

Evaluating Levels of Emotional Contagion with an Embodied Conversational Agent

Özge Nilay Yalçın (oyalcin@sfu.ca)

Interactive Arts and Technology, 250 -13450 102 Avenue
Surrey, BC, V3T 0A3 CANADA

Steve DiPaola (sdipaola@sfu.ca)

Interactive Arts and Technology, 250 -13450 102 Avenue
Surrey, BC, V3T 0A3 CANADA

Abstract

This paper presents an embodied conversational agent framework as a controlled environment to test components of empathy. We implement levels of emotional contagion which includes mimicry and affective matching along with necessary communicational capabilities. We further demonstrate an examination of these foundational behaviors in isolation, to better understand the effect of each level on the perception of empathy in a social conversational scenario with a human actor. We report three studies where the agent shows levels of emotional contagion behavior during (1) the listening act in comparison with baseline backchanneling behavior (2) additional verbal response matching simple emotional storyline (3) the verbal response to the human actor performing complex emotional behaviors. Results revealed that both mimicry and affective matching behaviors were perceived as more empathic than the baseline listening behavior, where the difference between these behaviors was only significant when the agent verbally responded to complex emotional behaviors.

Empathy; Emotional Contagion; Mirroring; Affect Matching; Affective Computing; Social Interaction; Embodied Conversational Agents

Introduction

Empathy, as the capability to understand and react to the emotions of another (Iacoboni, 2011; Coplan & Goldie, 2011), is a complex behavior that arises from the interaction of these basic affective mechanisms with higher-level cognitive functions (de Waal & Preston, 2017). Emotional contagion is said to be the foundation of empathic capacity, as it includes innate and automatic synchronization of the motor and affective responses during an interaction (Hatfield, Cacioppo, & Rapson, 1994). Behaviors such as mimicry and affective matching are levels of the emotional contagion that results from the innate capability of resonating with the other during social interaction.

The literature suggests the sustained act of mimicry results in a feeling of the mimicked emotion and affective matching through muscular feedback (Hatfield et al., 1994; Hatfield, Bensman, Thornton, & Rapson, 2014), while categorizing both behaviors as emotional contagion. Others use affect matching as a highly connected but distinct phenomenon to the mimicry, pointing out the differences between the subjective quality of experience in the emotional contagion and the automatic matching of expressions in mimicry (Hess & Fischer, 2014). However, both ideas converge on the foundational role of mimicry and affect matching in empathic behavior. This notion is consistent with the Perception-Action-

Model (PAM) (Preston & De Waal, 2002) and the Russian Doll model of empathy (de Waal, 2007), which integrates the neuroscience studies on mirror neurons as a baseline for the hierarchical levels of empathy mechanisms. However, it is difficult to study the levels of emotional contagion in isolation.

Research efforts often rely on behavioral experiments, neuroscientific techniques (EEG, fMRI) and pathology studies conducted to understand the effects of emotional contagion during social interactions (Hess & Fischer, 2014; Hatfield et al., 2014). As an alternative, computational empathy studies have recently gained attention in a way to simulating the empathy mechanism within the agent and examining empathic responses of the users towards the agent (Paiva, Leite, Boukricha, & Wachsmuth, 2017; Yalçın & DiPaola, 2018). The perception of empathy in artificial agents is shown to increase the length of the interaction (Leite, Castellano, Pereira, Martinho, & Paiva, 2014), user performance (Partala & Surakka, 2004), user satisfaction (Prendinger, Mori, & Ishizuka, 2005), and lead to more trust (Brave, Nass, & Hutchinson, 2005). These findings suggest that equipping interactive systems with empathic capacity would not only improve our understanding of the interaction between cognitive and affective processes in the human mind but may also help us enhance our interaction between artificial systems.

In this work, we use the simulation approach to study empathic behavior in virtual agents and try to understand the differences between the levels of emotional contagion behavior and the perception of empathy during a conversation. We examine the basic emotional contagion capabilities in an embodied conversational agent (ECA) in order to evaluate the perception of empathy during mimicry and affect matching behaviors. We present an agent framework and implementation with necessary communicational capabilities as a baseline. In the following section, we will present our implementation for an ECA that incorporates different levels of emotional contagion as a foundation for empathic capacity. Next, we will demonstrate three experiments that examine the effect of these levels on the perception of empathy during a social interaction scenario with a human actor. Our approach and results show the potential of computational empathy studies as a reliable alternative to test mechanisms for empathic behavior in isolation.

Agent Behavior

Our empathy framework is implemented in an embodied conversational agent that is capable of responding to an emotional conversation with the user using verbal and non-verbal behaviors. Our socially situated 3D virtual character system can perform a set of behavioral acts and context-specific dialogue in response to the speech and video input received from the user (see (Yalçın, in press) for a detailed explanation of the framework). Inputs are gathered using a standard webcam and a microphone. We use the Smartbody behavior realizer (Thiebaut, Marsella, Marshall, & Kallmann, 2008), that can provide face and body gestures, gaze, and speech output for virtual characters. We use the standard Behavior Markup Language (BML) (Kopp et al., 2006) as the basis for the two way connection between the framework and the behavior realizer.

The implementation includes mimicry and affect matching behaviors as the foundational capabilities of empathy in combination with basic conversational capabilities such as backchanneling. In order to achieve this, our system incorporates a perceptual module, a behavior controller and a behavior generation module. The visual and verbal input from the user is processed through the perceptual module, reasoned within the behavior controller according to the selected empathy mechanism and prepared for a behavioral output in the behavior manager before being displayed in the ECA.

Low-level empathic behaviors, such as mimicry and affective matching require a fast response to the emotional stimuli presented by the interaction partner. The fundamental components of this first level of empathic behavior include the perception of emotion, representation of emotion and expressing emotion. This cycle is realized with Perceptual Module and Controller and Behavior Generation modules of our system.

Perceptual Module

The perceptual module is responsible for handling the input received from the user and creating internal representations of these inputs to be used by the controller. Currently, our system is capable of handling audio, video and textual inputs to be used in recognition systems. The audio input includes verbal signals from the user to be recognized as speech and pauses. The initiation, pauses and termination in the speech signal are used to provide information about the dialogue state as well as backchannel timing.

Emotion recognition is a sub-module within the perceptual module that is specialized for emotion recognition and fusion processes. Here, three types of modalities can be used for further processing using the first level of recognition from the perceptual module: facial emotion recognition, tone analysis and speech emotion recognition. During listening, emotion recognition is based on the facial gestures and tone analysis, which is derived from the video and speech inputs for immediate listening feedback. After the speech signal from the user ended, the complete utterance is also being processed

in speech emotion recognizer for emotion detection based on the textual output of the speech recognizer. Outputs from this sub-module are used by the behavior controller depending on the dialogue state as well as the selected empathy mechanisms.

Behavior Controller

The behavior controller module is a central unit in the framework which provides a link between inputs and the outputs. It decides which input channel or information to be used depending on the state of the conversation, required empathy mechanisms and the behavioral capabilities of the agent. It is also responsible for providing the information necessary to the behavior manager module to prepare verbal and non-verbal behavior. The Controller acts as a decision-making component, which determines behavioral choices concerning the percepts of the agent and its internal state. Currently, the behavior controller provides a link between the perception-action mechanisms as a key component in computational empathy (de Waal, 2007). During a conversation, the agent should decide which behavioral state it is in depending on the user input: listening, thinking, speaking or waiting. According to the state of the interaction (listening, speaking, thinking and waiting) and the current emotional value (arousal, valence and emotion category), the controller assigns the proper behavior categories to the behavior generation component.

If the user is speaking, the agent should be in the listening mode. Here, the agent is expected to provide proper backchanneling to the user as well as the emotional feedback depending on the empathy mechanisms. After the speech of the user is over, the speech signal should be sent to the dialogue manager component through the controller with the assigned emotional value. The agent will be in thinking mode during the processing of this input by the dialogue manager component. The prepared output sentence will then be sent back to the controller to be sent to the behavior manager which will prepare the output behaviors including face gestures, body gestures and the verbal response to be presented in the speaking mode. After the speech behavior of the agent is done, the waiting or idle mode will be activated until the user speaks again. This cycle can be interrupted via the controller at any stage.

Behavior Generation

The behavior generation module is responsible for preparing the output for the virtual character depending on the emotion, dialogue state and speech information received from the behavior controller. During listening behavior, this module is relatively passive in preparing behaviors. It uses the backchanneling signal to select an appropriate head nod for the agent and a facial expression. When these behaviors are sent and consumed by the behavior realizer, the behavior generation module receives a signal back that indicates the behavior was successfully generated by Smartbody system.

Method

In this paper, we used the simulation approach to study low-level empathic behavior in virtual agents to show the differences between the levels of emotional contagion behavior in the perception of empathy. We examined the effect of mimicry and affect matching behavior on perceived empathy during conversational interaction using three studies.

Participants

Participants for all three experiments were recruited using Amazon's Mechanical Turk platform and were paid for their participation to the study. Because we were focusing on the emotional expressions during verbal communication, we only included participants who had English as their first language. Additionally, users that participate with mobile devices and tablets were excluded to ensure a consistency in the display quality.

A total of 84 subjects participated in the studies. 36 participants with ages ranging from 20 to 60 ($M=37.6$, $SD=10.7$) completed the first study. 19 of the participants were male and 16 of them female, while 1 participant defined themselves as 'other'. 24 subjects participated in the second study with ages ranging between 21 and 64 ($M=36.17$, $SD=10.82$). 10 of the participants were female and 13 of them male, while 1 participant defined themselves as 'non-binary'. The last study included 24 participants with ages ranging between 23 and 59 ($M=37.82$, $SD=10.64$), 12 Male and 12 Female.

Procedure

Studies followed the same procedure, where the participants are asked to evaluate the recorded interaction between the agent and a human (see Figure 1). The interaction scenario consists of a student/participant expressing an emotional story to the agent. We have chosen three stories inspired by the work of Omdahl (Omdahl, 2014), that includes three basic emotion categories: anger, joy and sadness. Other basic emotions such as fear, surprise and disgust were not considered for this study due to the involvement of facial action units that controlled mouth movements during the expression of these emotions. Furthermore, we selected the emotions that would be consistent with the facial emotions, that would not provide an advantage to the affective matching over mimicry.

All of the experiments were deployed in Amazon's Mechanical Turk environment using scripts written in Python 3.6 with `psiturk` and `jpsych` libraries. Each of the studies takes about 10 minutes to complete. Participants were first shown a test video and were asked to answer two questions about the visual and verbal content of the video, to make sure they can hear and see the videos that are displayed. This was required for the workers to participate in the study.

Each participant is then displayed a short video clip of an interaction, where the agent and a student are shown in a video-conferencing scenario in different conditions (see Figure 1). During the interaction, the student in the video talks about an emotional story in one of three basic emotions: joy,

sadness and anger. After displaying each video, the participant is asked to report what the story in the video was about, and also the main emotion of the user and the virtual agent. This is done to make sure the participants are paying attention to the video clips. The participants then evaluated the perceived empathy of the agent towards the student. The perceived empathy of the agent is evaluated by using a modified version of the Toronto empathy questionnaire (Spreng, McKinnon, Mar, & Levine, 2009) which is a 16-item survey that originally is used as a self-report measure. Each item on the questionnaire are scored in a 5-item likert scale (Never = 0; Rarely = 1; Sometimes = 2; Often = 3; Always = 4), where half of the items are worded negatively. Scores are summed to derive total for the perceived empathy and can be varied between -32 to +32. Similar evaluations were suggested by Paiva and colleagues (Paiva et al., 2017), as a modification of Davis's Interpersonal Reactivity Index (Davis et al., 1980).



Figure 1: An image from the video chat between the student and the avatar. Here, the student (left) converses with the avatar.

We used repeated measures design, where each participant is shown all levels of agent behavior in emotional contagion. The type of the interaction study and the order of the conditions are counterbalanced accordingly.

Experiment Conditions

Experiment conditions include three distinct agent behaviors that signifies levels of emotional contagion mechanisms in the empathy framework.

The baseline behavior of the agent is the backchanneling behavior, which is activated depending on the pauses during the speech signal from the audio input component in the perceptual module. In the following subsections, we will provide a detailed examination of three different listening behaviors depending on the level of empathic behavior of the agent: backchannel only, mimicry with backchannel, affective matching with backchannel.

Backchanneling as baseline behavior Listener behavior in humans include backchannels such as head nods, fa-

cial feedback, short vocalizations or a combination of them (Yngve, 1970). These behaviors might show information about listener agreement, acknowledgment, turn-taking or attitude (Schroder et al., 2012; Cassell, Bickmore, Campbell, & Vilhjálmsón, 2000). Backchannel feedback can occur due to change in pitch, disfluency or loudness of the speech signal, as well as shifts in speaker’s posture, gaze and head movements (Maatman, Gratch, & Marsella, 2005). In our current implementation we included backchanneling based on the pauses during speech, which is a form of disfluency in the speech signal (Maatman et al., 2005). Information about pauses are extracted from the perceptual module and sent to the controller, which in turn is used to trigger backchanneling as head nods. More advanced methods of adding backchannel that are compatible with the valence of the interaction partner or adding specific facial expressions such as smile, would have interfere with the empathy mechanisms that we would like to test. Therefore, we omitted these behaviors from the baseline behavior.

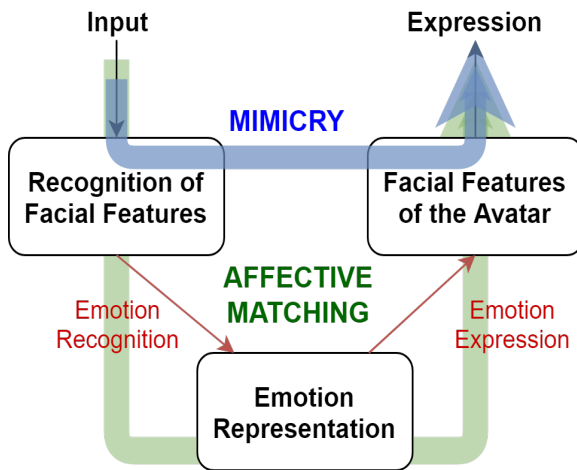


Figure 2: Two paths for emotional contagion. Basic emotional communication competence that results in low-level empathic capabilities of mimicry and affective matching by following distinct routes during the interaction process.

Mimicry Mechanism Mimicry is the lowest level of empathy behavior in our empathy model. It is achieved by a direct mapping between the gestures of the user to the gestures of the agent without being assigned to any type of emotional category. Facial mimicry behavior during listening is a result of mapping the perceived facial action units (AUs) extracted from the perceptual module, to the AUs of the embodied agent in the behavior generation module. The amount, duration and speed of these AUs match the perceived values of the interaction partner without any regulations. In order to avoid mimicking of the lip movements during the speaking of the user, we removed direct mapping of AU18 (lip puckerer), AU26 (jaw drop) and AU24 (lip pressor). As a side-effect of this modification, certain emotions that requires these AUs

(fear, surprise and disgust) were not properly expressed. In order to avoid bias, interactions that include these emotions were not used during the evaluation of the system for this study. However, this drawback should be noted for future studies.

After the listening cycle, the agent will sustain the mimicry behavior until it retrieves a response from the dialogue manager. The dialogue manager will then retrieve an emotionally neutral response, due to the lack of emotional representation that is needed to be acquired during the interaction.

Affective Matching Another type of low-level or affective empathy behavior is affective matching (de Waal & Preston, 2017). It is achieved by the emotion recognition and the emotion expression cycle that is connected through emotion representation. As it can be seen in Figure 2, the facial features are mapped to the representation of the basic emotion categories which in turn triggers the facial expressions of the agent that represents those emotions. The amount, duration and speed of these expressions depend directly on the values from the perceived emotions. In contrast to the mimicry behavior, this allows the agent to present and regulate emotions that are better perceived by the users. Moreover, excluded emotion categories in mimicry can be used without the disturbance of the AUs that control mouth muscles as explained in the previous section.

After the listening cycle, the agent will give an emotional feedback that reflects the overall emotion of the interaction partner until it retrieves a response from the dialogue manager. In the affective matching condition, the dialogue manager is able to use the representation of the interaction partner’s emotions to pick an emotional response. Without the effect of the higher level emotion regulation capabilities, the agent will pick a response that reflects the emotion of the interaction partner.

Study 1

In order to evaluate the perception of empathic behaviors we compared the listening behavior of the agent in backchannel, mimicry and affective matching conditions. For our study, we used within subjects design where three conditions of agent behavior are shown to the same subject for the evaluation. The conditions are baseline backchanneling behavior, mimicry with backchanneling and affective matching with backchanneling during only the listening act. We used three emotional stories told by the same person, which displays three different emotions as the main theme: joy, sadness and anger. Each video starts with a neutral remark, that is followed by the emotional story.

The experiment counterbalanced on the order of the type of interaction (backchannel, mimicry, affect matching), and the order of type of emotional story (angry, sad, happy). 36 (6x6) different conditions presented to subjects.

Evaluation In the evaluation of the first study, Mauchly’s Test of Sphericity indicated that sphericity had not been vio-

lated, $X^2(2) = 1.748, p = .417$. A one-way repeated measures ANOVA was conducted to compare the effect of (IV) level of emotional contagion behavior on (DV) the perception of empathy in backchanneling, mimicry, and affective matching conditions. The results showed that perceived empathy is significantly affected by the type of listening feedback $F(2, 70) = 16.721, p < .0001, 95\%CI$ (see Figure 3). Pairwise comparisons showed backchannel feedback only ($M = -5.47, SD = 12.45$) is perceived to have significantly lower empathy than both mimicry ($p < .001$) and affective matching ($p < .0001$). However, listening behavior with mimicry ($M = 5.16, SD = 10.64$) and affective matching ($M = 8.22, SD = 13.72$) did not have any significant difference ($p = .18$).




Figure 3: Results of our study showed significant differences in the perceived empathy levels between backchannel, mimicry and affective matching behavior (95%CI).

Study 2 and Study 3

Following up the first study, we further examined the effect of the verbal feedback produced by the dialogue manager in both mimicry and affect matching conditions. Our hypothesis is, due to the effect of emotional representation during the affect matching mechanism, the verbal response behavior of the avatar will be perceived as more empathic. However, this result might show difference when the interaction partner shows more complex emotions, where the context and information about the overall emotion representation is required to understand the semantics of the behavior. Therefore, we conduct two additional studies where one is focused on simple emotions and the other examines the effect of complex emotional behavior. For the following experiments, the participants were asked to evaluate the interaction stories, where the agent listens to different types of emotional stories told by the interaction partner and verbally reacts to it. As the first study showed significant differences over the baseline backchanneling behavior, the following studies did not compare the baseline behavior to emotional contagion.

In both conditions the listening behaviors of the agent will be the same as the first study, which showed no significant dif-

ference. The behavior of the agent between will differ from the first study in terms of verbal feedback during the conversational cycle. In the mimicry condition, the agent will produce an emotionally neutral feedback such as "I understand" or "I know what you mean" while sustaining the reflective facial expression of the interaction partner. In the affect matching condition, due to the additional information the dialogue manager will receive from the emotional representation of the interaction partner, the agent will produce an emotionally charged sentence. The emotional category of this sentence will be the same as the emotions of the interaction partner. For example, a happy story will trigger a happy remark such as "That sounds wonderful", an angry story will trigger a response such as "That is really frustrating", and a sad story will trigger a sad response such as "I am sorry to hear that".

The third experiment focused on more complex emotional stories, where the human actor will talk about two scenarios mentioning a dog and a plant. In the dog scenario, the actor will go through excitement, disgust, worry and happiness emotions while mentioning a story about their new pet dog. In the plant scenario, the actor will go through neutral, surprise, worry and happiness emotions while mentioning a story about their friend's plant. The listening behavior of the agent will be matching the emotions both in mimicry and affective matching conditions. Similar to the second study, mimicry condition will result in a generic verbal response from the agent while affective matching condition will give an emotionally charged feedback due to emotional representation.

The second experiment counterbalanced on the order of the type of interaction (mimicry, affect matching), and the order of the type of emotional story (angry, sad, happy). 12 (2x6) different conditions presented to subjects. The third experiment is also counterbalanced on the order of the type of interaction (mimicry, affect matching), and the order of the type of emotional story (dog, plant). 4 (2x2) different conditions presented to the subjects. Both experiments followed the same procedure as the first study.

Evaluations In the second study, one-way repeated measures ANOVA was conducted to compare the effect of (IV) level of emotional contagion behavior on (DV) the perception of empathy in mimicry, and affective matching conditions. The results showed that perceived empathy is not significantly different between mimicry ($M = 7.62, SD = 11.66$) and affect matching ($M = 9.5, SD = 8.03$) conditions $F(1, 23) = 1.030, p = .321$.

Following up these results, in the third study, one-way repeated measures ANOVA was conducted to compare the effect of (IV) level of emotional contagion behavior on (DV) the perception of empathy in mimicry, and affective matching conditions during the interaction with complex emotional behavior. The results showed that perceived empathy is significantly different between mimicry ($M = 0.75, SD = 10.45$) and affect matching ($M = 7.21, SD = 9.98$) conditions $F(1, 23) = 7.731, p = .011$ (see Figure 4).




Figure 4: Results of the third study showed significant differences in the perceived empathy levels between mimicry and affective matching behavior in complex emotional interaction (95%CI).

Discussion

The results of the studies showed a significant difference in the perception of empathy between the baseline backchanneling behavior and the emotional contagion behavior during the listening act. As expected, the perceived empathy was significantly higher in the emotional contagion behavior with respect to the baseline behavior in the first study. However, there was no significant difference between the different levels of empathic behavior (mimicry and affective matching) in this experiment. This result is a direct consequence of the similarity in the expressions of these two conditions during listening.

Even though mimicry and affect matching behaviors have important differences in terms of processing of input information, the real-time expressions of these behaviors during listening behavior show dramatic similarities. During the listening act, mimicry captures the facial expressions of the interaction partner and reflects them using the same facial muscles. In affective matching behavior, instead of copying the facial muscles, the system copies the emotions perceived from these facial expressions. As the emotions are expressed as a result of the facial muscles, these two behaviors are expected to show very similar expressions.

One advantage of affective matching that it allows the expression of emotions that are more suitable to the virtual agent, while any emotion will be expressed in terms of the virtual agent's repertoire instead of the expressions of the conversation partner. Moreover, affective matching allows processing of other input channels to conclude the emotion of the interaction partner, such as voice stress, the context of the speech and body expressions. However, the first and second studies only included simple emotions, and therefore such an effect was not present.

Another distinction between mimicry and affect matching conditions is present during the verbal response after the lis-

tening act is completed. This response is created by examining the overall emotion of the story told by the interaction partner. As mimicry behavior does not provide the representation of the emotions of the interaction partner, the virtual agent cannot generate a response that is aligned to that emotion. In contrast, the verbal response for the affective matching behavior can be generated from the emotion representation (see Figure 2 for a comparison of these two strategies). Study 2 and 3 are designed to show this distinction.

Interestingly, Study 2 did not show a significant difference between the mimicry and affect matching behaviors for simple emotional stories, where we see a significant difference in Study 3. In these studies, there are two main differences between mimicry and affective matching conditions: the content of the verbal response, and the facial emotions shown during the verbal response. In mimicry condition, the verbal response is generic where the affect matching condition generates an emotionally appropriate response. The facial expressions in mimicry response are sustained regardless of the overall emotions, where the affect matching condition generates facial expression based on the overall emotional representation for the whole story. The difference between the two studies was the emotional complexity of the overall story told by the interaction partner.

We argue that the mimicry response for the simple emotional stories in Study 2, did not show a significant difference on the perception of empathy due to the match between the overall emotion of the story and the sustained facial expression. Where in Study 3 the sustained emotion of the mimicry response was contrasting the overall emotions of the story, due to the complexity of the emotions presented by the interaction partner. We further examined the comments provided by the participants on how they perceived the behavior of the agent in response to the story told by the interaction partner. The comments of the participants in Study 2 showed that the mimicry condition is seen as "understanding" and "sympathy", where the affective matching behavior is seen as "concerned" and "empathy". In contrast, in Study 3, participant comments on the behavior of the virtual agent included descriptions such as "confused" and "indifferent", where the affect matching response was seen more as "attentive", "understanding" and "empathy". However, this distinction should be examined more systematically before reaching to a conclusion.

Overall, these results show that low-level emotional contagion behaviors of the agent during conversational interaction lead to an increased perception of empathy. Additionally, the results show that higher levels of emotional contagion behavior are perceived as more empathic behavior when the interaction includes more complex emotional behaviors. The proposed framework shows promise in providing a foundation to examine the perception of higher levels of empathic behavior during an interaction.

Conclusion and Future Work

Artificial systems provide means to test the empathy theories while allowing the manipulation of parameters in a controlled and isolated way. In this work, we proposed an embodied conversational agent framework to test empathy components and demonstrated three studies that evaluate the foundational empathy mechanisms along with basic communication behaviors. We found that during listening, mimicry and affective matching behaviors are perceived significantly more empathetic compared to backchannel behavior. We also found that the difference between the two levels of affective contagion only significant while the interaction involves complex emotional behaviors, where the context of the interaction is crucial for producing matching behavior. Our framework and the results of our initial study shows promising results that allows for easy integration and testing of higher level components of empathy. The suggested framework, study and evaluation methods shows the potential as a reliable alternative to test mechanisms for empathic behavior in isolation.

Our contributions were to provide a framework, implement the baseline behavior for real-time interaction with a highly realistic conversational avatar, and provide the first study for testing the theoretical assumptions. We hope this baseline for is useful the emerging community of researchers that study empathy in artificial agents and that it can be expanded through this framework and evaluation methods.

References

- Brave, S., Nass, C., & Hutchinson, K. (2005). Computers that care: investigating the effects of orientation of emotion exhibited by an embodied computer agent. *International journal of human-computer studies*, 62(2), 161–178.
- Cassell, J., Bickmore, T., Campbell, L., & Vilhjálmsón, H. (2000). Designing embodied conversational agents. *Embodied conversational agents*, 29–63.
- Coplan, A., & Goldie, P. (2011). *Empathy: Philosophical and psychological perspectives*. Oxford University Press.
- Davis, M. H., et al. (1980). A multidimensional approach to individual differences in empathy.
- de Waal, F. B. (2007). The ‘russian doll’ model of empathy and imitation. *On being moved: From mirror neurons to empathy*, 35–48.
- de Waal, F. B., & Preston, S. D. (2017). Mammalian empathy: behavioural manifestations and neural basis. *Nature Reviews Neuroscience*, 18(8), 498.
- Hatfield, E., Bensman, L., Thornton, P. D., & Rapson, R. L. (2014). New perspectives on emotional contagion: A review of classic and recent research on facial mimicry and contagion. *Interpersona: An International Journal on Personal Relationships*, 8(2), 159–179.
- Hatfield, E., Cacioppo, J. T., & Rapson, R. L. (1994). Emotional contagion. *Current directions in psychological science*, 2(3), 96–100.
- Hess, U., & Fischer, A. (2014). Emotional mimicry: Why and when we mimic emotions. *Social and Personality Psychology Compass*, 8(2), 45–57.
- Iacoboni, M. (2011). Within each other. *Empathy: Philosophical and Psychological Perspectives*, 45.
- Kopp, S., Krenn, B., Marsella, S., Marshall, A. N., Pelachaud, C., Pirker, H., ... Vilhjálmsón, H. (2006). Towards a common framework for multimodal generation: The behavior markup language. In *International workshop on intelligent virtual agents* (pp. 205–217).
- Leite, I., Castellano, G., Pereira, A., Martinho, C., & Paiva, A. (2014). Empathic robots for long-term interaction. *International Journal of Social Robotics*, 6(3), 329–341.
- Maatman, R., Gratch, J., & Marsella, S. (2005). Natural behavior of a listening agent. In *International workshop on intelligent virtual agents* (pp. 25–36).
- Omdahl, B. L. (2014). *Cognitive appraisal, emotion, and empathy*. Psychology Press.
- Paiva, A., Leite, I., Boukricha, H., & Wachsmuth, I. (2017). Empathy in virtual agents and robots: a survey. *ACM Transactions on Interactive Intelligent Systems (TiiS)*, 7(3), 11.
- Partala, T., & Surakka, V. (2004). The effects of affective interventions in human–computer interaction. *Interacting with computers*, 16(2), 295–309.
- Prendinger, H., Mori, J., & Ishizuka, M. (2005). Using human physiology to evaluate subtle expressivity of a virtual quizmaster in a mathematical game. *International journal of human-computer studies*, 62(2), 231–245.
- Preston, S. D., & De Waal, F. B. (2002). Empathy: Its ultimate and proximate bases. *Behavioral and brain sciences*, 25(1), 1–20.
- Schroder, M., Bevacqua, E., Cowie, R., Eyben, F., Gunes, H., Heylen, D., ... others (2012). Building autonomous sensitive artificial listeners. *IEEE Transactions on Affective Computing*, 3(2), 165–183.
- Spreng, R. N., McKinnon, M. C., Mar, R. A., & Levine, B. (2009). The toronto empathy questionnaire: Scale development and initial validation of a factor-analytic solution to multiple empathy measures. *Journal of personality assessment*, 91(1), 62–71.
- Thiebaux, M., Marsella, S., Marshall, A. N., & Kallmann, M. (2008). Smartbody: Behavior realization for embodied conversational agents. In *Proceedings of the 7th international joint conference on autonomous agents and multiagent systems-volume 1* (pp. 151–158).
- Yalçın, Ö. N. (in press). Empathy framework for embodied conversational agents. *Cognitive Systems Research*.
- Yalçın, Ö. N., & DiPaola, S. (2018). A computational model of empathy for interactive agents. *Biologically Inspired Cognitive Architectures*, 26, 20 - 25. doi: <https://doi.org/10.1016/j.bica.2018.07.010>
- Yngve, V. H. (1970). On getting a word in edgewise. In *Chicago linguistics society, 6th meeting, 1970* (pp. 567–578).