences of the child who is interacting with the robot. We also describe a preliminary study conducted in this scenario.

Keywords: social robots, empathy, affective user modeling, adaptive interaction.

1 Introduction

Emotions shape how we appraise the environment and our attitudes towards other people and tasks. Isen and Reeve [21] reported that positive affect influences people's intrinsic motivation, enjoyment and performance of enjoyable tasks. Conversely, while experiencing negative emotions, people tend to judge their environment and future events more negatively [24]. A growing amount of research studies the relations between emotions and learning. For example, students in a positive mood are more willing to dedicate more effort in a problem solving task [18]. In a study using an intelligent tutoring system, Craig et al. [12] found that learning correlates negatively with boredom and positively with flow [13]. Other studies showed that a positive student-teacher relationship increases student's trust, cooperation and motivation during the learning process, which can ultimately improve the learning experience [35].

In the last few years, there has been a growing interest in developing animated pedagogical agents with the aim of guiding the students through the learning environment while trying to keep them in a positive affective state [28]. Several studies suggest that pedagogical agents positively affect the way students perceive the learning experience due to their non-verbal behaviours [28], physical

L. Ardissono and T. Kuflik (Eds.): UMAP 2011 Workshops, LNCS 7138, pp. 135–147, 2012.

appearance or voice [17]. More recently, the ability to recognise and respond to the student's affective state has also been considered a very important characteristic of pedagogical agents [37,15,10]. In humans, the capacity of understanding and responding appropriately to the affective states of others is commonly designated as empathy [19]. It is argued that empathy facilitates the creation and development of social relationships [2]. A positive student-teacher relationship increases student's trust, cooperation and motivation during the learning process. For these reasons, empathy is often linked with effective teaching [11].

In this paper, we summarise our efforts towards the development of a social robot with empathic capabilities that acts as a chess companion for children. By endowing the robot with empathic capabilities, we expect to improve the relationship established between children and the robot, which can ultimately lead them to improve their chess abilities. To behave empathically, our robot needs to (1) *model* the child's affective states and (2) *adapt* its affective and prosocial behaviour in response to the affective states of the child. The model of the user is created by combining visual features (e.g., smiles and gaze) with contextual information of the chess game. Based on the user model and on previous user's responses to empathic reactions of the robot, the robot selects the most appropriate empathic response for a particular user. Such empathic responses can include facial expressions, encouraging comments or prosocial actions that can help the user to change the course of the game.

This paper is organised as follows. We start by presenting some related work on multimodal affect detection and agents that respond to user's affect in learning environments. Afterwards, we describe our application scenario and the proposed framework to model the user's affective state and react in an empathic manner. Then, we report a preliminary study to evaluate the influence of empathic behaviours on user's perception of the robotic companion in this scenario. Finally, we discuss the results of this study and present some future work directions.

2 Related Work

Research on artificial companions has recently started to address the issue of designing systems for the automatic recognition of scenario-dependent, spontaneous affect-related states emerging during the interaction with a robot or virtual agent. Examples include the system by Kapoor et al. [22], which can automatically predict frustration of users interacting with a learning companion based on multimodal non-verbal cues such as facial expressions, head movement, posture, skin conductance and mouse pressure data, and the work by Nakano and Ishii [29], who proposed an approach to estimate the user's conversational engagement with a conversational agent based on analysis of gaze patterns. A multimodal approach was also used by D'Mello et al. [16], who proposed a combination of features from different sensors (dialogue, body postures and facial features) to improve the classification accuracy of the affective states of users interacting with an Intelligent Tutoring System (Autotutor). In this domain there has been an increasing attention towards systems utilising contextual information to improve the affect recognition performance [10]. Kapoor and Picard [23] proposed an approach for the detection of interest in a learning environment by combining non-verbal cues and information about the learner's task (i.e., level of difficulty and state of the game). In our previous work on automatic detection of engagement with the iCat robot, we showed that a combination of task and social interaction-based features allows for the highest recognition rate to be achieved [9].

While many efforts are being made to detect user's affective and motivational states, another branch of research addresses the challenge of how affect-aware agents should react to those states, and in which ways empathic responses improve the interaction. For example, Robison et al. [30] studied the impact of empathic feedback on students interacting with a virtual agent in a narrative-centred learning environment. They found that affective and motivational states such as flow, delight and boredom seem to be more vulnerable to the quality of feedback given than other emotions such as frustration. In the same line of research, Pour and colleagues [1] investigated the relation between positive, neutral and negative feedback responses provided by an Intelligent Tutoring System and the learner's affective state (obtained both by self-report measures and physiological signals). The results suggest that the feedback from the system influenced the learner's affective state. Arroyo et al. [3] found that non-invasive interventions that help students reflect on their progress are beneficial to students learning, attitudes towards learning and towards the tutoring software. They also discovered that there is a higher probability of students to become re-engaged with the system immediately after seeing an intervention. In a scenario with a social robot, Saerbeck and colleagues [32] investigated the effects of the robot's social supportive behaviour on student's learning performance and motivation. The results indicate that simple manipulations in terms of the robot's supportiveness (while maintaining the same learning content) increased student's motivation and scores in a language test.

3 Towards an Empathic Chess Companion

Our application scenario consists of an iCat robot [36] that plays chess with children using an electronic chessboard (see Fig. 1). Chess can be considered an educational game, as it helps children develop their memory and problem solving skills [20]. The iCat provides feedback on the children's moves by employing facial expressions determined by the robot's affective state. The affective state is determined by an anticipatory mechanism that creates an expectation on the children's upcoming move, and then based on the evaluation of the actual move played by the child, an affective state is elicited, resulting in a different facial expression in the robot. In addition to these momentary emotions, the state of the game also influences the robot's mood, reflected in its overall positive or negative facial expression (for more details on the robot's affective system see [25]). While waiting for the user's move, the robot also exhibits "life-like" idle behaviours such as looking at the chessboard, blinking or looking to sides.

A previous study using this scenario showed that the affective behaviour expressed by the iCat increased user's understanding of the chess game [25].

Fig. 1. Child playing with the iCat

However, in another study, after repeated interactions with the robot, children started realising that the robot's behaviour did not take into account their own affective state [26]. The results of this study suggested that social presence decreased over time, especially in terms of perceived affective and behavioural interdependence. These dimensions refer to the extent to which users consider that the behaviour and affective state of the robot is influenced by their own behaviour and affective state. We believe that, by endowing the iCat with empathic capabilities, some of the limitations found in previous studies can be surpassed, and the relationship established between the user and the robot may improve. As described earlier, empathy requires the ability of understanding the user's affective state and responding accordingly. Thus, in the next section, we describe our current research in these two processes of empathy.

3.1 Modelling the User's Affective State

We aim to endow the robot with the ability to infer scenario-dependent user affective states, and specifically affective states related to the game and the social interaction with the robot. However, off-line analysis of videos recorded during several interactions between children and the iCat showed that children display prototypical emotional expressions only occasionally [8]. Therefore, we decided to measure the user's *valence of feeling*, a state that provides information about the overall feeling that the user is experiencing. We use this dimensional description of affect to describe the degree to which the user's affect is positive or negative [31]. In the remaining of this section we present an initial multimodal approach for the detection of the user's valence of feeling based on selected visual and task-related features.

A Multimodal Approach. People can infer the state of others by direct observation of emotional cues such as facial expressions and/or by inferring the contextual nature of the other's situation. For example, we can feel joy by simply watching someone we care about smiling, but also by knowing that he or she received a job promotion. Both ways are sufficient to generate empathic responses. Some theorists even argue that empathy cannot be detached from context, describing empathy as "an ability to understand a person's emotional reactions in consort with the context" [14].

In the related work section we presented some examples of affect detection systems that employed contextual information of the task to improve the recognition of user's affective states. We also follow a multimodal approach, by combining a set of visual and task-related features. The set of visual features considered for this scenario contain users' non-verbal behaviours exhibited while interacting with the iCat, such as *looking at the iCat, looking at the chessboard, looking elsewhere, smiling, mouth fidget, hand on mouth, scratching face or head, raising eyebrows, approaching* or *moving away.*

As for the task-related features, these are obtained using perspective taking [5], that is, the iCat uses the same mechanisms that it employs to generate its own behaviour, but in the user's point of view. Concretely, the robot analyses the game in the perspective of its opponent using the same chess evaluation function that it uses to evaluate and generate its own moves. An important component to achieve this is the Game Engine (see Fig. 2), which contains a chess heuristic function capable of evaluating the chess board not only in the perspective of the robot but also from the user's perspective. The Game Engine returns the *game state*, a value that indicates the advantage/disadvantage of the user in the game. The *game evolution* is a feature derived from the game state that indicates how much the user improved (or not) since the last move, by subtracting the current game state by the previous game state value. Another feature retrieved from the Game Engine indicates whether there were any *captured pieces*, either by the user or by the iCat, in the previous move. Note that the user might have a different perspective of the game (e.g., for not being able to correctly evaluate the state of the game). Nevertheless, the same can happen when humans try to put themselves in another person's place.

A Framework for Detecting Valence. In our previous work, we showed the key role of a subset of user's non-verbal behaviours and task-related features in

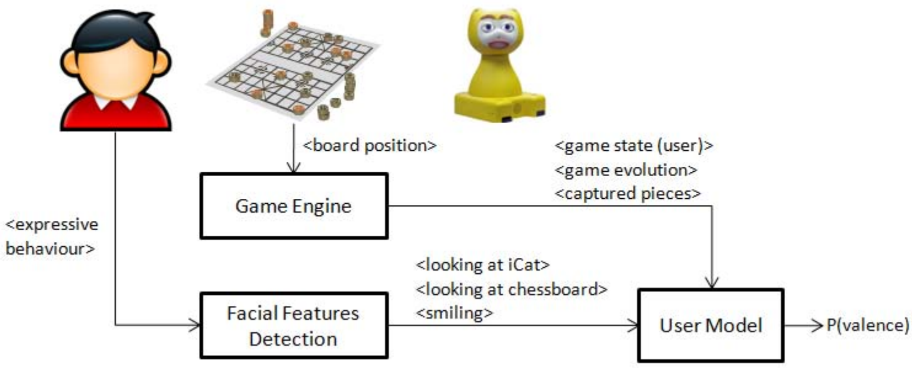

Fig. 2. Features and components relevant to the User Model

Table 1. Significant user's non-verbal behaviours and task-related features for the discrimination of user's valence

	User's non-verbal behaviours	Task-related features
Positive feeling	Looking at the iCat, smiling	User winning, user improving in the game, user capturing a piece
Negative feeling	Looking at the chess-board, looking elsewhere	User losing, user getting worse in the game, iCat capturing a piece

the discrimination of the user's valence [7,8]. Table 1 summarises the identified key non-verbal and task-related features for this scenario.

As illustrated in Fig. 2, the system automatically extracts visual features (based on face tracking performed with the FaceAPI software[1]) and contextual features of the game using the Game Engine. These features are combined in the User Model, which returns the probability of the user being in a positive or negative feeling. The User Model is based on Support Vector Machines trained using the Inter-ACT corpus, an affective and contextually rich multimodal video corpus including affective expressions of children playing chess with the iCat [33]. The affect probability values are then be used by the robot's action selection mechanism, which selects the most appropriate empathic response at a certain moment.

3.2 Adaptive Empathic Responses

After modelling the affective state of the user, the iCat should be able to select the empathic responses that are most effective to keep the user in a positive affective state. The empathic responses of the robot comprise two main steps. First, the iCat reacts emotionally to the move played by the user by displaying a facial expression that reflects that move. For example, if the user played a good move and the robot was not expecting that (considering the history of previous moves played by the user), the iCat displays a facial expression that reflects "surprise". In addition to the facial expressions displayed after every user's move, when the user is experiencing a negative feeling, the iCat also employs an empathic strategy. The empathic strategies implemented in this scenario were inspired on the literature of empathy and prosocial behaviour (e.g., the work of Cooper et al. [11]), and were translated into game-related actions. Some of these strategies were proven to be successful in a previous study that investigated the influence of empathic behaviours on people's perceptions of a social robot [27]. Currently, the empathic strategies implemented in the robot are the following:

1. Encouraging comments, for example, "don't be sad, I believe you can still recover your disadvantage".

[1] http://www.seeingmachines.com/product/faceapi/

2. Scaffolding, by providing feedback on the user's last move and, if the move is not good, let the user play again.
3. Suggesting a good move for the user to play in his or her next turn.
4. Intentionally playing a bad move, for example, playing a move that allows the user to capture an important piece of the robot.

But how should the robot decide, among this set of possible empathic strategies, which one is more appropriate for a certain user? We follow an approach based on Reinforcement Learning (RL), so that the robot can learn by trial and error the best strategies for a particular user, and adapt its empathic behaviour accordingly. For example, consider a situation where the user is experiencing a negative feeling for loosing an importance piece in the game and the iCat responds with an encouraging comment. If the user's valence changes from negative to positive, then encouraging behaviours will become part of the user's preferences in that particular situation. The RL policy for selecting the empathic strategies is based on an algorithm that solves the multi-armed bandit problem [4]. In this problem, a gambler chooses a slot machine to play with the overall goal of maximising the sum of rewards he receives from a sequence of pulls. Every time the gambler chooses a machine, he receives a reward that can be positive, negative or zero. Formalising this problem to our application scenario, when the robot detects that the user is experiencing a negative affective state, it selects an empathic strategy from a set of strategies S_i. When a strategy i is selected, it yields a reward R_i. The policy for selecting the next strategy to employ is thus based on the history of previous rewards:

if NOT all strategies initialised **then**
  for $i = 1 \to num\_strategies$ **do**
    $strategy = strategies[i]$
  end for
else
  loop
    Select strategy i that maximises: $\overline{x_i} + \sqrt{\frac{2 \ln n}{n_i}}$
    Where:
    $\overline{x_i}$ = average reward obtained from strategy i
    n_i = number of times strategy i was selected so far
    n = overall number of strategies selected so far
  end loop
end if

According to this algorithm, first each strategy is selected once. After that, the strategy with the maximum average rewards is selected. This algorithm falls into the category of the online learning algorithms. These algorithms attempt to optimise the learning performance without requiring a large amount of trials. This is appropriate for our scenario, since the robot should adapt its behaviour to the preferences of the user from an early stage of the interaction. The main goal of the robot is thus to maximise the user's positive valence, and so the rewards that classify each strategy must reflect that. In other words, to classify the success (or failure) of a certain empathic strategy, the reward function captures how much the user's positive feeling improved after the robot employed such strategy:

$$R_i = P(positive feeling)_{\text{after}} - P(positive feeling)_{\text{before}}$$

In this function, *P(positive feeling)*$_{before}$ represents the probability of the user being in a positive feeling before the iCat employing an empathic strategy, and *P(positive feeling)*$_{after}$ is the new probability of the user's valence, measured a while after the iCat employed S_i. This means that, a few seconds after the iCat employs the strategy, the Affect Detection system is called again to update the User Model, which will contain P(positive feeling)$_{\text{after}}$. With this information, the reward value for the employed strategy can be calculated and, consequently, the corresponding average reward value is updated.

This policy addresses the "exploration versus exploitation" dilemma, a problem defined as "the search for a balance between exploring the environment to find profitable actions while taking the empirically best action as often as possible" [4]. This means that the robot will select the empathic strategy with the maximum value of average rewards, but with an additional decay factor that enables the selection (and subsequent reward update) of other strategies after some iterations. Note that, at the beginning, the algorithm selects each strategy once in order to obtain an initial reward value for each strategy. As the same users are expected to interact with the robot for several games, the preferences for a particular user are updated even over different interaction sessions.

4 Experiment

To preliminary test our approach for modelling empathy in a social robotic companion, we performed an exploratory study using the scenario described earlier. The objectives of this exploratory study were twofold. First, we wanted to evaluate how users react to the different empathic strategies displayed by the robot. Second, our goal was to test the viability of using adaptive strategies in real time, while collecting data that could be used in the future to improve the model.

4.1 Procedure

In this study, we had three different conditions concerning the robot's behaviour: *neutral*, *random_empathic* and *adaptive_empathic*. In the *neutral* condition, the robot does not display any empathic strategies to the user. Instead, it only provides feedback through facial expressions, but in a competitive way, that is, when the user has disadvantage in the game, the iCat will be happier. In the *random_empathic* condition, when the robot detects that the user's is experiencing a negative feeling, it selects an empathic strategy randomly. Finally, in the *adaptive_empathic* condition, the empathic strategies are chosen according to the adaptive algorithm described in the previous section.

The three conditions of the robot's behaviour were evaluated with 40 students of an elementary school where children have two hours per week of chess lessons. From the 40 participants, 19 were male and 21 were female, with ages between 8 and 10 years old. Each child only interacted with one version of the robot. Therefore, 14 subjects played with the *adaptive_empathic* version of the robot and the two other conditions, *neutral* and *random_empathic*, had 13 subjects each.

During the study, each participant was alone in a room with the experimenter and played a chess exercise with the iCat for approximately 10 to 15 minutes[2]. After that time, depending on the state of the game, the iCat either gave up (if it was in disadvantage) or proposed a draw (if the child was loosing or if none of the players had advantage), arguing that it had to play with another user. Participants were then oriented to another part of the room where they filled in a questionnaire. Finally, they were guided to another room and were interviewed by another experimenter. The questionnaire used in this study measured user's **engagement with the robot**, **help** and **self-validation**. For each measure, users expressed their opinion on a set of assertions using a 5-point Likert scale, where 1 meant "totally disagree" and 5 meant "totally agree".

Engagement is a metric that has been extensively used both in human-robot and human-agent interaction and has been defined from several perspectives [34,6]. For example, Sidner et al. defined engagement as "the process by which two (or more) participants establish, maintain and end their perceived connection" [34]. The questionnaire items regarding engagement in our study are based on the questions used by Sidner et al. to evaluate users' responses towards a robot capable of using social capabilities to attract users' attention.

We also decided to measure two of the friendship dimensions that were used in a previous empathy study with the iCat [27], **help** and **self-validation**. We decided not to use the complete friendship questionnaire since we wanted to keep the overall number of questions in the questionnaire reasonable to 8-10 years old children. For this reason, we selected only help and self-validation, because we believe that these dimensions may be affected by the empathic strategies employed by the robot in this scenario. According to this questionnaire, help measures how the robot provided guidance and other forms of aid to the users, whereas self-validation measures the degree of reassuring, encouraging and helping the other to maintain a positive self-image.

4.2 Results

For each measure (**engagement**, **help** and **self-validation**), we considered the average ratings of all the questionnaire items associated to it. We verified that the distribution of the data was not normal by applying Kolmogorov–Smirnov test, so non-parametric tests were applied.

The Kruskal–Wallis test only yielded significant results in terms of **help**, which indicates that empathy affected how helpful participants found the robot

[2] The exercise was suggested by the chess instructor, so that the difficulty was appropriate for all the students.

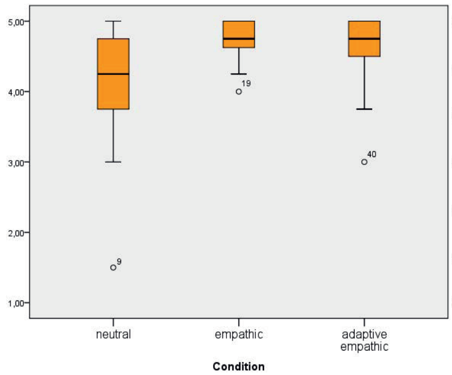

(a) Engagement

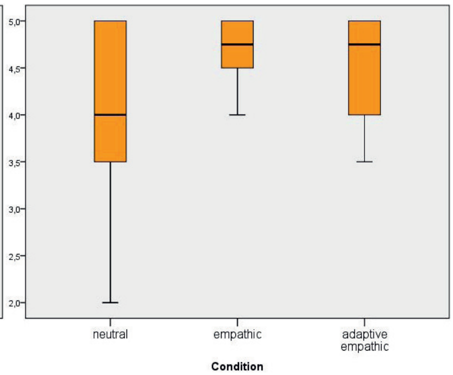

(b) Self-Validation

Fig. 3. Boxplot charts for Engagement and Self-Validation in the three conditions

(*H(2)*=10.53, p<0.05). Mann–Whitney tests were used to follow up this finding. A Bonferroni correction was applied and so all effects are reported at a 0.025 level of significance, since two further tests were performed. These results show that participants in the *random_empathic* (U=33, r=-0.53) and *adaptive_empathic* (U=34, r=-0.56) conditions significantly found the robot more helpful than participants in the *neutral* group.

Regarding **engagement** and **self-validation**, the Kruskal–Wallis test did not yield significant results. However, the distributions of the empathic conditions for these two measures are very similar, especially in terms of medians and upper quartiles (see Boxplot charts of Figure 3). Therefore, we decided to combine the two empathic distributions and perform a new analysis in terms of engagement and self-validation. Mann-Whitney tests were run to assess the significance of the differences observed across the *neutral* group and another group that covers the two empathic conditions (*all_empathic*)). In this case, the *all_empathic* group differed significantly from the *neutral* group, both for engagement (U=108, r=-0.32, p<0.05) and self-validation (U=106, r=-0.32, p<0.05), with a medium effect size. These results suggest that participants from the two empathic conditions significantly found the robot more engaging than participants from the control condition, and also that they gave higher ratings in terms of self-validation.

We also tested whether the end-game result had any effects on the measures of the study using Kruskal-Wallis test. Considering the subjects from all conditions (N=40), the end game result affected how engaging participants found the robot (*H(2)*=8.17, p<0.05), with participants who lost the game rating the interaction as less engaging. Similarly, participants who lost the game against the iCat also found the robot less helpful (*H(2)*=8.65, p<0.05). However, the end game result did not have any significant influence on self-validation. As opposed to what happened with engagement and help, participants found that the robot encouraged them during the game, despite the final result. Since the empathic behaviours of the robot are more often employed when the user is loosing the game, this suggests that the iCat's empathic behaviour had a positive effect on users.

5 Discussion and Future Work

In this paper, we described our work towards endowing a social robot with empathic capabilities. A multimodal system for predicting and modeling some of the children's affective states in real time was trained using a corpus of videos previously collected in another experiments using this scenario. With this model of the user, we personalise the learning environment by adapting the robot's empathic responses to the particular preferences of the child interacting with the robot. We have described the procedure and main results of a preliminary study conducted with this scenario.

The results of the study suggest that children perceived the robot in both empathic versions as more engaging, helpful and also provided higher ratings in terms of self-validation. However, no significant results were found between the two different empathic conditions, *empathic* and *adaptive_empathic*. This may have happened for two main reasons. First, we believe that a long-term study is needed to properly evaluate the adaptive empathic model. During the whole interaction, each child played between 10 to 20 moves, which was also the number of times that the robot assessed the children's affective state. The amount of times that an empathic strategy is employed is only a small part of this (when the user is experiencing a negative valence). For this reason, users might not have had enough time to realise that adaptation was taking place or, also very likely, the inherent noise of the reward function did not allow for the robot to really capture the users' preferences in such a short number of interactions. Another possible reason concerns with the time between employing the empathic strategy and measuring the user's affect again for calculating the reward function. A fixed time for every strategy might not be appropriate, as children might need more time to react to certain strategies and less time for others. To address this issue, we are planning to analyse the videos collected during this study to determine an average time, for each strategy, that users take to react to it.

Acknowledgements. This research was supported by EU 7th Framework Program (FP7/2007-2013) under grant agreement n^{o} 215554 and 3 PhD scholarships (SFRHBD/41358/2007, SFRH/BD/41585/2007, SFRHBD/62174/2009) granted by FCT. The authors are solely responsible for the content of this publication.

References

1. Aghaei Pour, P., Hussain, M.S., AlZoubi, O., D'Mello, S., Calvo, R.A.: The Impact of System Feedback on Learners' Affective and Physiological States. In: Aleven, V., Kay, J., Mostow, J. (eds.) ITS 2010. LNCS, vol. 6094, pp. 264–273. Springer, Heidelberg (2010)

2. Anderson, C., Keltner, D.: The role of empathy in the formation and maintenance of social bonds. Behavioral and Brain Sciences 25(01), 21–22 (2002)
3. Arroyo, I., Ferguson, K., Johns, J., Dragon, T., Meheranian, H., Fisher, D., Barto, A., Mahadevan, S., Woolf, B.P.: Repairing disengagement with non-invasive interventions. In: Proc. of AIED 2007, pp. 195–202. IOS Press, Amsterdam (2007)
4. Auer, P., Cesa-Bianchi, N., Fischer, P.: Finite-time analysis of the multiarmed bandit problem. Machine Learning 47(2), 235–256 (2002)
5. Batson, C.D., Early, S., Salvarani, G.: Perspective taking: Imagining how another feels versus imaging how you would feel. Personality and Social Psychology Bulletin 23(7), 751 (1997)
6. Bickmore, T., Schulman, D., Yin, L.: Maintaining engagement in long-term interventions with relational agents. Applied Artificial Intelligence 24(6), 648–666 (2010)
7. Castellano, G., Leite, I., Pereira, A., Martinho, C., Paiva, A., McOwan, P.W.: It's all in the game: Towards an affect sensitive and context aware game companion. In: Proceedings of ACII 2009, pp. 29–36. IEEE (2009)
8. Castellano, G., Leite, I., Pereira, A., Martinho, C., Paiva, A., McOwan, P.W.: Affect recognition for interactive companions: Challenges and design in real-world scenarios. Journal on Multimodal User Interfaces 3(1-2), 89–98 (2010)
9. Castellano, G., Pereira, A., Leite, I., Paiva, A., McOwan, P.W.: Detecting user engagement with a robot companion using task and social interaction-based features. In: Proceedings of ICMI 2009, pp. 119–126. ACM Press (2009)
10. Conati, C., Maclaren, H.: Empirically building and evaluating a probabilistic model of user affect. User Modeling and User-Adapted Interaction 19(3), 267–303 (2009)
11. Cooper, B., Brna, P., Martins, A.: Effective Affective in Intelligent Systems: Building on Evidence of Empathy in Teaching and Learning. In: Paiva, A.M. (ed.) IWAI 1999. LNCS, vol. 1814, pp. 21–34. Springer, Heidelberg (2000)
12. Craig, S., Graesser, A., Sullins, J., Gholson, B.: Affect and learning: an exploratory look into the role of affect in learning with AutoTutor. Learning, Media and Technology 29(3), 241–250 (2004)
13. Csikszentmihalyi, M.: Flow: The psychology of optimal experience. Harper Perennial, New York (1991)
14. Deutsch, F., Madle, R.A.: Empathy: Historic and current conceptualizations, measurement, and a cognitive theoretical perspective. Human Development 18(4), 267–287 (1975)
15. D'Mello, S., Craig, S., Fike, K., Graesser, A.: Responding to learners' cognitive-affective states with supportive and shakeup dialogues. In: Human-Computer Interaction. Ambient, Ubiquitous and Intelligent Interaction, pp, 595–604 (2009)
16. D'Mello, S., Picard, R.W., Graesser, A.: Toward an affect-sensitive autotutor. IEEE Intelligent Systems, 53–61 (2007)
17. Domagk, S.: Do pedagogical agents facilitate learner motivation and learning outcomes?: The role of the appeal of agent's appearance and voice. Journal of Media Psychology 22(2), 84–97 (2010)
18. Efklides, A., Petkaki, C.: Effects of mood on students' metacognitive experiences. Learning and Instruction 15(5), 415–431 (2005)
19. Hoffman, M.L.: Empathy and moral development: Implications for caring and justice. Cambridge Univ. Press (2001)
20. Horgan, D.D., Morgan, D.: Chess expertise in children. Applied Cognitive Psychology 4(2), 109–128 (1990)

21. Isen, A.M., Reeve, J.: The influence of positive affect on intrinsic and extrinsic motivation: Facilitating enjoyment of play, responsible work behavior, and self-control. Motivation and Emotion 29(4), 295–323 (2005)
22. Kapoor, A., Burleson, W., Picard, R.W.: Automatic prediction of frustration. International Journal of Human-Computer Studies 65(8), 724–736 (2007)
23. Kapoor, A., Picard, R.W.: Multimodal affect recognition in learning environments. In: ACM International Conference on Multimedia, pp. 677–682 (2005)
24. Keltner, D., Ellsworth, P.C., Edwards, K.: Beyond simple pessimism: Effects of sadness and anger on social perception. Journal of Personality and Social Psychology 64, 740–740 (1993)
25. Leite, I., Martinho, C., Pereira, A., Paiva, A.: iCat: an affective game buddy based on anticipatory mechanisms. In: Proceedings of AAMAS 2008, pp. 1229–1232. IFAAMAS (2008)
26. Leite, I., Martinho, C., Pereira, A., Paiva, A.: As Time goes by: Long-term evaluation of social presence in robotic companions. In: Proceedings of RO-MAN 2009, pp. 669–674. IEEE (2009)
27. Leite, I., Mascarenhas, S., Pereira, A., Martinho, C., Prada, R., Paiva, A.: "Why Can't We Be Friends?" An Empathic Game Companion for Long-Term Interaction. In: Safonova, A. (ed.) IVA 2010. LNCS, vol. 6356, pp. 315–321. Springer, Heidelberg (2010)
28. Lester, J.C., Converse, S.A., Kahler, S.E., Barlow, S.T., Stone, B.A., Bhogal, R.S.: The persona effect: affective impact of animated pedagogical agents. In: Proceedings of CHI 1997, pp. 359–366. ACM, NY (1997)
29. Nakano, Y.I., Ishii, R.: Estimating user's engagement from eye-gaze behaviors in human-agent conversations. In: Proceeding of IUI 2010, pp. 139–148. ACM, New York (2010)
30. Robison, J., McQuiggan, S., Lester, J.: Evaluating the consequences of affective feedback in intelligent tutoring systems. In: Proceedings of ACII 2009. IEEE (2009)
31. Russell, J.A.: A circumplex model of affect. Journal of Personality and Social Psychology 39, 1161–1178 (1980)
32. Saerbeck, M., Schut, T., Bartneck, C., Janse, M.: Expressive robots in education: varying the degree of social supportive behavior of a robotic tutor. In: Proceedings of CHI 2010, pp. 1613–1622. ACM (2010)
33. Sanghvi, J., Castellano, G., Leite, I., Pereira, A., McOwan, P.W., Paiva, A.: Automatic analysis of affective postures and body motion to detect engagement with a game companion. In: ACM/IEEE International Conference on Human-Robot Interaction, Lausanne, Switzerland (2011)
34. Sidner, C.L., Kidd, C.D., Lee, C., Lesh, N.: Where to look: a study of human-robot engagement. In: Proc. of IUI 2004, pp. 78–84. ACM (2004)
35. Tiberius, R.G.: The why of teacher/student relationships. Essays on Teaching Excellence 31 (2003)
36. van Breemen, A., Yan, X., Meerbeek, B.: iCat: an animated user-interface robot with personality. In: Proc. of AAMAS 2005, pp. 143–144. ACM (2005)
37. Woolf, B., Burleson, W., Arroyo, I., Dragon, T., Cooper, D., Picard, R.: Affect-aware tutors: recognising and responding to student affect. International Journal of Learning Technology 4(3), 129–164 (2009)