

Increasing the Engagement of Conversational Agents through Co-Constructed Storytelling

Cristina Battaglini

College of Computer and Information Science,
Northeastern University, Boston, MA, United States

Timothy Bickmore

College of Computer and Information Science,
Northeastern University, Boston, MA, United States

Abstract

Storytelling can be used by conversational agents in a wide variety of domains to maintain user engagement, both within a single interaction and over dozens or hundreds of interactions over time. The majority of agents designed with this ability to date deliver their stories as monologues without user input. However, people rarely tell stories in conversations this way, and instead rely on listener contributions to guide the storytelling process. Corpus-based studies of human-human conversational storytelling have demonstrated greater engagement, in the form of longer stories, when listeners co-construct stories this way.

We describe a research framework for the generation and evaluation of co-constructed social stories in the context of task-based conversations, and a study on the effects of degree of user-agent story co-construction on user engagement. We find that users are more engaged with storytelling agents that allow them to co-construct stories in a contentful manner by asking questions, compared to co-construction through acknowledgments only.

Introduction

The design of engaging characters is important not only in entertainment, but in serious applications of conversational agents. In domains such as exercise promotion (Bickmore et al. 2013; King et al. 2013), medication adherence (Bickmore et al. 2010), or social support and companionship for elder-care (Vardoulakis et al. 2012), keeping users cognitively engaged and motivated to keep talking to an agent for dozens or hundreds of conversations is a prerequisite for other system objectives: if a user stops interacting with the agent, then it cannot have any further impact.

Conversational storytelling provides a communication channel that an agent can use to maintain engagement. While stories can be used to convey instrumental information about the task being performed (e.g., relating how another user met their exercise goals), “off-task” stories, in the guise of “social chat” or “small talk”, can be used to expand the agent’s conversational repertoire to an alternate universe of nearly unbounded conversational topics. When used in

a natural, context-appropriate manner, such story-telling by the agent has been shown to significantly increase user engagement. In a study of conversational “backstories” told by a conversational agent playing the role of an exercise trainer, it was demonstrated that users logged in significantly more frequently over the course of a month when the agent told stories about itself, compared to the same stories told about someone else (Bickmore, Schulman, and Yin 2009).

We are developing a conversational agent that tells off-task stories to users to keep them engaged in a year-long exercise promotion program. In order to maintain the naturalness and therapeutic integrity of the health counseling intervention, we are designing the agent to tell these stories in the same way a human exercise trainer would, in the context of an overall counseling session.

Following conversation analysis studies by Jefferson and Sacks, conversational stories are: *locally occasioned*, that is, related in some way to the conversation that preceded it (Jefferson 1978); *recipient designed*, in that the basic facts of a story can be told in very different ways, depending on the audience; and *co-constructed*, in that the listener is almost never a passive recipient of a monologue, but rather collaborates with the storyteller in a variety of subtle (and not so subtle) ways to influence the storyteller’s behavior (Sacks and Jefferson 1995). Some past research has investigated how a conversational agent can tell locally occasioned stories, for example by modeling topic coherence relations among task and non-task discourse segments and using discourse markers to transition between task and off-task talk (Cassell and Bickmore 2003). Others have investigated automated tailoring of stories to particular audiences, for example manipulating linguistic style features based on rhetorical goals (Hovy 1987). In our current work, we are focused on developing the ability for our exercise trainer agent to tell conversational stories that are co-constructed with the user. Our working hypothesis is that stories that are told in a co-constructed manner are more engaging to users than stories that are delivered as monologues, with minimal opportunity for user input.

Related Work

There have been several agents developed in the research community that incorporate storytelling into a task-based conversation. Cassell, et al, used topic coherence relations

and a spreading-activation network-based dialogue planner to interleave task talk and social chat in the context of a real-estate sales conversation (Cassell and Bickmore 2003). The social chat was shown to significantly increase user trust in the agent, but only for more extroverted users. Bickmore, et al, developed a simple story generator for an agent to tell stories about itself or others, used in the context of exercise counseling conversations (Bickmore, Schulman, and Yin 2009). In (Lombardo and Damiano 2012), a virtual conversational agent presents cultural information about an historical site by conveying it through storytelling. Ring, et al, use social storytelling, and other methods, to decrease loneliness in isolated older adults (Ring et al. 2015).

Other research efforts have focused on methods for making storytelling in user-agent interactions more coherent and/or effective. Swanson and Gordon developed “Say Anything”, a system that used story fragments gleaned from Internet blog posts to co-construct stories with users (Swanson and Gordon 2012). Smith developed a story listening agent that used active listener behaviors (e.g., backchannel feedback) to elicit stories from older adults, and was shown to elicit significantly longer stories compared to a no feedback condition (Smith 2000). Gratch, et al, created a “rapport agent” that exhibited listener behaviors, although they found that users spoke significantly fewer words to this agent when telling stories, compared to the same agent without listener behavior (Gratch et al. 2007).

Co-Constructed Stories in Human-Human Conversation

Stories are an important element of human life. As argued by J. Bruner, stories enforce the canonicity of a culture’s value and humans share their experiences through stories (Bruner 1991; Ochs 1997). According to Bruner, “stories achieve their meanings by explicating deviations from the ordinary in a comprehensible form” (Bruner 1991). Humans use stories for communicating with other people in an effective way and the mechanism of stories is one of the most ancient form of entertainment (Jefferson 1978). They have the referential function of giving information by referring to an experience or by recapitulating an experience (i.e., reporting what it is happened) and they have the evaluative function of communicating the meaning of the narrative by establishing some point of personal involvement (Labov 1972).

Story can be expressed through different modes, e.g., books, movies, games, etc. One of the most basic universal form in which stories are expressed and shared is through ordinary conversation (Bruner 1990). When stories are enacted in conversations, we refer to them as narrative conversation (Rühlemann 2014).

Rühlemann conducted a corpus analysis of 531 naturally-occurring conversational stories, involving 385 individuals, comprising a total of 149,520 words, (Rühlemann 2014). Interlocutors were labeled with one of six roles that describe how narrators and recipients respond to the development of a story (Table 1).

The participant who is the narrator of the story can be marked with the following roles: *Participant Narrator Un-*

Narrative Status	Narrator Roles	Recipient Roles
Monologue	PNU Narrator Unsupported	Inactive
Story is co-constructed	PNS Narrator Supported	PRR Recipient Responsive
	PNP Narrator Primary	PRC Recipient Co-constructive

Table 1: Participants roles in narrative conversation (Rühlemann 2014).

PRR Responses: *yeah, oh, and, well, so, okay, you, know, um, mm, right, uh, really, ah, wow, oh.*

PRC Responses: *wh-questions (what, why, where, who, whom, when, etc.), yes/no questions, comments or extension based on general knowledge.*

Figure 1: Examples of PRR and PRC responses (Rühlemann 2014; Norrick 2009).

supported (PNU), *Participant Narrator Supported* (PNS) and *Participant Narrator Primary* (PNP) (Table 1). The participant who is listening the story (i.e., the recipients) can be assigned to the following roles: *Participant Recipient Responsive* (PRR) and *Participant Recipient Co-constructive* (Table 1). Participants’ roles are dynamic; they change over the narrative conversation. The PNU role is an unsupported narrator that does not receive any feedback from recipients for continuing the story (Table 1). When the narrator receives feedback from a responsive recipient, she becomes a PNS (Table 1). Responsive Participants (PRR) signal their interest in the story but they don’t add or elicit new information on story back-ground. PRR recipients in general provide backchannel feedback (Ward and Tsukahara 2000), including head nods, and interjections (Norrick 2009) (see Figure 1).

Rühlemann found that recipient behavior has impact on the storyteller and the overall length of the story told: PRR and PRC recipients lead to significantly longer stories being co-constructed. Thus, we hypothesize that a conversational agent that allows users to participate in agent story-telling in a more significant manner, as co-constructive (PRC) recipients for the stories that it tells, will be more engaging than agents that either do not allow for any user input or only user acknowledgment (PRR).

Towards a Computational Model

As we have shown, narrative conversations should not just be monologues in which an agent tells something that has happened, but are, ideally, dialogues in which users are the




Figure 2: Conversational Agent User Interface.

recipients of the story and also play an active role in the construction of the narrative itself. One of the issues that we need to address, and that arises in building a conversational agent that can engage a user in a long-term relationship through narrative conversation, is how we can design the interactive experience in a way that users feel engaged in the construction of the story.

We are conducting an exploration of co-constructed conversational storytelling using an existing conversational agent framework (Figure 2). In this system, dialogue is modeled using hierarchical transition networks. Agent utterances are produced using template-based text generation and synthetic speech accompanied by animated nonverbal behavior (hand gestures, facial displays, posture shifts, etc) generated using BEAT (Cassell, Vilhjálmsón, and Bickmore 2004). User utterances are made via multiple-choice menus updated each turn of conversation. The agents are used to counsel users on health topics over sequential conversations that can last years, supported by a persistent memory saved between conversations.

There are several ways that co-constructed storytelling could be computationally modeled. PRR contributions for the user are straightforward to generate, since they do not depend upon the content of the story. PRC contributions, however, are more challenging. One approach is to start with a monologue, and automatically generate “WH”- questions and “Y/N”-questions (PRC contributions) for the user to ask preceding each agent contribution to the study.

Others have used automated rhetorical structure theory (RST) parsers to generate such questions for tutorial dialogues (Kuyten et al. 2012). More sophisticated approaches (and ones that would better engage the user) would be for the agent to tell stories about topics it knew the user had knowledge of, so that the user could contribute to the telling in a meaningful way. Examples include stories about past interactions between the user and the agent, and stories about shared context (weather conditions or current events, gleaned from the Internet).

	Dialogue 1	Dialogue 2	Dialogue 3
Topics	Story: Visiting Cape Cod; Story: Grandfather; Exercise tips; Story: First 10k race	Story: Spa experience; Story: Friend’s dog; Diet tips; Story: Moving in Boston	Story: Friend’s social life; Story: Friend’s family adventures; Story: Friend’s losing weight
Turns	48	48	116 max
Time	~5 mins.	~5 mins.	~20 mins. max

Table 2: Participants roles in narrative conversation (Rühlemann 2014).

Narrative Conversational Agent Study

We conducted an empirical study to determine whether degree of co-construction in agent storytelling (in the context of a task-oriented conversation) would have a significant impact on user engagement.

For the purpose of this study we used *manually-constructed dialogues* in order to control all aspects of the experimental stimuli. The study was a randomized, counter-balanced, within-subjects experiment with two treatments: PRR, in which the user could only make backchannel acknowledgments after each agent contribution to a story; and PRC, in which the user could “ask” a “WH”-question or “Y/N”-question prior to each agent story contribution or can extend the agent contribution with general comments (e.g., “*yeahs, kids are fun! They always do funny things.*”). The user contributions, both PRR and PRC response types, are manually designed as part of the manually-constructed dialogues in a script-like manner. The user contributions are prompted as a menu on the right of the user interface (see Figure 2). As the user contributions, the agent utterances are also manually designed, scripted and then generated through text-to-speech technology.

We created two dialogues that consisted of an agent greeting and role-setting statements (“*I’m going to talk to you about exercise...*”, 6 turns), social storytelling (22 turns), task-oriented talk consisting of health tips (8 turns), more social storytelling (10 turns), and a farewell (2 turns) (Table 2).

One dialogue addressed exercise and the other healthy diet. We also created a third dialogue that was significantly longer and was comprised of a re-greeting, and 110 turns of social storytelling, each of which offered the user a polite way to end the conversation (e.g., “*Sorry, I have to run.*”)(Table 2). Each of these three dialogues could be used with either PRR or PRC user contributions after each agent utterance (Figure 3).

The manually-constructed dialogues are human authored so that the user can perceive that the agent is co-constructing

NARRATIVE CONVERSATION 1	NARRATIVE CONVERSATION 2
<p>...</p> <p>Agent: I can't wait until the weekend. I am going with my partner to visit Cape cod.</p> <p>[User Menu Responses]</p> <p>PRR> Really?</p> <p>PRC> Are you driving?</p> <p>Agent: We are renting a car. We plan to tour several of the towns on the Cape.</p> <p>[User Menu Responses]</p> <p>PRR> OK.</p> <p>PRC> Are you going to the beach?</p> <p>Agent: We are going to get some rest on the beach and enjoy the weather. The forecast says the weather should be nice this week end.</p> <p>[User Menu Responses]</p> <p>PRR> Yes.</p> <p>PRC> Aren't you glad Spring has finally arrived?</p> <p>...</p> <p>[Task-oriented conversation – exercise]</p> <p>...</p> <p>[End task-oriented conversation – exercise]</p> <p>Agent: And remember. After physical activity, you should always do 5 minutes of stretching.</p> <p>[User Menu Responses]</p> <p>PRR> Ok.</p> <p>PRC> Ok.</p> <p>[Narrative conversation]</p> <p>Agent: One of my friends, Kathy, has just started jogging to lose weight because she is getting married this fall.</p> <p>[User Menu Responses]</p> <p>PRR> Really?</p> <p>PRC> What happened?</p> <p>Agent: She came in today and could barely walk after her first run.</p> <p>Stretching is important!</p> <p>[User Menu Responses]</p> <p>PRR> Yes.</p> <p>PRC> I understand.</p> <p>...</p>	<p>...</p> <p>Agent: I met my friend Katie.</p> <p>[User Menu Responses]</p> <p>PRR> Oh?</p> <p>PRC> What did you do?</p> <p>Agent: We went for coffee at the Pavement Coffee House.</p> <p>[User Menu Responses]</p> <p>PRR> Nice.</p> <p>PRC> What did you have?</p> <p>Agent: I had my favorite, a double caramel cappuccino. Yum!</p> <p>[User Menu Responses]</p> <p>PRR> ooh.</p> <p>PRC> Tell me about your friend.</p> <p>Agent: I've known Katie from high school. She's actually getting married, and is planning a honeymoon in Italy!</p> <p>...</p> <p>[Task-oriented conversation – diet]</p> <p>...</p> <p>[End task-oriented conversation – diet]</p> <p>Agent: And remember that breakfast is the most important meal of the day. Don't underestimate it and do always have your breakfast.</p> <p>[User Menu Responses]</p> <p>PRR> I will.</p> <p>PRC> I will.</p> <p>[Narrative conversation]</p> <p>Agent: Oh speaking of breakfast, that reminds me that I need to buy my yogurt for tomorrow and prepare for making pancakes. I get crazy if I don't have my morning routines. And my daughter has been asking me to make pancakes for a couple of days now and I always forget to do it. I am so busy lately.</p> <p>[User Menu Responses]</p> <p>PRR> Really?</p> <p>PRC> Why are you so busy?</p> <p>Agent: Well, we are moving outside of Boston. We found a lovely house at a great price in Newton. So, I've been packing all our things every evening. Its hard. Its incredible how many things we can accumulate in our life!</p> <p>[User Menu Responses]</p> <p>PRR> Yeah!</p> <p>PRC> Yeah!</p> <p>...</p>

Figure 3: Examples of PRR and PRC story responses in the first conversation.

the narrative conversation with her, following findings and narrative conversation examples in (Rühlemann 2014; Norrick 2009). The consequence is that the turns of the conversational agent are scripted as longer turns than user PRR and PRC options prompted on the menu on the right side of the dialogues. Table 2 provides more details on the dialogues.

Procedure

Following informed consent and assessment of socio-demographics, participants were randomized to either conduct a PRR or PRC dialogue first with one of two agents (Figure 4). Following completion of this conversation, they filled out a questionnaire with 7-point scale items (Table 3). They then conducted the other style of conversation (PRR vs. PRC) and filled out the rating questionnaire again. The order of agent (Figure 4) and dialogue topics (Table 2) were fixed, but PRR/PRC were varied, to counterbalance any confound of agent or dialogue topic on results.

After completing the two conversations, we asked participants to choose one of the agents they had just interacted with to have a third conversation (their choice). They then conducted the long storytelling dialogue with the chosen agent (using the same PRR/PRC format that the chosen agent had used earlier). The PRR/PRC style of the preferred agent and the number of turns of social dialogue the user conducted before terminating the conversation were used as additional measures of engagement.

Following the third conversation, we conducted a semi-structured interview with participants about their experience



Figure 4: The two conversational agents used in the study. Angela on the right, Tanya on the left.

and preferences.

Participants

We recruited participants via ads on an online job listing site. Participants were required to be 18 years of age or older and speak and read English. All participants were compensated for their time.

Question	Anchor 1	Anchor 7	PRR M(SD)	PRC M(SD)	p- value
How much would you like to continue with the agent?	Not at all	Very much	4.6 (1.8)	5.0 (1.6)	.196
How satisfied are you with the agent?	Not at all	Very satisfied	4.9 (1.6)	5.2 (1.5)	.278
How natural was your conversation with the agent?	Not at all	Very natural	4.3 (1.7)	4.3 (1.5)	.752
How much do you like the agent?	Not at all	Very much	4.8 (1.3)	4.8 (1.3)	.907

Table 3: Self-Report Measures and Outcomes (Wilcoxon Signed-Rank Tests).

Measures

We collected three measures in order to measure the engagement of participants and understand if participants were more engaged in the PRC condition than in the PRR condition. We asked participants to fill a satisfaction questionnaire after each two conversations, while for the third conversation we recorded the number of turns spent before the user stopped the interaction. After finish the two conversation, we collected the *desire to continue* self-report measure through a questionnaire by asking participants “How much would you like to **continue** working with the agent?” with a 7-item Liker scale (anchors: not at all, very much). The other two measures of engagement we used are the *choice* made by participants for the third conversation and the *number of turns* (Behavioral Outcome Measure spent in the third conversation).

Results

A total of 36 participants (11 female, 25 male), aged 19-65 (mean 33.8), took part in the study.

Engagement: Self-Report. Participants in the PRC condition reported a greater desire to continue interacting with the PRC agent compared to the PRR agent, but this difference was not significant (Table 3).

Engagement: Agent Choice. The majority of participants (61.1%) chose to continue their conversation with the PRC agent, although this difference was not significantly different from random, $X^2(1)=1.78$, n.s.

Engagement: Number of turns. Participants conducted significantly more turns of dialogue in the third conversation with the PRC agent compared to the PRR agent, 88.0 (SD 32.5) vs. 57.0 (SD 30.0), Mann-Whitney $U=77.5$, $p=.05$.

Attitude towards agent. There were no significant differences on any of the other self-report measures regarding attitudes towards the PRC vs. PRR agent (Table 3).

Qualitative results. During the interviews, most of the participants who chose the PRC agent confirmed that their choice was due to engagement. Participants also perceived the differences between the two styles of conversation. For example, one participant who chose the PRC agent for the third conversation reported:

“and it respects my perspective. The other agent [PRR condition] is not really listening”

“I prefer user options [PRC responses] in which it seems that the agent is responding in a more empathic way”

Some participants also commented on the positive impact of the social chat on their engagement with the task talk:

“Social chat helps me remember what the agent said to me about exercise because I have the feeling of building a relationship with the agent...”

Conclusion

As we hypothesized, participants were significantly more engaged when they felt they were co-constructing stories with the agent to a greater degree (via PRC contributions), as measured by the number of turns of storytelling they were willing to engage in. Although most other measures of engagement also favored PRC co-construction, the differences did not reach statistical significance. Regarding the self-report measures and the choice measure, despite we randomized the two agents and the conditions PRR and PRC, it is emerged during the interview that participants choices are driven by some bias related to the topic of the stories (i.e., participants rated and chose the agent based on their preferences on the diet or exercise task topic of Story 1 and Story 2) and/or by the appearance of the agent. This suggests that we need a greater sample in order to overcome these bias.

The hypothesis that Rühlemann findings (i.e., the impact of PRR and PRC responses on the length of the stories told by the narrator) hold also in conversational agent it is confirmed by the Behavioral Outcome Measure. Notwithstanding participants’ choices, participants who chose the PRC condition interacted longer in the third story than participants who chose the PRR condition for the third story.

We can conclude that, in modeling conversational storytelling by a conversational agent, we need to take into account the user role in the process.

Future Work

As future work we want to concentrate on endowing the conversational agent with first experience stories based on a fictitious background life and the current context of the interaction. We are driven in this choice by two main reasons. First, during the interviews many participants stated

that they preferred first person narratives over third person narrative, confirming the findings in (Bickmore, Schulman, and Yin 2009). Second, in his analysis of human conversations, Rühlemann found that 63% of narrative conversations were first person accounts (Rühlemann 2014).

We plan to explore procedural generation of first experience narratives using planning techniques (Porteous, Cavazza, and Charles 2010) and the notion of plot graphs (Riedl and Young 2006). In order to overcome the author bottleneck of encoding domain knowledge for stories, plot graphs can be also crowd-sourced from the web (Li et al. 2013). Narrative prose can then be generated from the plot graphs by following, for example, the work in (Callaway and Lester 2002).

Another approach for content generation that we want to explore is the acquisition of structured events from the crowd. The basic idea is to personalize acquired events based on a storyteller’s profile and user preferences, as proposed in (Sina, Rosenfeld, and Kraus 2013). Furthermore, by using this approach, we can acquire “Wh”-questions and “Y/N”-questions by submitting semi-structured questionnaires to crowd workers.

The main issues that we need to address in building a narrative conversation agent framework is how we can generate narratives in such a way that the stories lead to a natural and coherent conversation. As a next step, we want to explore how we can adopt NLP techniques for addressing two problems:

1. *the generation of PRC responses*. This problem can be addressed by analyzing the event that the agent is telling with NLP techniques (Jurafsky and Martin 2000) and “WH”-questions and “Y/N” questions can be generated by adopting template-based NLG techniques (Reiter, Dale, and Feng 2000);
2. *linking topics*. We need to trigger new topics based on the current context of the conversation. For example, by generating a list of name entities (Nadeau and Sekine 2007) that can enact narratives based on the correlated topics (and that the user can decide to follow or not). This can resolve the issue of maintaining coherence in the conversation, but the problem of transitions between topics in a fluent way still needs to be addressed.

Finally, we plan to explore how we can use the longitudinal record of user-agent interactions to produce narratives of past interactions, so that the user is engaged in perceiving that the agent has memories of the story that they are building together day by day. Narratives of human personal relationships are key in determining how those relationships are perceived (Frost 2013). We are also investigating how we can include real-time weather reports and daily news in the agent’s social conversation in a way that the user perceives that the agent is aware of the world, and produces a narrative that user is actually able to contribute to in the production of a co-constructed story.

References

- Bickmore, T. W.; Puskar, K.; Schlenk, E. A.; Pfeifer, L. M.; and Sereika, S. M. 2010. Maintaining reality: relational agents for antipsychotic medication adherence. *Interacting with Computers* 22(4):276–288.
- Bickmore, T. W.; Silliman, R. A.; Nelson, K.; Cheng, D. M.; Winter, M.; Henault, L.; and Paasche-Orlow, M. K. 2013. A randomized controlled trial of an automated exercise coach for older adults. *Journal of the American Geriatrics Society* 61(10):1676–1683.
- Bickmore, T.; Schulman, D.; and Yin, L. 2009. Engagement vs. deceit: Virtual humans with human autobiographies. In *Intelligent Virtual Agents*, 6–19. Springer.
- Bruner, J. S. 1990. *Acts of meaning*. Harvard University Press.
- Bruner, J. 1991. The narrative construction of reality. *Critical inquiry* 1–21.
- Callaway, C. B., and Lester, J. C. 2002. Narrative prose generation. *Artificial Intelligence* 139(2):213–252.
- Cassell, J., and Bickmore, T. 2003. Negotiated collusion: Modeling social language and its relationship effects in intelligent agents. *User Modeling and User-Adapted Interaction* 13(1-2):89–132.
- Cassell, J.; Vilhjálmsdóttir, H. H.; and Bickmore, T. 2004. Beat: the behavior expression animation toolkit. In *Life-Like Characters*. Springer. 163–185.
- Frost, D. M. 2013. The narrative construction of intimacy and affect in relationship stories implications for relationship quality, stability, and mental health. *Journal of Social and Personal Relationships* 30(3):247–269.
- Gratch, J.; Wang, N.; Gerten, J.; Fast, E.; and Duffy, R. 2007. Creating rapport with virtual agents. In *Intelligent Virtual Agents*, 125–138. Springer.
- Hovy, E. 1987. Generating natural language under pragmatic constraints. *Journal of Pragmatics* 11(6):689–719.
- Jefferson, G. 1978. Sequential aspects of storytelling in conversation. *Studies in the organization of conversational interaction* 219–248.
- Jurafsky, D., and Martin, J. H. 2000. *Speech & language processing*. Pearson Education India.
- King, A. C.; Bickmore, T. W.; Campero, M. I.; Pruitt, L. A.; and Yin, J. L. 2013. Employing virtual advisors in preventive care for underserved communities: results from the compass study. *Journal of health communication* 18(12):1449–1464.
- Kuyten, P.; Bickmore, T.; Stoyanchev, S.; Piwek, P.; Prendinger, H.; and Ishizuka, M. 2012. Fully automated generation of question-answer pairs for scripted virtual instruction. In *Intelligent Virtual Agents*, 1–14. Springer.
- Labov, W. 1972. *Sociolinguistic patterns*. Number 4. University of Pennsylvania Press.
- Li, B.; Lee-Urban, S.; Johnston, G.; and Riedl, M. 2013. Story generation with crowdsourced plot graphs. In *AAAI*.

- Lombardo, V., and Damiano, R. 2012. Storytelling on mobile devices for cultural heritage. *New Review of Hypermedia and Multimedia* 18(1-2):11–35.
- Nadeau, D., and Sekine, S. 2007. A survey of named entity recognition and classification. *Linguisticae Investigationes* 30(1):3–26.
- Norrick, N. R. 2009. Interjections as pragmatic markers. *Journal of pragmatics* 41(5):866–891.
- Ochs, E. 1997. Narrative. *Discourse as structure and process* 1:185–207.
- Porteous, J.; Cavazza, M.; and Charles, F. 2010. Applying planning to interactive storytelling: Narrative control using state constraints. *ACM Transactions on Intelligent Systems and Technology (TIST)* 1(2):10.
- Reiter, E.; Dale, R.; and Feng, Z. 2000. *Building natural language generation systems*, volume 33. MIT Press.
- Riedl, M. O., and Young, R. M. 2006. From linear story generation to branching story graphs. *Computer Graphics and Applications, IEEE* 26(3):23–31.
- Ring, L.; Shi, L.; Totzke, K.; and Bickmore, T. 2015. Social support agents for older adults: longitudinal affective computing in the home. *Journal on Multimodal User Interfaces* 9(1):79–88.
- Rühlemann, C. 2014. *Narrative in English Conversation: A Corpus Analysis of Storytelling*. Cambridge University Press.
- Sacks, H., and Jefferson, G. 1995. Lectures on conversation.
- Sina, S.; Rosenfeld, A.; and Kraus, S. 2013. Social narrative adaptation using crowdsourcing. In *CMN*, 238–256.
- Smith, J. 2000. *GrandChair: conversational collection of grandparents’ stories*. Ph.D. Dissertation, Massachusetts Institute of Technology.
- Swanson, R., and Gordon, A. S. 2012. Say anything: Using textual case-based reasoning to enable open-domain interactive storytelling. *ACM Transactions on Interactive Intelligent Systems (TiiS)* 2(3):16.
- Vardoulakis, L. P.; Ring, L.; Barry, B.; Sidner, C. L.; and Bickmore, T. 2012. Designing relational agents as long term social companions for older adults. In *Intelligent virtual agents*, 289–302. Springer.
- Ward, N., and Tsukahara, W. 2000. Prosodic features which cue back-channel responses in english and japanese. *Journal of pragmatics* 32(8):1177–1207.