

# Engagement vs. Deceit: Virtual Humans with Human Autobiographies

Timothy Bickmore, Daniel Schulman, and Langxuan Yin

Northeastern University College of Computer and Information Science,  
360 Huntington Ave, WVH202, Boston, MA 02115  
{bickmore, schulman, yinlx}@ccs.neu.edu

**Abstract.** We discuss the ethical and practical issues involved in developing virtual humans that relate personal, fictitious, human autobiographical stories (“back stories”) to their users. We describe a virtual human exercise counselor that interacts with users daily to promote exercise, and the integration of a dynamic social storytelling engine used to maintain user engagement with the agent and retention in the intervention. A longitudinal randomized controlled experiment tested user attitudes towards the agent when it presented the stories in first person (as its own history) compared to third person (as happening to humans that it knew). Participants in the first person condition reported enjoying their interactions with the agent significantly more and completed more conversations with the agent, compared to participants in the third person condition, while ratings of agent dishonesty were not significantly different between the groups.

**Keywords:** Embodied Conversational Agent, Relational Agent, Longitudinal Study.

## 1 Introduction

One design issue faced by all developers of conversational virtual human agents that interact with users in non-entertainment domains is to what extent the agents should present themselves as actually being human. The decision as to *whether* the agents should be presented as humans at all is moot, since fidelity to human appearance and behavior is the overarching objective of this field of research. However, many researchers feel that they are somehow crossing an ethical boundary if their agents start discussing their childhood home or the fight they just had with their (presumably human) spouse. Just as Deckard in the movie *Blade Runner* was shocked when he learned that replicants (bioengineered anthropomorphic beings) were being created with autobiographical memories, many people seem to recoil at the thought of a computer being designed to actually present itself *as* human, without any fictional or “as if” framing. However, there has been no systematic exploration of this topic from an empirical perspective. How would users actually react to agents that present themselves with human autobiographical memories compared to the same agents that make no such pretense? Do users feel cheated and deceived, as many researchers

contend, or do they take it in stride as part of their “suspension of disbelief”? Are there any user benefits to giving agents human personal histories? These are the research questions we sought to address in this work.

Aside from their ethical and intellectual merits, answers to these questions have practical ramifications as well. Many applications in healthcare, education, entertainment and other fields require designing voluntary-use interfaces for long-term use. Designing such systems requires novel approaches to maintaining user engagement over dozens, if not thousands, of interactions. Social chat by agents in these applications provides a mechanism for maintaining user engagement over arbitrary lengths of time, provided that the stories the agent tells are, in fact, entertaining and engaging. Within this context, first person stories may provide the additional engagement required to make a longitudinal application successful.

A number of empirical studies suggest that users actually want agents to be more like them, whether they are conscious of this desire or not. For example, in the Media Equation studies, Reeves and Nass demonstrated that users prefer agents that match them in personality (along the introversion/extroversion dimension) compared to agents that do not [1]. Van Vugt, et al, demonstrated that users prefer characters that match them in body shape [2]. Finally, Bickmore related anecdotes from study participants in which they stated their desire for the animated exercise coach they had worked with for the prior month to have a more human back story [3]. For example:

*“I wish she could imitate a real person's life in her answers rather than sticking to the reality and saying things like she is limited to that box. Maybe this has something to do with trainees wanting to have role model to achieve their own physical fitness roles by taking the trainer as a role model. Or maybe it is just about having a richer conversation helping getting connected to the other person.”*

## 1.1 Ethical Issues

Deception and its negative consequences have been widely studied in ethics [4, 5]. User trust in conversational agents that tell fictitious stories (as well as trust in their developers and marketers) can be greatly damaged if users actually thought the stories agents told were true and later discovered they were not. Widespread use of such deceptive agents could begin to erode generalized trust towards all agents, all technology or universally within a community.

This condemnation of deception extends into the human-computer interaction and agents research communities as well. For example, Fogg states that deception, used in the context of persuasive technology, is “almost always” unethical [6]. Shneiderman contends that computers must clearly relate their capabilities and limitations to users, rejecting any notion of anthropomorphization of the interface [7].

However, deception is rarely a black and white phenomenon. Even ethicists argue whether there are absolute truths, without which deception loses its meaning. Deception is both common in all societies and a necessary component of many professions [4]. One could argue that virtual humans or anthropomorphic robots of any kind represent a kind of deception. Perhaps the degree of deception lies solely in the degree to which such agents are presented without explicit messages or cues that

they are not really human, regardless of the number of messages or cues they present to the contrary (e.g., anthropomorphic body, natural language, etc.).

Docents who provide historical re-enactments at living history museums provide a good analogy to the current issue. Good actors will go to great lengths to stay “in character” even in the face of in-depth questioning and explicit questions about their authenticity (“You’re not *really* Abraham Lincoln, are you?”). However, the larger context of the museum is intended to provide the meta-message that this kind of deceit is not only tolerable, but done for the engagement and benefit of the visitors. Most virtual human researchers who are not working in entertainment-related domains similarly dismiss any accusations of deceit by saying that, obviously, users know they are only interacting with a computer.

Other researchers justify their deceit by saying that people engage in deceitful behaviour similar to the one they are modelling, therefore it must be acceptable for their agents to do the same thing. For example, Klein, in his work on artificial caring, argues that computers that exhibit empathy, sympathy and caring for users are no less authentic than people who express caring for others without really understanding their feelings, or pets who seem to respond in comforting ways to their owner’s negative moods [8].

Finally, some researchers would argue that if their deceit is ultimately to the benefit of the user, then the ends justify the means, and it is sanctioned within a utilitarian ethical framework. For example, Bickmore justifies the possible deceit and manipulation effected by his health promotion agents by the fact that they result in users leading healthier lives [9].

## 1.2 Related Work

Bates, et al, conducted some of the earliest research into the development of virtual characters in the “Oz” project at CMU [10]. The explicit objective in this work was to create a “believable character”, which is not “an honest or reliable character, but one that provides the illusion of life, and thus permits the audience’s suspension of disbelief”. Mateas argues that believability is not the same as realism, and that characters are artistic abstractions of people, which have been exaggerated in order to engage users [11]. He states that believable agents are “designed to strongly express a personality, not fool the viewer into thinking they are human.” Unlike our work, the overarching goal of the Oz project was entertainment, and the work was always presented to users as such. This stance is continued in the majority of work in the growing field of interactive drama and narrative, in which systems are only used to present fictional autobiographies within the explicit framework of make believe.

In contrast, most researchers investigating human-agent interactions in non-entertainment domains carefully avoid giving their agents human back stories. Examples include the Reeves & Nass Media Equation studies [1], studies by Moon [12], Klein [8], and Bickmore [13]. For example, in the Moon study on reciprocal self-disclosure exchanges between a user and a computer, she explicitly states that the computer never referred to itself as “I” to avoid creating the impression that the computer regarded itself in human terms [12]. Self-disclosures for the computer were also scripted to avoid any hint of human back story:

*“This computer has been configured to run at speeds up to 266 MHz. But 90% of computer users don’t use applications that require these speeds. So this computer rarely gets used to its full potential. What has been your biggest disappointment in life?”*

There are a few exceptions, of course. The earliest, and most famous, being the ELIZA system, created intentionally to demonstrate how easy it is to trick people interacting with a computer into thinking they are interacting with a person [14]. This tradition has been continued in the development of many “chatterbots” and the institution of the Loebner prize [15]. Valerie, a robotic receptionist at CMU, was given a running human back story that was continuously updated [16]. However, there have been no experimental investigations into the impact of these back stories on users. We are also unaware of prior investigations in which users were even asked whether they felt they were being deceived by a conversational agent they had interacted with, regardless of how the agent presented itself.

Another related area of investigation is the use of autobiographical memory for virtual agents as a way of making them more adaptive and socially intelligent (e.g., [17]). However, these memories are typically not seeded with a fictitious past for the purpose of relating to a user in a task-oriented context.

### **1.3 An Empirical Investigation**

In order to investigate reactions of actual users to agents that relate personal human (“first person”) back stories, we conducted a randomized longitudinal experiment in which users conducted daily conversations with an agent that related such stories. In the remainder of this paper we describe the experimental framework in which the study was conducted, the narrative generation system that was used to produce the stories, and finally present findings from the experiment itself before concluding and discussing future work.

## **2 The Virtual Laboratory System**

To answer the empirical questions about user reactions to autobiographical agents and how these change over time, we constructed a longitudinal experiment in the “Virtual Laboratory” system [18]. This system provides a framework for running longitudinal studies of ongoing interactions between humans and conversational virtual humans, in which a standing group of study participants interacts periodically with a computer agent that is remotely manipulated to effect different study conditions, with outcome measures also collected remotely. This architecture allows new experiments to be dynamically defined and immediately implemented in the continuously-running system without delays due to recruitment and system reconfiguration. In the current instantiation, 30 older adults interact daily with a virtual human who plays the role of an exercise counselor to promote walking behavior. Older adults were selected as the target population because of their particular need for physical activity and their lower levels of computer literacy [19].



**Fig. 1. Virtual Laboratory Exercise Counselor Agent**

The Virtual Laboratory has been running continuously over the last year, with a total of 36 study participants aged 55 or older conducting a total of over 3,500 conversations with the animated exercise counselor (Fig. 1). The subject pool has had 24 participants on average, with participants staying in the intervention between 18 and 308 days. Participants are on average 60 years old (range 55-75), 73% female, and 54% married.

Fig. 2 shows the virtual laboratory architecture. The client side of this architecture features a virtual agent, web browser, and user input windows (Fig. 1). The server features the following components: an agent database for storing all user data and information about previous user-agent interactions; a measures database for storing all experimental results (e.g., from questionnaires remotely administered to users); an experiment database that contains specifications for all experiments to be run; a dialogue engine that manages conversational interaction between the agent and a user; a web server that provides users with web content (e.g., multimedia educational material and study questionnaire forms); the dialogue engine parameters to instantiate for a particular user on a particular day; an experiment planner that schedules requested experiments; and an experiment evaluator that produces data files and web-based summaries of experimental results.

For the virtual laboratory, we have developed a new dialogue engine—RADIUS (relational agent dialogue system)—which subsumes both augmented transition network-based and task-decomposition-based models of dialogue. In contrast to more complex systems, such as COLLAGEN, RADIUS models a recipe as a state machine,

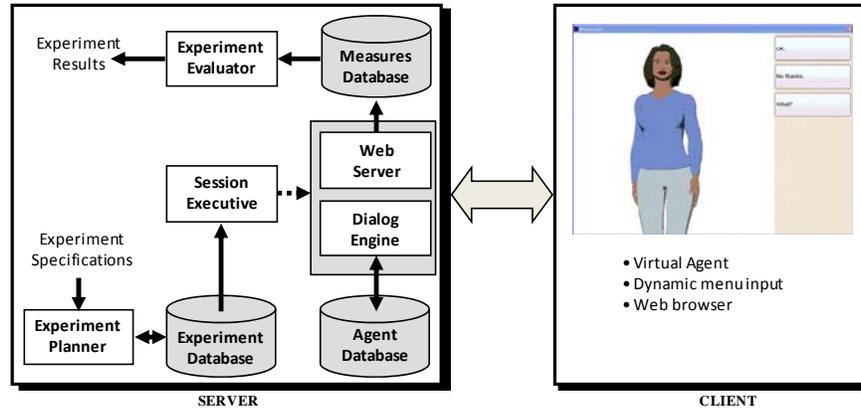


Fig. 2. Virtual Laboratory Architecture

in which agent utterances are states, and user utterances are state transitions. A state transition may invoke a sub-task by specifying a goal, which will cause the dialogue engine to find an appropriate recipe and execute it, before continuing to the next state. In practice, this provides increased modularity and reuse with only a small increase in complexity for authors. Dialogue may still be written as state machines. However, when modifications are required in order to reuse a dialogue fragment, this may be implemented by providing additional recipes for those portions of dialogue.

### 3 Dynamic Social Story Generation

Providing social dialogue in daily conversations between a user and an agent over months or years requires a considerable number of narratives as the agent's background stories. While these could be manually scripted in their entirety, a less laborious alternative is to generate the stories at runtime with a narrative generation system.

#### 3.1 Narrative Generation Technology

A number of interactive narrative generation systems have been developed over the last two decades, such as Façade [22], FearNot! [23], and those developed in the Oz project [10, 11]. These systems employ different levels of natural language generation to create dynamic content that is used to fabricate interactive experiences. Interactive narrative systems, however, are generally domain specific and depend on large scale domain knowledge. Furthermore, in many of these interactive narrative systems, such as Façade and FearNot!, users are allowed to make their contributions using unconstrained typed text input. Narratives generated in response to unconstrained input may fall significantly short of human generated narratives (e.g., lacking in coherence), resulting in loss of believability by the user.

A different approach to narrative generation exists in “Say Anything” [20, 21], which collaborates with users in constructing narratives by contributing sentences extracted from tens of thousands of weblogs. Although this approach creates unique narratives in almost every interaction, and studies have shown that users rate these narratives as being more coherent than ones generated randomly [21], these narratives still fall far short of human-generated stories and do not provide longitudinal coherence (subsequent stories that are logically consistent with earlier ones).

### 3.2 Our Approach to Agent Back Story Generation

We have developed a method for generating social narratives that avoids manually scripting every day’s conversation while providing significant day-to-day variability and maintaining coherence throughout each story. Our approach is similar to Swanson and Gordon’s [21], in that it involves run-time linking of pre-authored story fragments, but differs in several significant ways. We begin with a set of story fragments, each just one to three utterances in length that conveys a complete event or thought. We then manually tag particular words and phrases within each story fragment as mentioned and elaborated concepts, and then we create a link from every story A to story B, where story A has a mentioned or elaborated concept which is also an elaborated concept in story B, following the notions and methodologies from Cleary and Bareiss [22]. This process provides a set of links that point from one story fragment to another, based on common concepts. Finally, we annotate each link with a transition utterance. Fig. 3 is an example of an annotated story fragment, where the utterances between the `<link>` tags point to the other four story fragments. N12 and N13 are two other stories about more of the storyteller’s high school life, and N22 and N23 are stories about sports games.

During a conversation with a user, the system randomly picks one of the story fragments and tells it to the user. Following this, the agent selects a linked fragment (at random if there are several), speaks the transition utterance associated with the link, and then begins telling the linked fragment. Between each story fragment and linking utterance the agent pauses and gives the user the choice to continue to the next utterance, or to repeat the previous one. Each conversation consists of two or three

```
N7:
Mentioned: biking, hiking, walk, picnic, health
Elaborated: family, sports, exercise, outdoor activities
When my family was living in Falmouth, my parents always had us doing outdoor stuff.
So especially when it was nice out I would go biking or hiking or we just go for a walk
and have a picnic, things like that. And I think I really developed an appreciation for
exercise and being outdoors and just staying healthy and moving around all the time.
<link>
N12: And this made my high school life much more interesting.
N13: And this made my high school life much more interesting.
N22: Although I love sports, I don't watch sports games very much.
N23: Although I love sports, I don't watch sports games very much. I go to a few Red
Sox games though.
</link>
```

Fig. 3. Example Story Fragment Representation

story fragments, and thus is composed of seven or eight utterances, including the linking utterances. An example of part of a storytelling interaction can be found in Fig. 4.

<b>1<sup>st</sup>-person</b>	<b>3<sup>rd</sup>-person</b>
1. I'm not quite sure if I told you about this before.	1. I'm not quite sure if I told you about this before.
2. When my family was living in Falmouth, my parents always had us doing outdoor stuff.	2. When her family was living in Falmouth, her parents always had them doing outdoor stuff.
3. So especially when it was nice out I would go biking or hiking or we would just go for a walk and have a picnic, things like that.	3. So especially when it was nice out she would go biking or hiking or they would just go for a walk and have a picnic, things like that.
4. And I think I really developed an appreciation for exercise and being outdoors and just staying healthy and moving around all the time.	4. And I think she really developed an appreciation for exercise and being outdoors and just staying healthy and moving around all the time.

**Fig. 4. Example Narrative Dialogue Showing the Same Story Fragments in 1ST-PERSON and 3RD-PERSON Conditions**

In order to maintain global and longitudinal coherence, we developed an initial set of story fragments for the exercise advisor agent based on autobiographical stories told by a professional exercise trainer. The stories were verbally related to a member of our research staff, recorded, and transcribed. The transcript was then partitioned into fragments and annotated following the scheme above.

#### **4 Longitudinal Evaluation Study**

In order to compare the effects of the use of 1<sup>st</sup>-person and 3<sup>rd</sup>-person narrative dialogue by an agent, we conducted a brief longitudinal study using participants enrolled in the virtual laboratory system. The agent conducted daily conversations about exercise identical to those used in earlier studies with the system [18], with the addition of narrative dialogue generated using the social story generation system described above. Participants were randomized into one of two conditions: In the first (1ST-PERSON), the agent presented the narrative as its own life story, while in the second (3RD-PERSON) the agent presented the narrative as stories about a friend.

We expected that the use of 1<sup>st</sup>-person narrative would promote greater engagement with the agent due to a perception of self-disclosure by the agent, leading to more consistent usage of the system. However, we were also concerned that users would perceive the agent as dishonest when it presented a life story for itself that was not plausibly true for a computer character. Participants were administered daily questionnaires to assess their enjoyment of the stories, their engagement with the system, and their belief that the agent was dishonest.

**Hypothesis 1:** Participants in the 1<sup>st</sup>-person condition will use the system significantly more than those in the 3<sup>rd</sup>-person condition.

**Hypothesis 2:** Participants in the 1<sup>st</sup>-person condition will report greater enjoyment of the stories and greater engagement with the agent than those in the 3<sup>rd</sup>-person condition.

**Hypothesis 3:** Participants in the 1<sup>st</sup>-person condition will report greater perceived dishonesty by the agent than those in the 3<sup>rd</sup>-person condition.

#### 4.1 Participants

A total of 26 participants (21 female, 5 male, aged 54-67, 80% Caucasian, 20% African American) took part in the study, all recruited via ads placed on craigslist.com. The sample was well-educated (92% had some college education), computer literate (12% self-identified as computer experts, the other 88% said they use computers regularly), and had positive attitudes towards computers overall (64% said they enjoyed working with computers). Fifteen had previously been interacting with the system at the start of the study, while 11 were newly recruited. All participants were compensated \$1 per day for each day they completed a conversation with the agent. Exactly half of the participants were randomized into each arm of the study (1ST-PERSON and 3RD-PERSON). Participants were exposed to these study conditions for varying periods of time, ranging from 5 to 37 days (mean 28.8 days).

#### 4.2 Measures

To assess system usage, we recorded whether or not each participant had a complete conversation with the agent each day. Following each complete conversation, after the agent walked off the screen, participants were given three single-item measures in randomized order, asking how much they (1) “enjoy the stories that the counselor tells”, (2) “look forward to talking to the counselor”, and (3) “feel that the counselor is dishonest”. Each item was assessed on a 5-point rating scale ranging from “not at all” to “very much”.

#### 4.3 Narrative Dialogue

Narrative social dialogue was generated using the dynamic social story generation described above. In the first-person condition, the narratives were initially introduced as being part of the agent’s own life story (“I’d like to tell you some stories about myself”). In the third-person condition, the narratives were introduced as being from the life story of a human friend of the agent with a similar role and occupation (“I’d like to tell you some stories about a friend of mine. She’s an exercise counselor too.”).

The differences between the first- and third- person variants of the dialogue were minimal, and consisted mainly of replacing pronouns. Fig. 4 shows an example of the narrative dialogue, in both variants.

#### 4.4 Results

The 3 self-report items were analyzed by fitting linear mixed-effect regression models<sup>1</sup> to the data, while system usage was analyzed as a binary outcome with a logistic mixed-effect regression model. Analysis was performed using R 2.9.0, with the “nlme” and “lme4” packages [23].

For all outcomes, models were used which included fixed effects of study day and study condition. Initially, we considered models which included an additional fixed effect modeling the interaction of day and condition, thus allowing for a different rate of change in the outcomes between the two conditions. However, both inspection of the data and model selection procedures indicated that any interaction effects were minimal, most likely due to the short duration of the study. All models include random effects of intercept and study day. Table 1 shows the results of the analysis.

**Table 1. Mixed-Effect Regression Estimates of Effects of Study Day and Condition on Outcomes**

Condition 0 = 1ST-PERSON, 1=3RD-PERSON  
\* p < 0.5; \*\* p < 0.01; \*\*\* p < 0.001

		Look Forward	Enjoy Stories	Dishonest	System Usage
Random Effects	Intercept	0.676 ***	1.127 ***	0.794 ***	1.477 ***
	Day	0.031 ***	0.038 ***	0.034 ***	0.012 ***
Fixed Effects	Intercept	4.410 *** (0.198)	3.384 *** (0.326)	1.688 *** (0.236)	3.207 *** (0.478)
	Day	-0.017 * (0.007)	-0.035 *** (0.009)	0.272 (0.326)	-0.046 *** (0.010)
	Condition	0.145 (0.281)	-1.059 * (0.461)	0.002 (0.008)	-1.148 * (0.560)

**Enjoyment of the Stories.** Participants in the 1ST-PERSON condition reported significantly greater enjoyment of the stories compared to those in the 3RD-PERSON group. There was also a significant effect of study day on the degree to which participants reported enjoying the stories; participants reported decreasing enjoyment of the stories over time (approximately 0.035 per day).

**Engagement.** There were no significant differences between conditions on degree to which participants said they “looked forward” to working with the agent. However,

<sup>1</sup> Linear mixed-effect regression is a generalization of ordinary linear regression, which adds random effects in order to account for clustered data, such as multiple measurements per subject in a longitudinal study. Similarly, logistic mixed-effect regression is a generalization of logistic regression, suitable for analyzing repeated binary measurements.

the average participant (both groups) reported significantly decreasing levels of engagement over time (approximately 0.017 per day).

**Perceived Dishonesty.** Participants, overall, did not perceive the agent as very dishonest. Average perceived dishonesty (both groups) following the first conversation was 1.69 (on a 1="not at all" to 5="very much" scale). There was no significant effect of study day or of study condition on this measure.

**System Usage.** Participants in the 1ST-PERSON condition had a significantly greater probability of talking to the agent on any given day, compared to those in the 3RD-PERSON group (Fig. 5). There was also a significant effect of study day on this measure; for the average participant, the probability of completing a session on any given day decreased over time.

**Continuing vs. New Participants.** The 11 participants who were newly recruited for this study did use the system significantly more compared to participants who had already been interacting with the system at the start of the study,  $p=.01$ . Including old vs. new participant as a covariate in the regression analysis does not change the significance status of any of the results above.

#### 4.5 Discussion

As hypothesized, participants who interacted with an agent that used first-person stories reported greater enjoyment of the stories, and were more likely to use the system. Therefore, we can conclude that the first-person stories led to greater engagement with the system, at least during the short duration of time studied here. Both measures had significant decreases over time (in both conditions). This is likely due to increasing repetitiveness, as the agent had only a small set of story

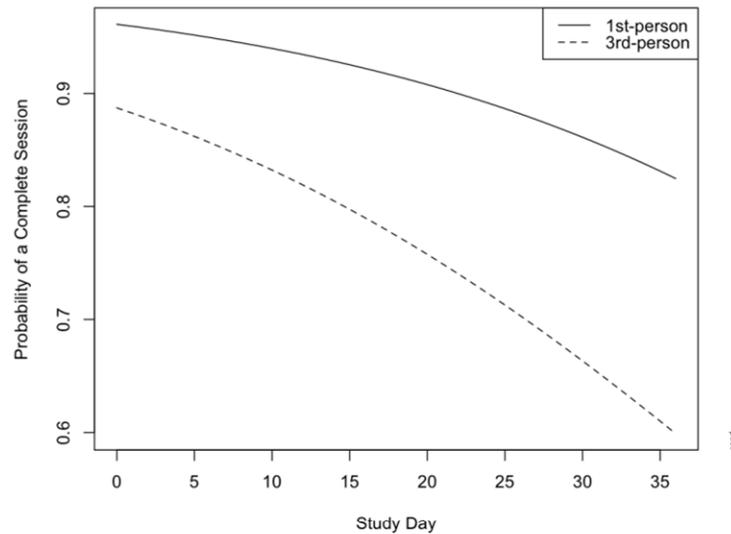


Fig. 5. Effects of Study Day and Condition on System Usage

fragments to draw from in generating each day's story.

However, participants were not significantly more likely to report that they looked forward to working with the agent in the first-person condition. We consider two possible explanations: First, scores on this measure were all quite high (mean 4.22 on a 5-point scale), so ceiling effects may be hiding any difference caused by study conditions. Second, this result may indicate that our self-report measure of engagement does not reflect actual behavior; this raises methodological issues for future studies.

Participants were not significantly more likely to report that the agent was dishonest when it used first-person narrative, despite the fact that these stories could not possibly be true stories about a virtual character. This result suggests that users are willing to accept a fictional narrative that would be plausible for the character if the character were human.

Finally, we note that there were highly significant random effects of intercept and study day on all outcomes: there was a large amount of variability between participants both on outcomes at the start of the study, and on the rate of change over time. This suggests that there are predictors we have not examined, such as personality traits, which play an important role in these outcomes. This may represent a fruitful area of research in the future. Other limitations of the study include the small sample size, the recruitment methodology that resulted in a relatively well-educated and computer literate sample, the relatively short time span of the intervention, and the use of single-item questionnaires.

## 5 Conclusion

We believe this is the first systematic investigation into user reactions to human autobiographical stories told by a virtual human. We find that, rather than rejecting such an agent because of the perception of being deceived, the stories led to greater user engagement with the agent, and users did not rate the agent as being any less dishonest than an equivalent agent that did not relate the stories as its own history. Whether this behavior should be called deceitful is an open question.

Of course, participants in the study knew that they were interacting with a virtual human from the outset: they were told during enrollment that they would be interacting with a computer character; the agent is rendered using a cartoon rendering style (Fig. 1); and most participants had already experienced weeks or months of daily conversations with the agent that most report as repetitive and robot-like before beginning the study described above. How users would feel if one or more of the above cues regarding the agent's authenticity were removed also remains an open question.

These results are significant for designers of "serious" virtual humans that engage users in counseling, pedagogical or health care conversations over long periods of time. Maintaining user engagement with these systems is a pre-requisite for achieving any intervention outcomes, since users who stop using such a system or use it at a sub-optimal frequency do not receive the therapeutic and informational messages required to achieve the desired results. The autobiographical stories evaluated in our

study perform what Jakobson defined as the “phatic” function of dialogue, to simply keep the communication channel open so that the primary functional messages can be conveyed [24].

In our ongoing work in this area we are developing virtual human-based health counseling interventions that span a year or more of daily conversations with a user. In addition to procedural dialogue content generation (e.g., based on weather data or sports scores) we see autobiographical conversational storytelling by the agent as one of the most important methods available for maintaining user engagement in the intervention over time. We plan to continue exploring ways in which such stories can be dynamically generated and integrated into the counseling conversation in a coherent and natural manner [25].

**Acknowledgments.** Thanks to Thomas Brown, Jenna Zaffini and the rest of the Relational Agents Group at Northeastern for their help, and to Jennifer Smith and Laura Pfeifer for their helpful comments on this paper. This work was supported by NSF CAREER IIS-0545932.

## References

1. Reeves, B., Nass, C.: *The Media Equation*. Cambridge University Press, Cambridge (1996)
2. van Vugt, H., Konijn, E., Hoorn, J., Veldhuis, J.: *Why Fat Interface Characters Are Better e-Health Advisors*. *Intelligent Virtual Agents (IVA)*, Marina Del Rey, CA (2006) 1-13
3. Bickmore, T.: *Relational Agents: Effecting Change through Human-Computer Relationships*. PhD Thesis in Media Arts & Sciences. Massachusetts Institute of Technology, Cambridge, MA (2003)
4. Nyberg, D.: *The varnished truth: truth telling and deceiving in ordinary life*. University of Chicago Press, Chicago (1993)
5. Ford, C.: *Lies! Lies!! Lies!!! The psychology of deceit*. American Psychiatric Publishing, Washington, DC (1999)
6. Fogg, B.J.: *Persuasive Technology: Using Computers to Change What We Think and Do*. Morgan Kaufmann, New York (2003)
7. Schneiderman, B.: Looking for the bright side of user interface agents. *interactions* **2** (1995) 13-15
8. Klein, J., Moon, Y., Picard, R.: This Computer Responds to User Frustration: Theory, Design, Results, and Implications. *Interacting with Computers* **14** (2002) 119-140
9. Bickmore, T.: *Ethical Issues in Using Relational Agents for Older Adults*. AAAI Fall Symposium on Caring Machines: AI in Eldercare, Washington, DC (2005)
10. Bates, J.: The Role of Emotion in Believable Agents. *Communications of the ACM* **37** (1994) 122 - 125
11. Mateas, M.: *An Oz-Centric Review of Interactive Drama and Believable Agents*. In: Wooldridge, M., Veloso, M. (eds.): *AI Today: Recent Trends and*

- Developments: Lecture Notes in AI 1600. Springer-Verlag, Berlin, New York (1999) 297-328
12. Moon, Y.: Intimate self-disclosure exchanges: Using computers to build reciprocal relationships with consumers. Harvard Business School report, Cambridge, MA (1998)
  13. Bickmore, T., Picard, R.: Establishing and Maintaining Long-Term Human-Computer Relationships. *ACM Transactions on Computer Human Interaction* **12** (2005) 293-327
  14. Weizenbaum, J.: Eliza - a computer program for the study of natural language communication between man and machine. *Communications of the ACM* **9** (1966) 36-45
  15. Loebner, H.: The Loebner Prize in Artificial Intelligence. <http://www.loebner.net/Prize/loebner-prize.html>
  16. Gockley, R., Bruce, A., Forlizzi, J., Michalowski, M., Mundell, A., Rosenthal, S., Sellner, B., Simmons, R., Snipes, K., Schultz, A.C., Wang, J.: Designing Robots for Long-Term Social Interaction. *IEEE/RSJ International Conference on Intelligent Robots and Systems* (2005)
  17. Ho, W., Watson, S.: Autobiographic Knowledge for Believable Virtual Characters. *Intelligent Virtual Agents* (2006) 383-394
  18. Bickmore, T., Schulman, D.: A Virtual Laboratory for Studying Long-term Relationships between Humans and Virtual Agents *Proceedings of Autonomous Agents and Multi-Agent Systems, Budapest, Hungary* (2009)
  19. Bickmore, T., Caruso, L., Clough-Gorr, K., Heeren, T.: "It's just like you talk to a friend" - Relational Agents for Older Adults. *Interacting with Computers* **17** (2005) 711-735
  20. Gordon, A., Cao, Q., Swanson, R.: Automated Story Capture From Internet Weblogs. *Fourth International Conference on Knowledge Capture, Whistler, BC* (2007)
  21. Swanson, R., Gordon, A.: A Comparison of Retrieval Models for Open Domain Story Generation. *AAAI 2009 Spring Symposium on Intelligent Narrative Technologies II, Stanford, CA* (2009)
  22. Cleary, C., Bareiss, R.: Practical Method for Automatically Generating Typed Links. *Hypertext'96* (1996) 31-41
  23. R Core Development Team: R: A language and environment for statistical computing. 2009. R Foundation for Statistical Computing, Vienna, Austria (2009)
  24. Jakobson, R.: Linguistics and Poetics. In: Sebeok, T.A. (ed.): *Style in language*. MIT Press, Cambridge, MA (1960) 130-144
  25. Jefferson, G.: Sequential aspects of storytelling in conversation. In: Schenkein, J. (ed.): *Studies in the organization of conversational interaction*. Academic Press, New York (1978) 219-248