

Modeling Parallel and Reactive Empathy in Virtual Agents: An Inductive Approach

Scott W. McQuiggan¹, Jennifer L. Robison¹, Robert Phillips¹², and James C. Lester¹

¹Department of Computer Science
North Carolina State University
Raleigh, NC 27695 USA

²Applied Research Associates
8540 Colonnade Center Dr., Suite 301
Raleigh, NC 27615-3052 USA

{swmcquig, jlrobiso, rphilli, lester}@ncsu.edu

ABSTRACT

Humans continuously assess one another's situational context, modify their own affective state, and then respond based on these outcomes through empathetic expression. Virtual agents should be capable of similarly empathizing with users in interactive environments. A key challenge posed by empathetic reasoning in virtual agents is determining whether to respond with parallel or reactive empathy. *Parallel empathy* refers to mere replication of another's affective state, whereas *reactive empathy* exhibits greater cognitive awareness and may lead to incongruent emotional responses (i.e., emotions different from the recipient's and perhaps intended to alter negative affect). Because empathy is not yet sufficiently well understood, it is unclear as to which type of empathy is most effective and under what circumstances they should be applied. Devising empirically informed models of empathy from observations of "empathy in action" may lead to virtual agents that can accurately respond in social situations.

This paper proposes a unified inductive framework for modeling parallel and reactive empathy. First, in training sessions, a trainer guides a virtual agent through a series of problem-solving tasks in a learning environment and encounters empathetic characters. The proposed inductive architecture tracks situational data including actions, visited locations, intentions, and the trainer's physiological responses to generate models of empathy. Empathy models are used to drive runtime situation-appropriate empathetic behaviors by selecting suitable parallel or reactive empathetic expressions. An empirical evaluation of the approach in an interactive learning environment suggests that the induced empathy models can accurately assess social contexts and generate appropriate empathetic responses for virtual agent control.

Categories and Subject Descriptors

H.5.1 [Multimedia Information Systems]: Artificial, augmented, and virtual realities; Evaluation/methodology.

General Terms

Algorithms, Design, Experimentation, Human Factors.

Keywords

Intelligent Virtual Agents, Affective Reasoning, Empathy, Machine Learning, Human-Computer Interaction.

Cite as: Modeling Parallel and Reactive Empathy in Virtual Agents, McQuiggan, et al., *Proc. of 7th Int. Conf. on Autonomous Agents and Multiagent Systems (AAMAS 2008)*, Padgham, Parkes, Müller and Parsons (eds.), May, 12-16., 2008, Estoril, Portugal, pp. 167-174. Copyright © 2008, International Foundation for Autonomous Agents and Multiagent Systems (www.ifaamas.org). All rights reserved.

1. INTRODUCTION

Recent years have seen significant advances in cognitive and behavioral models for virtual agents [1, 4, 13, 16, 24, 28]. Affective modeling in virtual agents [13] is the subject of increasing attention because of its role in motivating users, supporting them through stressful tasks, and increasing users' abilities to recognize and regulate emotions. Agents with affective capabilities can form social relations with users to motivate them [2], reduce stress levels in job interview role-players [26], teach children to deal with frustration [3], and elicit emotional responses to teach children how to deal with bullying [24]. A critical component of each of these social interactions is the ability to use empathy.

Empathy is the expression of emotion based on another's situation and not merely one's own [8, 14, 15]. Its expression can demonstrate that the target's (the recipient of empathetic expression) feelings are understood or shared. In the case of *parallel empathy*, an individual exhibits an emotion similar to that of the target [8]. This is typically based on an understanding of the target's situation and shows the empathizer's ability to identify with the target. *Reactive empathy*, in contrast, focuses on the target's affective state, in addition to her situation [8]. Reactive empathizers will display emotions that are different from the target's, often in order to alter or enhance the target's own affective state. This type of empathy is focused on the target whereas parallel empathy is more self-oriented. As such, reactive empathy can be viewed as a higher level of empathetic behavior.

This paper presents CARE, an inductive framework for learning empirically grounded models of empathy from observations of human-agent social interactions. During training sessions, the inductive framework monitors user situation data (actions, visited locations, and intentions), affective states, physiological responses, and other characteristics (e.g., age and gender) while a training user (the *target*) directs her virtual agent to perform a sequence of tasks in a virtual learning environment. Meanwhile, virtual characters (the *empathizers*) empathetically respond to target (user) situations with either parallel or reactive empathy. During conversations with virtual characters, users are able to evaluate character empathetic responses. The model induction stage of the data-driven approach learns models of empathy from "good" examples. At runtime, induced models drive situation-appropriate empathetic behaviors (parallel or reactive) in virtual agents as they interact with actual users. Because the preference for particular types of empathy in a given situation may vary considerably between users, the capability of adapting virtual agents' empathetic behaviors to individual differences holds much appeal.

The paper is structured as follows. Section 2 provides a brief review of related work and background on empathetic virtual agents. Section 3 presents the inductive empathy modeling framework, CARE. Section 4 introduces the interactive learning environment (CRYSTAL ISLAND), and Section 5 describes the experiment and presents the results. Concluding remarks and directions for future work follow in Section 6.

2. EMPATHETIC VIRTUAL AGENTS

The social-psychological study of empathy is a relatively recent development [8]. Defined as an awareness of another's affective state that generates emotions in the empathizer that reflect more than their own situation, empathy is formalized in a tripartite tableau: an antecedent, an assessment and an empathetic outcome [8]. The *antecedent* captures the affective and situational context of the target (the recipient of empathetic expression) that is then assessed by the empathizer. This *assessment* yields an empathetic *outcome* that can be cognitive (e.g., greater awareness of the target's situation) or affective (e.g., flow, frustration, delight, etc.). As noted above, two types of empathetic outcomes can be distinguished: parallel outcomes and reactive outcomes. In parallel outcomes, the empathizer mimics the affective state of the target. For example, the empathizer may become fearful when the target is afraid. In reactive outcomes, empathizers exhibit a higher cognitive awareness of the situation and react with empathetic behaviors that do not necessarily match those of the target's affective state. For example, empathizers may become encouraging when the target is frustrated with the problem-solving tasks.

Empathy represents one of several constructs that have been investigated in affective computing [25], a field which has seen the appearance of computational models of emotion [10, 13], automated techniques for recognizing user affect [12, 18, 25], and explorations of the role of affect in learning [3, 5, 6, 11, 17, 19, 20, 21, 23].

Recent years have seen a growing interest in empathetic reasoning in virtual agents. Bickmore [2] showed how embodied agents can form social relationships with users by employing empathy and thereby improving the users' motivation. The Empathic Companion [26] tracks a user's bio-signals (e.g., GSR and heart rate) in order to assess the effect of empathetic interventions within stressful job interview scenarios. It was found that empathetic feedback successfully reduced the user's arousal. Bursleson [3] studied empathetic interventions in frustrating learning environments and explored their effect on meta-affective-strategy learning. In a similar vein, Paiva et al. [24] studied the requirements for eliciting user empathy and showed that psychological proximity (e.g., gender, shared qualities) is important for generating empathy. Finally, McQuiggan and Lester [22] extracted empathetic behavior protocols mimicking human empathetic behavior. While significant advances have been made in modeling empathy, previous work has not addressed the problem of parallel and reactive empathetic reasoning.

3. EMPATHY INDUCTIVE FRAMEWORK

3.1 Architecture

The CARE inductive framework operates in two modes: *empathy model induction*, in which the architecture interacts with a student trainer to gather data, and *runtime operation*, in which it monitors student situations, student affective states, student characteristics

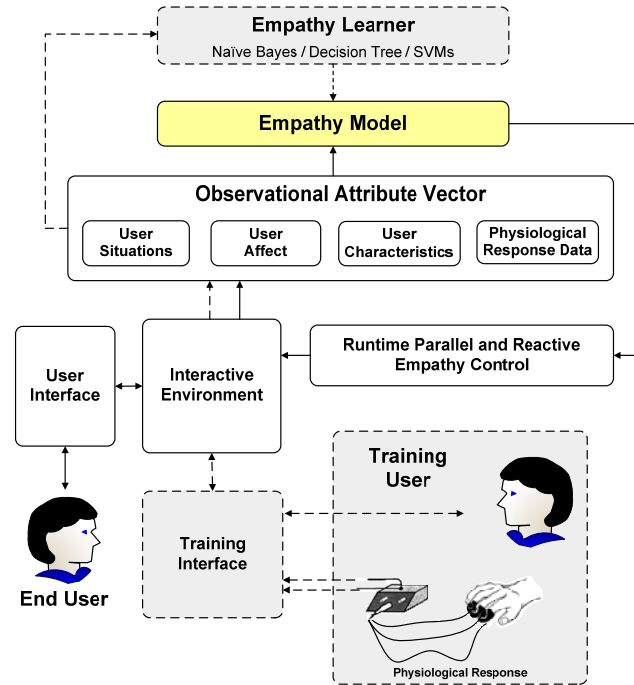


Figure 1. The CARE Inductive Framework.

(i.e., age, gender, etc), and student physiological response to respond empathetically.

- Empathy Model Induction.** During model induction (depicted in Figure 1 with dashed arcs), the framework acquires training data and learns models of empathy from training users interacting with the learning environment and the characters that inhabit it. The student trainers are outfitted with biofeedback equipment that monitors their heart rate and galvanic skin response. Biofeedback signals are recorded in training logs via the interactive environment, which also records an event stream produced by the students' behaviors in the environment. The event stream includes information about actions, locations and intentions of the student in the 3D interactive environment. Together, the biofeedback signals and the corresponding elements in the event stream are assembled in temporal order into the observational attribute vector. After training sessions (typically involving multiple training users) are complete, the empathy learner induces models from the observed situational data, reported student affect, other student characteristics, and physiological data. Characters interacting with students respond to student situations with either parallel or reactive empathy. The type of empathy used by each character serving as the class label.
- Runtime Operation.** During runtime operation (represented in Figure 1 with solid arcs), which is the mode employed when students interact with fielded learning environments, the induced models inform virtual agent behavior control by predicting whether parallel or reactive empathy is best suited for the current social situation and learning context. The learning environment again tracks all activities in the world and monitors the same observable attributes reported to the empathy learner during empathy model induction. The

induced model is used by the empathy controller to assess the situation in order to determine if empathy is called for, and if so, whether parallel or reactive empathy would be most effective for the given situation.

3.2 Training Data Acquisition

Accurately modeling parallel and reactive empathy requires a representation of the situational context that satisfies two requirements: it must be sufficiently rich to support empathetic assessment, and it must be encoded with features that are readily observable at runtime. Affect is fundamentally a product of cognitive processes in which the user appraises the relationship between herself and her environment [13, 27]. Similarly, empathy draws heavily on appraisal of the situation at hand in addition to user affect. Thus, empathy models should take into account environmental information, user affective states, and user physiological information. For interactive learning environments, empathy models can leverage knowledge of problem and learning task structures as well as the state of the student in the environment to assess if empathy is called for, and if so in what capacity (parallel or reactive). The proposed inductive empathy modeling framework therefore employs an expressive representation of all activities in the learning environment, including those controlled by users and the interactive system, by encoding them in an *observational attribute vector*, which is used in both the model induction mode and in the runtime mode of operation. During model induction, the observational attribute vector is passed to the empathy learner for model generation; during runtime operation, the attribute vector is monitored by an empathy model-enhanced runtime component that utilizes empathy predictions to inform effective virtual character response. The observable attribute vector represents four interrelated categories of features for making decisions:

- **User situations.** Empathy models can observe users' actions in the world and their relationship to achieving particular goals. Empathy models also have access to auxiliary information about the interactions, e.g., any artifacts manipulated, as well as locations visited and the characters with which users have interacted.
- **User affect.** Empathy models can observe user affective states. In this case, virtual characters query users to report on emotional states. One can also imagine the use of affect recognizers [5, 12, 18] to automatically detect user affect without requiring the user to self-report emotions.
- **Other user characteristics.** Empathy models can account for other user characteristics such as age and gender. In this case, the models also consider users' empathetic nature (measured by the Interpersonal Reactivity Index [7]) and their goal orientation (measured by Elliot and McGregor's goal inventory [9]).
- **User physiological response data.** Empathy models can observe changes in physiological response via biofeedback apparatus. For instance, the CARE inductive framework monitors changes in user heart rate and galvanic skin response.

During model induction, a continuous stream of physiological data is collected and logged approximately 30 times per second. In addition, an instance of the observational attribute vector is logged every time a significant event occurs, yielding, on average, several hundred vector instances each minute. An event is

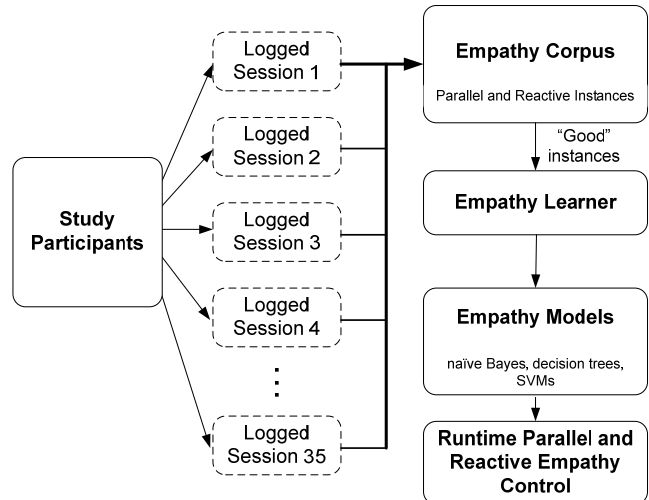


Figure 2. Empathy Data Acquisition.

considered to be significant when a user's manipulation of the environment causes one or more features of the observational attribute vector to take on new values. At runtime, the same features are continuously monitored by the environment.

3.3 Learning Empathy Models

Naïve Bayes, decision tree, and support vector machine classifiers can be used to generate predictive empathy models. Bayes classification approaches produce probability tables that can be implemented in runtime systems and used to continually update probabilities for predicting parallel or reactive empathetic responses. Decision trees provide interpretable rules that support runtime empathy modeling. At runtime, models can monitor the condition of the attributes in the rules to determine when conditions are met for diagnosing whether an empathetic response should be parallel or reactive in nature. Support vector machines (SVMs) are yet another classification method that is particularly effective at handling high-dimensional data. SVMs search for hyperplanes that linearly separate data into classes (parallel or reactive).

These classification techniques are particularly useful for inducing models with large multidimensional data, such as the data gathered in the user study described below. Because it is unclear precisely which runtime variables are likely to be the most predictive, naïve Bayes and decision tree modeling provide useful analyses that can inform more expressive machine learning techniques (e.g., Bayesian networks) that also leverage domain experts' knowledge. We have used the WEKA machine learning toolkit [29] to analyze naïve Bayes, decision tree, and SVM approaches for generating models of empathy that predict whether virtual agents should respond to student situations with parallel or reactive expressions.

The induction procedure uses the following steps to generate models of empathy (Figure 2):

- **Data construction.** Each training log is first translated into a full observational attribute vector. For example, temporal windows monitoring trends in physiological response data are computed and combined with situational data, user affect, and other user characteristics.

- **Data cleaning.** After data is converted into observation attribute vector format, a dataset is generated containing only instances of virtual character empathetic expression. Each instance in the vector then contains the observational attributes, a user rating evaluating the characters' empathetic response (see Section 5.1 for details on how ratings were collected), and the class label (parallel or reactive).
- **Model induction.** Once the dataset is prepared, it is passed to the learning systems. Each dataset is loaded into the WEKA machine learning toolkit [29], a naïve Bayes classifier, decision tree, and SVM were learned, and tenfold cross-validation analyses were run on the resulting models.

The following section presents CRYSTAL ISLAND, the narrative learning environment test bed utilized in the virtual agents empathy study described in Section 5.

4. CRYSTAL ISLAND

The virtual agents empathy study was conducted in a narrative-centered inquiry-based learning environment, CRYSTAL ISLAND (Figure 3). This environment is designed to teach the domains of microbiology and genetics to middle school students. It features a science mystery set on a recently discovered volcanic island where a research station has been established to study the unique flora and fauna. The user plays the protagonist, Alex, attempting to discover the genetic makeup of the chickens whose eggs are carrying an unidentified infectious disease at the research station. The story opens by introducing the student to the island and the members of the research team for which her father serves as the lead scientist. As members of the research team fall ill (Figure 4), it is her task to discover the cause and the specific source of the outbreak. She is free to explore the world and interact with other characters while forming questions, generating hypotheses, collecting data, and testing her hypotheses. Throughout the mystery, she can walk around the island and visit the infirmary, the lab, the dining hall, and the living quarters of each member of the team. She can pick up and manipulate objects, and she can talk with characters to gather clues about the source of the disease. In the course of her adventure, she must gather enough evidence to correctly choose which breeds of chickens need to be banned from the island.

The virtual world of CRYSTAL ISLAND, the semi-autonomous characters that inhabit it, and the user interface were implemented with Valve Software's Source™ engine, the 3D game platform for Half-Life 2. The Source engine also provides much of the low-level (reactive) character behavior control. The character behaviors and artifacts in the storyworld are the subject of continued work as evidenced by the current study. Students direct their character through CRYSTAL ISLAND by using the keyboard controls (WASD) and mouse movements (typical game controls).

To illustrate the behavior of CRYSTAL ISLAND, consider the following situation. Suppose a student has been interacting within the storyworld and learning about infectious diseases, genetic crosses and related topics. In the course of having members of her research team become ill, she has learned that an infectious disease is an illness that can be transmitted from one organism to another. As she concludes her introduction to infectious diseases, she learns from the camp nurse that the mystery illness seems to be coming from eggs laid by certain chickens and that the source of the disease must be identified. During problem-solving activities the student is introduced to several characters. Some



Figure 3. Overview of CRYSTAL ISLAND.



Figure 4. The user, Alex, with Jin, the camp nurse on CRYSTAL ISLAND.

characters are able to help identify which eggs come from which chickens while other characters, with a scientific background, are able to provide helpful genetics information. The student discovers through a series of tests that the bad eggs seem to be coming from chickens with white-feathers. The student then learns that this is a codominant trait and determines that any chicken containing the allele for white-feathers must be banned from the island immediately to stop the spread of the disease. The student reports her findings back to the camp nurse.

5. EMPATHY CORPUS COLLECTION

To empirically investigate “empathy in action,” a study was conducted with subjects interacting with virtual agents. The subjects of the study consisted of 35 college students ranging in age from 21 to 60 ($M = 24.4$, $SD = 6.41$) including 9 females and 26 males. Among these students, 60% were Asian ($n = 21$), approximately 37% were Caucasian ($n = 13$) and one participant chose not to respond.

5.1 Experimental Design

Participants entered the experiment room where they completed informed consent documentation. They were randomly assigned to either the control condition or the empathy condition and were seated in front of a laptop computer. They were then given an overview of the experiment agenda, and they completed the pre-experiment questionnaires including the demographics survey, the interpersonal reactivity index survey [7], and the goal orientation survey [9].

The interpersonal reactivity index [7] includes 28 items that measure subjects' empathetic nature by asking them to rate the degree to which each statement describes them. These items are assessed on a 5-point Likert scale (0 - does not describe me well to 4 - describes me very well). The IRI is divided into four subscale measurements quantifying the following components of empathy: fantasy, perspective taking, empathetic concern, and personal distress [7]. The achievement goals questionnaire [9] measures four achievement goal constructs (mastery-approach, performance-approach, mastery-avoidance, and performance-avoidance goals). Subjects indicate the extent to which each statement is true of them on a 7-point Likert scale (1 - not at all true of me to 7 - very true of me).

Upon completing the pre-experiment questionnaires, participants were instructed to review CRYSTAL ISLAND instruction materials. These materials consisted of the backstory and task description, the character overviews, the map of the island, the control sheet, and definition sheet of the self-report emotions. Participants were then further briefed on the controls via a presentation summarizing the task and explaining each control in detail.

Participants were given 35 minutes to solve the mystery. Solving the mystery consisted of completing 15 goals including learning about various diseases, compiling the symptoms of the sickened researchers, testing a variety of possible sources, and reporting the solution (cause and source) back to the camp nurse.

Six characters (Audrey, Elise, Jin, Quentin, Robert, and Teresa), each play distinct roles in the CRYSTAL ISLAND environment. When subjects decided to interact with the characters, the following schema was used to direct subject-character interactions and virtual character empathetic responses:

1. The virtual character queries the subject for a self-reported affective state. The subject is presented with a dialog box asking the question, "Hi Alex, how are you feeling?" The subject may respond by selecting one of the 10 available emotions (*anger, anxiety, boredom, confusion, delight, excitement, fear, flow, frustration, and sadness*).
2. The virtual character then responds to the subject's reported affective state with an empathetic response. The empathetic response is determined from the subject-reported emotion and the character's empathizer type, i.e., whether the character is a reactive empathizer or a parallel empathizer. Empathetic responses are short, consisting of 1 to 2 sentences. Parallel responses consist of the character expressing the same emotion as the user through text responses (i.e., *I feel frustrated by this as well*); alternatively, reactive responses demonstrate advanced cognitive

processing on the character's part by providing responses designed to be more motivating, revealing the character's desire for the user to be in a positive emotional state (i.e., *I really feel that you can get it!*).

3. A follow-up dialog box is then presented to the subject asking her to respond with the prompt, "...and you respond". The subject is able to choose from 4 Likert-scaled responses designed to evaluate the appropriateness and effectiveness of the virtual character's empathetic response. Subjects can issue responses ranging from (1) "That does not help me at all." to (4) "Thanks, I feel a lot better!"
4. The virtual character responds with a one-word quip (e.g., "Thanks", or "Great!") directed towards the subject's evaluation response (Step 3). In addition, the virtual character provides narrative and problem-solving information.
5. The virtual character then asks the subject how she feels one final time before concluding the interaction. The subject is presented a dialog box similar to the one described in Step 1 without the character greeting. Here, the character prompts the subject with, "How are you feeling now?"
6. Finally, the virtual character again empathetically responds to subject-reported affective states in the same manner as described in Step 2.

Immediately after solving the science mystery of CRYSTAL ISLAND (or after 35 minutes of elapsed interaction time for subjects who had not solved the mystery), subjects completed the post-experiment questionnaire. This researcher-designed questionnaire assessed perceptions of individual CRYSTAL ISLAND characters. The results of this instrument are outside the scope of this discussion.

5.2 Results

In Step 3 of the user-agent interaction schema presented above, subjects evaluated virtual character responses as part of the "conversation" with the character. The distribution of empathetic responses (parallel and reactive) and the associated ratings are detailed in Table 1. The first two rows show the number of empathetic responses that were found to be appropriate by subjects (i.e., the instances of parallel and reactive empathetic behaviors that were given an evaluative rating of 3 or 4), while the next two rows indicate empathetic responses that were found to be inappropriate by subjects (i.e., the instances of parallel and reactive empathetic behaviors that were given an evaluative rating of 1 or 2).

Naïve Bayes, decision tree, and support vector machine (SVM) models were induced from data collected in the training sessions described above using the WEKA machine learning toolkit [29]. All models were constructed using a tenfold cross-validation scheme for producing training and testing datasets. In this scheme, data is decomposed into ten equal partitions, nine of which are used for training and one used for testing. The equal parts are swapped between training and testing sets until each partition has been used for both training and testing. Tenfold cross-validation is widely used for obtaining an acceptable estimate of error (Witten and Frank, 2005).

Table 1. Distribution of empathetic responses and associated subject evaluation ratings (4 = high to 1 = low).

		Empathetic Responses		
		Parallel	Reactive	Total
Evaluative Rating	4	1200	1305	2505
	3	1663	1600	3263
	2	1558	1498	3056
	1	1033	1033	2066
	Total	5454	5436	10890

Two distinct datasets were used. The first dataset was comprised only of empathetic responses receiving high ratings of 4 ($n = 2505$) from subjects during conversations with virtual characters. The second dataset was comprised of empathetic responses receiving a rating of either a 3 or 4 ($n = 5768$). Within each dataset, several versions of models were learned from the various types of observational attributes (situation data, user affect reports, user characteristics, and physiological response data). Table 2 provides results of all induced models and baselines.

A baseline measure determines the most frequent class label (in this case reactive empathy) and predicts all empathetic responses to call for reactive responses. For the dataset containing only empathetic responses rated level 4 (*Highest-Rated*), reactive empathy accounted for 52% of the instances. Reactive empathy occurred in 50% of the instances in the dataset containing empathetic responses rated as level 3 or 4 by subjects (*Favorably-Rated*).

All induced models outperformed baseline models. The improvement of induced models over baselines is statistically significant. For instance, the least accurate induced model from the Highest-Rated dataset is the naïve Bayes model (73% accuracy), which was learned from situation attributes only. The accuracy of this naïve Bayes model was statistically significant from baseline accuracy, $\chi^2(1, N = 2505) = 238.94, p < 0.001$. Also, in the Highest-Rated dataset, decision tree performance was statistically significant from naïve Bayes performance, except for the models induced from situation attributes only. For instance, in the models learned from situation attributes and user affective states, decision tree accuracy (87%) was statistically significant from naïve Bayes accuracy (74%), $\chi^2(1, N = 2505) = 136.42, p < 0.001$.

Similar results appear in models learned from the *Favorably-Rated* dataset. All induced models again outperformed the baseline (reactive empathy) in which 50% of instances were reactive responses. The worst performing induced model's accuracy (64%) was the naïve Bayes model learned from situation attributes only. This model was statistically significant from baseline accuracy, $\chi^2(1, N = 3263) = 131.04, p < 0.001$. Decision tree model accuracies were statistically significant from their corresponding naïve Bayes models for each category of data in the

Table 2. Empathy model results by data type used for learning (left column bold items) and dataset used. Highest-Rated refers to the dataset containing empathetic responses rated a 4. Favorably-Rated refers to the dataset containing empathy responses rated either a 3 or 4.

Induced Empathy Models	Dataset	
	Highest-Rated	Favorably-Rated
Baseline (Reactive)	0.52	0.50
Situation Attributes Only		
Naïve Bayes	0.73	0.64
Decision Tree	0.75	0.70
SVM	0.75	0.67
Situation + Affect		
Naïve Bayes	0.74	0.67
Decision Tree	0.87	0.88
SVM	0.83	0.80
Situation + Affect + Bio		
Naïve Bayes	0.75	0.68
Decision Tree	0.95	0.96
SVM	0.84	0.81
Situation + Affect + Characteristics		
Naïve Bayes	0.74	0.70
Decision Tree	0.98	0.98
SVM	0.97	0.92
Situation + Affect + Bio + Characteristics		
Naïve Bayes	0.75	0.71
Decision Tree	0.98	0.98
SVM	0.97	0.92

Favorably-Rated dataset. For instance, the accuracy of the decision tree model learned from situation data only (70%) was statistically significant from the naïve Bayes model (64%) induced from the same data, $\chi^2(1, N = 2505) = 26.65, p < 0.001$.

5.3 Design Implications

The study found that models of empathy induced from knowledge of the user's situation and the user's affective state can effectively determine which type of empathy is most appropriate for interactions requiring empathetic expression. The results suggest that designers of empathetic virtual agents should consider both parallel and reactive empathy in virtual agent architectures. The study has the following implications, each of which is discussed below:

- Empathetic response modeling should be integrated with other virtual agent response functionalities.
- Individual differences of users should be accounted for when determining empathetic responses.

- Empathy models induced from “good” examples (examples rated highly by subjects) should improve the quality of interaction with virtual agents.

Approximately one-half of the empathetic responses appearing in the user study received high marks from subjects (ratings of 3 or 4) leaving the other half as “bad” examples of empathetic responses. While we have not fully explored which attributes of the interaction correspond to low ratings, it seems reasonable to conclude that in instances where a parallel response received a low rating, it would not be improved by replacing the response with an instance of reactive empathy. To deal with this, virtual agent architectures should combine empathy models with other response strategies. For example, cases in which subjects reported feelings of frustration or confusion and subsequently rated empathetic responses poorly might be better addressed not by emotional responses, but by directed content feedback, e.g., by addressing the obstacle that is the source of user frustration or providing hints that may relieve user confusion. Certainly, emotional empathetic responses are not always appropriate and should be combined with a broad range of agent response strategies.

The models of empathy described in the previous section account not only for knowledge of user situation and user emotion, but also other demographic information such as gender, age, user empathetic nature, and goal orientation. This information is typically accounted for subconsciously in human-human interaction but often discounted in human-agent interaction. Agents that understand who users are, male or female, young or old, mastery- or performance-oriented, may be able to more effectively determine how to empathetically respond to user emotional situations.

Finally, models of empathy that account for both parallel and reactive empathy should lead to more effective human-agent interactions. Since subjects rated each empathetic response on a Likert scale, models were able to be induced solely from “good” examples, instances receiving ratings of 3 or 4. While future investigations will consider the effects of empathy models on learning and task performance, the addition more flexible empathetic abilities yields an immediate improvement in an agents’ abilities to respond appropriately to social situations.

5.4 Limitations

The results of the study are affected by the virtual characters that interacted empathetically with participants. First, it is possible that the gender, narrative role, and pedagogical role of the empathetic characters may not generalize to other characters and across domains. Further investigation is required to assess the effect of character persona on perception of character empathy. Second, the population participating in this study is a small group of college students studying computer science. Different empathy models may be appropriate for different demographic segments. To determine the generalizability of empathy models, additional user studies are required. The procedure used in the study utilized only empathetic responses. It may be the case, particularly for instances of empathetic behavior that were not favorably evaluated, that empathy was inappropriate and that other types of responses, such as those providing pedagogical assistance, or other content would be more effective.

6. CONCLUSION

Recent advances in affective reasoning demonstrate the important role that emotion plays in cognitive accounts of social interaction and suggest that it should therefore play an equally important role in virtual agents. Because empathy is a natural extension of the appraisal process and appears continuously in human-human interaction, it is important to endow virtual agents with the capability to respond with the parallel and reactive empathetic expressions that are most appropriate for the user, her situation and her affective state.

This paper has presented CARE, an inductive approach to learning empathy models that accounts for both parallel and reactive empathetic expression. The data-driven approach centers on the observation of “empathy in action” and acknowledges the psychological understanding of empathetic assessment and appraisal processes by including appropriate information in model induction. Such data include situational contexts (e.g., user actions, visited locations, goals), user affect and affective responses (as measured through physiological changes), and user demographics. While previous work has either focused solely on parallel empathy or not distinguished between the two forms, the inductive approach proposed here was evaluated in a user study designed to examine the effects of parallel and reactive empathy upon the recipient. By allowing virtual agents to employ parallel and reactive empathy that is appropriate for the social situation, it is hoped they will become more effective and engaging.

Several directions for future work appear promising. First, it is important to create integrated models of empathetic behavior and automatic affect detection. This integration may support more natural interactions and enable virtual agents to infer user affective state more unobtrusively. Second, because empathetic preferences may vary across agent characteristics as well as user characteristics, empathy models should be extended to consider “agent demographics.” For example, the gender of an agent, in addition to the gender of the user, may significantly affect the manner in which empathy is most appropriately expressed. Third, empathy expression may benefit from enriched models of social context. For example, modeling the social roles played by virtual agents in a narrative, perhaps in conjunction with the role played by the user, could significantly increase the effectiveness of human-agent interactions. Lastly, it is necessary to investigate the impact of parallel and reactive empathetic expressions on users in various contexts in order to ensure the pedagogical effectiveness of such interventions.

7. ACKNOWLEDGMENTS

The authors wish to thank the members of the IntelliMedia Center for Intelligent Systems at North Carolina State University for their assistance in implementing CRYSTAL ISLAND, Omer Sturlovich and Pavel Turzo for use of their 3D model libraries, and Valve Software for access to the Source™ engine and SDK. This research was supported by the National Science Foundation under Grants REC-0632450 and CNS-0540523. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the authors and do not necessarily reflect the views of the National Science Foundation.

8. REFERENCES

- [1] André, E., and Müller, M. (2003) Learning affective behavior. In *Proceedings of the 10th International Conference on Human-Computer Interaction* (Heraklion,

- Crete, Greece, June 22-27). Lawrence Erlbaum, Mahwah, NJ, 512-516.
- [2] Bickmore, T. (2003) Relational Agents: Effecting Change through Human-Computer Relationships. *Ph.D. Thesis*, Massachusetts Institute of Technology, Cambridge, MA.
- [3] Bursleson, W. (2006) Affective learning companions: Strategies for empathetic agents with real-time multimodal affective sensing to foster meta-cognitive and meta-affective approaches to learning, motivation, and perseverance. *PhD thesis*, Massachusetts Institute of Technology, Cambridge, MA.
- [4] Cavazza, M., Charles, F., and Mead, S. (2002) Interacting with virtual characters in interactive storytelling. In *Proceedings of the 1st International Conference on Autonomous Agents and Multi-Agent Systems* (Bolonga, Italy, July 15-19). ACM Press, New York, NY. 318-325.
- [5] Conati, C. (2002) Probabilistic assessment of user's emotions in educational games. *Applied Artificial Intelligence*, 16:555-575.
- [6] D'Mello, S., Craig, S., Sullins, J. and Graesser, A. (2006). Predicting Affective States expressed through an Emote-Aloud Procedure from AutoTutor's Mixed-Initiative Dialogue. *International Journal of Artificial Intelligence in Education*. 16:3-28.
- [7] Davis, M. (1983) Measuring individual differences in empathy: Evidence for a multidimensional approach. *Journal of Personality and Social Psychology*, 44, 113-126.
- [8] Davis, M. (1994) *Empathy: A Social Psychological Approach*. Brown and Benchmark Publishers, Madison, WI.
- [9] Elliot, A., and McGregor, H. (2001) A 2 x 2 achievement goal framework. *Journal of Personality and Social Psychology*, 80, 3, 501-519.
- [10] Elliott, C. (1992) The affective reasoner: A process model of emotions in a multi-agent system. *PhD thesis*, Northwestern University.
- [11] Elliott, C., Rickel, J., and Lester, J. (1999) Lifelike pedagogical agents and affective computing: An exploratory synthesis. *Artificial Intelligence Today*, Lecture Notes in Artificial Intelligence, Vol. 1600, Wooldridge, M. and Veloso, M. (eds.), Springer-Verlag, Berlin, 195-212.
- [12] Graesser, A., McDaniel, B., Chipman, P., Witherspoon, A., D'Mello, S., and Gholson, B. (2006) Detection of emotions during learning with AutoTutor. *Twenty-eighth Annual Conference of the Cognitive Science Society*, 285-290.
- [13] Gratch, J., and Marsella, S. (2004) A domain-independent framework for modeling emotion. *Journal of Cognitive Systems Research*, 5(4):269-306.
- [14] Hoffman, M. (2000) *Empathy and Moral Development: Implications for Caring and Justice*. Cambridge University Press, Cambridge, UK.
- [15] Ickes, W. (1997) *Empathic Accuracy*. Guilford Press, New York, NY.
- [16] Johnson, W., and Rickel, J. (1998) Steve: An animated pedagogical agent for procedural training in virtual environments. *SIGART Bulletin* 8:16-21.
- [17] Johnson, W., and Rizzo, P. (2004) Politeness in tutoring dialogs: "Run the factory, that's what I'd do". In *Proceedings of the 7th International Conference on Intelligent Tutoring Systems* (Maceio, Alagoas, Brazil, August 30 – September 3). Springer-Verlag, New York, NY, 67-76.
- [18] Kapoor, A. and Picard, R. (2005) Multimodal Affect Recognition in Learning Environments. In *Proceedings of the 13th annual ACM International Conference on Multimedia* (Singapore, November 6-11). ACM Press, New York, NY. 677-682.
- [19] Kort, B., Reilly, R., & Picard, R. (2001). An affective model of interplay between emotions and learning: Reengineering educational pedagogy—building a learning companion. In *Proceedings IEEE International Conference on Advanced Learning Technology: Issues, Achievements and Challenges*. Madison, Wisconsin, 43-48.
- [20] Lester, J., Towns, S., Callaway, C., Voerman, J., and FitzGerald, P. (2000) Deictic and emotive communication in animated pedagogical agents. In *Embodied Conversational Agents*, Cassell, J., Prevost, S., Sullivan, J., and Churchill, E. (Eds.), MIT Press, Cambridge, 123-154.
- [21] Litman, D., & Forbes-Riley, K. (2004). Predicting student emotions in computer-human tutoring dialogues. In *Proceedings of the 42nd annual meeting of the association for computational linguistics*. East Stroudsburg, PA, 352-359.
- [22] McQuiggan, S., and Lester, J. (2006) Learning empathy: A data-driven framework for modeling empathetic companion agents. In *Proceedings of the 5th International Joint Conference on Autonomous Agents and Multi-Agent Systems*. Hakodate, Japan, 961-968.
- [23] McQuiggan, S., Mott, B., and Lester, J. (to appear) Modeling self-efficacy in intelligent tutoring systems: An inductive approach. *User Modeling and User-Adapted Interaction*.
- [24] Paiva, A., Dias, J., Sobral, D., Aylett, R., Woods, S., Hall, L., and Zoll, C. (2005) Learning by feeling: Evoking empathy with synthetic characters. *Applied Artificial Intelligence*, 19:235-266.
- [25] Picard, R. (1997) *Affective Computing*. MIT Press, Cambridge, MA.
- [26] Prendinger, H., and Ishizuka, M. (2005) The empathic companion: a character-based interface that addresses users' affective states. *Applied Artificial Intelligence*. 19:267-285.
- [27] Smith, C., and Lazarus, R. (1990) Emotion and adaptation. In Pervin (Ed.), *Handbook of Personality: theory & research* (pp. 609-637). Guilford Press, New York, NY.
- [28] Swartout, W., Gratch, J., Hill Jr., R., Hovy, E., Lindheim, R., Marsella, S., Rickel, J., and Traum, D. (2004) Simulation meets Hollywood: Integrating graphics, sound, story and character for immersive simulation. In Stock, O., and Zanczarno, M. (eds.), *Multimodal Intelligent Information Presentation*, Kluwer.
- [29] Witten, I., and Frank, E. (2005) *Data Mining: Practical Machine Learning Tools and Techniques*. 2nd Edition, Morgan Kaufman, San Francisco, CA