

“How Was Your Weekend?” A Generative Model of Phatic Conversation

Hannah Morrison, Chris Martens

{hmmorris3, crmarten}@ncsu.edu
Principles of Expressive Machines Lab
Department of Computer Science
North Carolina State University
Raleigh, NC

Abstract

Unspoken social rules, such as those that govern choosing a proper discussion topic and when to change discussion topics, guide conversational behaviors. We propose a computational model of conversation that can follow or break such rules, with participant agents that respond accordingly. Additionally, we demonstrate an application of the model: the Experimental Social Tutor (EST), a first step toward a social skills training tool that generates human-readable conversation and a conversational guideline at each point in the dialogue. Finally, we discuss the design and results of a pilot study evaluating the EST. Results show that our model is capable of producing conversations that follow social norms.

Introduction

Social interactions are complex, each regulated by a set of social norms. In successful interactions, these norms are followed. However, these norms vary. What might be considered a violation of a social norm in one situation may be acceptable, and even expected, in another. Context, which determines the type of interaction, matters (Ogan et al. 2012). Social interaction is further complicated due to its reliance on social cues: facial expressions, body language, and other nonverbal behaviors contribute information to conversations (Giri 2009).

We are interested in building computational models of social-norm-following dialogue so that we can generate conversations and use them in order to simulate believable interactions between virtual agents. Given the inconsistency of conversational norms, we focused our attention on a common, yet complex, interaction: a short encounter between two acquaintances. We created a rule-based, generative model that simulates conversation between two agents, who may follow and violate social norms and generate emotional responses to the conversation.

This work contributes to the landscape of AI tools available for social skills training. We present a prototype social skills training tool (SSTT) as an example application of our model, called the Experimental Social Tutor (EST). We designed and ran a pilot study in order to evaluate the EST’s usability and quality of conversations generated by the model.

Results suggest that although the rendering of natural language between characters needs improvement, the EST’s model is able to generate conversations that follow social norms closely.

One of our long-term goals is to develop computational tools that can assist people learning social interactions and that may be useful for people on the autism spectrum. We view the EST as a promising first step towards this goal.

Background: Defining Conversational Rules

We seek to model interactions between two acquaintances. We base our model on the following rules, taken primarily from axioms of interpersonal interaction presented by Berger and Calabrese (Berger and Calabrese 1975), which represent social norms and attitudes: the goal of casual conversation among acquaintances is to improve rapport; following the social norms of greeting and exchanging pleasantries is favored because conversational participants know what to expect; participants ask more questions during the beginning of conversation, and offer more detail as conversation continues; participants value self-disclosure; participants favor mutual participation; participants view their conversational partner more favorably if they have similar interests or opinions (thus, participants should seek to find commonalities).

Berger and Calabrese state that the goal of conversation between two strangers is to reduce the amount of uncertainty between them—this is why unexpected conversational moves (i.e. social gaffes), such as dominating a conversation, make conversation less likely to continue. Because commonalities between conversational participants reduce uncertainty, participants prefer speaking with those that are similar to them. The importance of similarity among participants is echoed by Byrne and Griffitt (Byrne and Griffitt 1973), who conclude that attraction can be estimated by the number of similar opinions between the participants. They found that this holds among children, hospital patients, and the elderly: that is, that the importance of similarity between conversational participants is ubiquitous.

Similarity can be found through the establishment of the common ground between participants. Cassell et al. (Cassell and Bickmore 2003) state that familiarity among participants can be increased by talking about “topics that are obviously in the common ground such as the weather, physi-

cal surroundings, and other topics available in the immediate context of utterance.” Participants may also increase familiarity by introducing personal information, thereby moving it into the common ground. Zegarac (Zegarac 1998) states that the goal of phatic (small-talk-like) communication is to manage the common ground between participants.

We use the sources above to model socially normative conversation as conversation that engenders rapport among participants. Conversation that violates social norms, then, would decrease rapport. A deviation from the rules we present above would constitute this. For the scope of our task, however, we limit the social violations in the model to the following: participants that do not involve the other participant (that is, those that dominate their conversation), or participants that vocalize a strong disagreement in opinion, will decrease the chance that the conversation will continue, and lower rapport between them and their conversational partner.

Related Work

PsychSim, created by Marsella et al. (Marsella, Pynadath, and Read 2004), is a social simulation tool that can be used to generate a variety of social situations. Agents in the simulation may communicate with each other and each have their own goals, opinions, and beliefs regarding others. The purpose of this tool is for the user to explore social situations in a casual, low-pressure environment. While the tool we present in this paper is much simpler—the model’s agents do not hold beliefs about others beyond whether or not they are liked, nor do they have explicitly encoded goals—we see our simple, logic-based encoding as having several advantages, such as ease of authoring and extension, the ability to trace causal links between utterances, and the separability of the model from its specific implementation in a SSTT.

Bickmore and Cassell (Bickmore and Cassell 2005) describe an intelligent virtual agent, REA, capable of engaging in small-talk with a user. The agent’s conversational choices, and the underlying logic that motivates them, take into account the common ground between participants, the norm of staying on-topic in a conversation, and trust between participants. Our model uses similar social norms, but where REA uses small talk in the context of selling real estate, our model is general over topics and may be used to generate conversations for many settings.

Previous work by Even et al. (Even et al. 2016) introduced an interactive, turn-based SSTT for schizophrenia patients that can “teach the patients how to introduce themselves, start, maintain or interrupt a discussion, request or refuse something, receive or give a compliment, or criticism.” Tartaro and Cassell (Tartaro and Cassell 2007) describe a virtual peer with which the user can interact. Users engage in a ‘collaborative storytelling task’ with a life-sized avatar and may also manipulate the nonverbal behaviors of the virtual peer in order to “observe the effects on interaction.” Our proposed SSTT differs from these in several ways, most notably in its focus on the casual relationships between conversational events and our more explicit presentation of the social rules we aim to teach. While the EST may have the potential to be useful on its own, we envision the EST in its current state as



Figure 1: A diagram of the project’s architecture. The model produces an abstract conversation, which is then converted into text by a natural language generator (NLG) and guideline generator (GLG).

a complement to these avatar-containing SSTTs, being used, perhaps, prior to these in order to “warm up” participants for conversation with the virtual avatar.

The Model

We take inspiration from several existing models of social interaction. Richard Evans (Evans 2016), in presenting a model based on the Game of Giving and Asking for Reasons, describes agents that follow the conversational norm of turn-taking and are able to express emotions such as worry. Dixon et al. (Dixon, Smaill, and Tsang 2009) present a framework for modeling complex agents that introduces the concept of dialogue-related obligations, encoding the social norms of requesting and receiving information in an exchange between two individuals.

Our executable model generates conversations that may either follow or break social norms, with participant agents that respond accordingly. When the model is run, it non-deterministically generates one of many distinct conversations between two acquaintance agents. Agents must follow the norms of greeting one another and saying goodbye before exiting the conversation, must stay on topic, and must make small talk in some form before moving onto more varied topic discussion¹. Agents are, however, free to offend, bore, or annoy their conversation partner, and the partner, if made sufficiently upset, may leave the conversation. In order to simulate a time constraint, each utterance spends a turn. Conversation will transition to a natural ending after a set number of turns.

We formalize these conversational moves and conventions as an unordered collection of rules describing what is possible for the agents to do under which circumstances. Conversation states are encoded as collections of logical predicates: for example, when Alice becomes annoyed, the predicate `feels(alice, annoyed)` will be added to the conversational state.

We represent the following conversational states that may change as conversation proceeds:

- The current topic
- Each agent’s current feeling (happy, sad, annoyed, or content)
- The number of times each agent has spoken so far
- Affinity between agents (how much they like each other)

¹Initial topics are restricted to the *most* common ground (this idea is taken from (Cassell and Bickmore 2003)): the agent’s weekend or the weather; additional topics are sports (baseball, soccer, or running) and music (pop, country, or rock).




Figure 2: A finite state machine depicting the conversation flow constraints in the model. Each state represents a conversational rule (edges between the **make topic talk** state and **question** and **change topic** states have been omitted. Edges between the **continue talking** state and **question** and **change topic** states have also been omitted.) A socially inappropriate conversation might consist of the path **greet**→**make small talk**→**make topic talk**→**continue talking**→**continue talking**→**continue talking**→**continue talking**→**goodbye**, while a socially normative conversation might consist of the path **greet**→**make small talk**→**make topic talk**→**question**→**reciprocate question**→**change topic**→**topic talk**→**question**→**goodbye**.

In addition, we represent agents having opinions (positive or negative) about topics, and relatedness between topics. These may influence conversation, but are not themselves altered as conversation proceeds.

Conversational Moves

Ryan et al. (Ryan, Mateas, and Wardrip-Fruin 2016) present a procedural system of dialogue generation in which one line of dialogue is generated per conversational turn; each line of dialogue has a corresponding *dialogue move* or moves in the system. The EST works in a similar manner: each line of displayed dialogue corresponds to a rule in the model that has fired. The model defines dialogue moves for greeting, making small talk, making more advanced conversation (“topic talk”), continuing to speak on the current topic, asking a question, reciprocating questions (e.g. “What about you?”), responding with a typical level of enthusiasm and a high level of enthusiasm (e.g. “I like soccer” vs “I love soccer”), changing a topic, and saying goodbye. These may, in combination, produce conversation that is either socially appropriate or inappropriate; in either case, agents will respond emotionally. There are no restrictions on the social appropriateness of the generated conversations.

Conversation is constrained to follow a particular global structure so that, for example, participants may not continue talking before responding to a question, or change the subject without it being related to the current topic. We depict

Interpreted Utterance	Ceptre Rule	State Change
Bob: Good morning, Alice!	<code>greet_bob_alice</code>	<code>greet bob alice</code>
Alice: Good morning, Bob!	<code>greet_alice_bob</code>	<code>greet alice bob</code>
Bob: This weather today is really nice—good for playing sports	<code>small_talk_weather</code>	current topic: weather, times spoken by bob: 1
Bob: I did a lot of playing baseball on Saturday. It was nice out, just like today.	<code>change_topic_weather_baseball</code>	current topic: baseball, times spoken by bob: 2
Bob: I think baseball is a lot more interesting than people give it credit for.	<code>topic_talk_baseball_typical_positive</code>	times spoken by bob: 3
Bob: Some of the people I know like baseball.	<code>continue_talking_baseball</code>	times spoken by bob: 6, Alice feels annoyed
Alice: Uh-huh, well...I have to go now. Goodbye.	<code>annoyed_by_unfair_participation_terminate_conversation</code>	Alice feels annoyed, conversation is ending
Bob: Take care.	<code>say_goodbye_bob_alice</code>	<code>say_goodbye_bob_alice</code>

Figure 3: An example of an abbreviated model-generated sequence of statements. Some of Bob’s statements have been omitted.

the global conversation structure as a state machine in Figure 2, which approximates these constraints as precedence ordering between conversational moves.

Here are two examples of the conversational moves we implement:

- **Make topic talk:** If the current topic is T and agent C has an opinion O about T , then C states that their opinion about T is O (in either a typical or enthusiastic way). Increment the number of times C has spoken.
- **Reciprocate question:** If C has just asked C' a question about topic T , C states their opinion on T and asks the same question in response. Increase affinity between C and C' and increment the number of times C' has spoken.

Emotional and Affinity Changes.

Rules that can change agent emotions and affinity between agents include:

- **Like from agreement:** If C has just stated their opinion on topic T , and C' shares this opinion, increase affinity of C' towards C .
- **Annoyance from unbalanced participation:** If C is feeling content, but the number of times C' has spoken is more than $2/3$ the total conversation length, change C ’s feeling to annoyed.

Implementation

We implemented the model using the linear-logic-based modeling language Ceptre (Martens 2015), which allows for a declarative representation of conversational moves. Each move is represented with a name, a set of preconditions, and a description of how it modifies the conversational state. The set of moves is unordered, but ordering constraints may be

See a Conversation:

The goal of casual conversation is to make people more comfortable around each other. If people are comfortable with each other, they're more open to conversation as the rapport between them increases. Therefore, generally, you want to avoid controversial or highly-specific topics until you have a higher level of rapport. This simulation consists of 2 classmates talking one morning before class, and assumes the conversants only know each other casually, so for now you are only able to "discuss" benign topics.

Press the "Render" button to generate a new conversation.

Refresh to clear the conversations from the screen.

Render

Figure 4: A partial screenshot of the EST

introduced through preconditions, e.g. to enforce the structure described in Figure 2. Once the author provides an initial state (a set of facts in first-order logic), Ceptre will run the state forward according to the moves that apply, choosing nondeterministically whenever multiple moves apply, and stop when no more moves apply. The nondeterminism inherent in Ceptre can generate highly varied conversational output. Figure 3 shows an example conversation generated by the model.

The Experimental Social Tutor

Socialization is sometimes problematic for people on the autism spectrum. Autism is a neurodevelopmental condition characterized by difficulty interacting and communicating with others (e.g. having trouble interpreting facial expressions) (Tartaro and Cassell 2007) as well as restricted, repetitive behavior (e.g. becoming anxious if one's routine is not followed) and interests (e.g. having an intense focus on a special interest or topic) (Lam and Aman 2007). While those on the spectrum can often identify nonverbal and social cues in isolation, they may have difficulty interpreting them in practice, which can lead to problems such as loneliness, social isolation, and avoidance of social interaction (Myles 2003). We are interested in developing computational tools that may be useful in helping people on the spectrum learn about social interactions. The EST is an early prototype of such a tool.

The EST separates the concerns of structure and presentation of conversation. In addition to our rule-based *model* of social interaction described above, which generates the abstract structure of conversation in terms of conversational moves, we implemented a front-end *interface* to render and contextualize the conversations. The front end transforms the abstract representation into concrete dialogue and an accompanying social guideline. This text is displayed to the user by the EST, our early-stage, web-based SSTT. The architecture of our system is depicted in Figure 1.

Dialogue and Guideline Generation

The EST uses conversations generated by the model described above to display a natural-language rendering of the conversation and relevant conversational guidelines. It is intended for middle-and-high-school aged people on the autism spectrum. Figure 4 displays the prompt given to the user when beginning use of the application. Figure 5 displays sample output of the EST.

Dialogue	Guideline
Bob: Good morning, Alice.	Good places to start a conversation: waiting in line, a club meeting, on the bus.
Alice: Good morning, Bob	Greeting someone acknowledges them and lets them know you are open to conversation.
Alice: I love time at home listening to my music. Wish the weekend didn't go by so quickly	Small talk makes people more comfortable around each other.
Bob: How was your weekend, Alice?	Avoid only asking questions and never giving information about yourself; this makes the conversation one-sided.

Figure 5: A table depicting a partial example of EST output.

Dialogue Generation Our dialogue generator takes as input a trace from Ceptre abstractly representing a conversational exchange, and produces human-readable text output for that conversation. It maps each step of the trace to a line of dialogue, where the rule in the model that determined the step determines an utterance to represent it. In order to increase diversity of phrasing, some cases have multiple potential utterances, which are displayed at random.

Guideline Generation As with agent dialogue, each guideline is generated by case-analyzing the rule used in the model. In order to communicate numerous relevant ideas, some cases have multiple potential guidelines, which are selected to appear at random.

Guidelines were gleaned from several sources. Our primary reference was a social-skills book written by an autistic adult (Wendler 2014). We also referenced several guides describing therapeutic sessions for social-skills improvement (Laugeson and Frankel 2011; Myles 2003) and a guidebook intended for autistic teenagers (Patrick 2008). Rules in the model that involved the initial stages of conversation were matched with guidelines about when and where to initiate conversation, as well as the importance of beginning conversation; rules involving subsequent dialogue moves were matched with guidelines about the importance of making one's conversational partner comfortable through near-equal exchange of information, as well as potential conversation topics; rules involving questions were matched with guidelines explaining the role of questions as signals of interest, in addition to advising how often questions should be asked.

Experiment

Next, we describe the design and results of a pilot study evaluating the EST.

Design

Because the EST is in an early stage of development, our primary goal in study design was to evaluate the model's capability of generating conversation that follows social norms. The questions asked included three 0-10 scales asking participants to judge difficulty of using the EST, realism of the




Figure 6: Histograms depicting the range of participant ratings of social norm adherence, conversation realism, and usability, on a scale from 0 to 10.

conversations, and how closely the conversations follow social norms.

Due to the early stage of development of the EST, we did not test the EST’s intended users: middle-and-high schoolers on the spectrum. However, we were able to get feedback from 3 autistic college-aged participants. Participants (N=35) were recruited through advertising (flyers, in-class announcements, etc.) on campus. The study was conducted entirely online: participants were e-mailed instructions that included a link to the tool and a link to the survey (which, when completed, re-directed to another survey where participants could enter their e-mails for compensation purposes, in order to ensure participant anonymity). Participants were e-mailed \$10 Amazon gift-cards for their participation.

Results

Most participants found the tool easy to use and thought that the generated conversations followed social norms; however, they found the dialogue generated by the tool to be awkward and repetitive. See Figure 6 for histograms of ratings for all three questions asked.

Social Norm Adherence Most participants rated the generated conversations as adhering to social norms closely (median: 8). Many mentioned² the one-sided conversation between the two agents, stating “People will most likely back away from a conversation if it is one sided or confusing” and “The way they awkwardly end the conversation if the other person is talking too much is really good.”

Realism of Generated Conversation Most participants found the “speech” between the two agents to be moderately realistic (median: 6). Some participants felt the speech between the classmates was “too formal,” or “stiff and unnatural.” Despite this, the dialogue still evoked emotional

²For our purposes, ‘mentioned’, ‘said’, etc. refer to the text response that was submitted by the participant.

responses from users. One participant said, “[The conversation] reminds me of a casual conversation I have with someone like a study partner.” One participant felt an emotional reaction in response to the conversations: “Honestly I had some pain from seeing social norms breached by one person. It was like watching a movie where someone messes up really badly.”

Usefulness Interestingly, several neurotypical participants found the EST educational, stating that it prompted them to think about conversation in a way they typically do not. One participant said “the ‘guideline’ [sic] section of the tool described things I do in conversation that I didn’t even realize I was doing.” Another said “The ‘guideline’ part made me think about how some people...would view social interaction as a set of rules and steps instead of something that happens somewhat naturally and organically.”

One autistic participant said that, while the user experience needed significant work, “the underlying principles are sound.” Another said that they had trouble applying social rules rather than identifying them, so “adding more detail to the guidelines or more specific strategies might be more helpful.” (We discuss a response to this in the Future Work section.)³ Another participant said that the guidelines given by the tool were not useful because “I have done a lot of studying on [how you are supposed to act during a conversation].” We think the EST might have the potential to be helpful during the studying and learning of basic social rules that the participant mentions, although the tool is too basic for use by most adults. Considering the feedback from the first participant we described, we think the EST might be useful for adults if we add more complex interactions; however, we do think it has the potential to be useful for its intended audience (those of middle-school and high-school ages) with some modifications. Most autistic participants expressed that, due to difficulties with nonverbal communication, a video-based component would have been helpful. This supports our vision of the EST in its current state being used in combination with a SSTT that makes use of a virtual avatar.

Summary

We have presented a model capable of generating varied, norm-adhering conversation between two acquaintance agents. We used a declarative, rule-driven system to encode conversational moves, which supports compositional dialogue generation and the ability to track causal relationships between changes in the conversational state. The EST is an early-stage social skills training tool that is based on our model. The results of a pilot study suggest the model on which the EST is based generates conversation that follows social norms closely. The model on which the EST is based can easily be extended.

³While we intended for the EST’s conversations to act as examples of applications of these rules, they were generally too unrealistic to be helpful in this regard.

Future Work

The ultimate goal of future work is to create an interactive environment that uses our model of conversation, allowing the user to both observe casual conversations between simulated agents and engage in interactions with them, while receiving useful conversational guidelines. Such agents might also be useful for increasing the social believability of characters in games (Morrison and Martens 2017).

We plan to improve the natural-language rendering of the conversations generated by the EST by integrating stylistic variations and studying how these affect perception of the conversation. We would also like to add *explanations* for the causal relationships between lines of dialogue and the “hidden variables” in the conversational state, such as agent emotions. Finally, we plan to integrate more context and reasoning into the model, such as differentiating between conversations between participants with different relationships (e.g. a peer vs. an authority figure).

References

- Berger, C. R., and Calabrese, R. J. 1975. Some explorations in initial interaction and beyond: Toward a developmental theory of interpersonal communication. *Human communication research* 1(2):99–112.
- Bickmore, T., and Cassell, J. 2005. Social dialogue with embodied conversational agents. *Advances in natural multimodal dialogue systems* 30:23–54.
- Byrne, D., and Griffitt, W. 1973. Interpersonal attraction. *Annual review of psychology* 24(1):317–336.
- 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):89–132.
- Dixon, L.; Smaill, A.; and Tsang, T. 2009. Plans, actions and dialogues using linear logic. *Journal of Logic, Language and Information* 18(2):251–289.
- Evans, R. P. 2016. Computer models of constitutive social practice. In *Fundamental Issues of Artificial Intelligence*. Springer. 389–409.
- Even, C.; Bosser, A.-G.; Ferreira, J. F.; Buche, C.; Stéphan, F.; Cavazza, M.; and Lisetti, C. L. 2016. Supporting social skills rehabilitation with virtual storytelling. In *FLAIRS Conference*, 329–334.
- Giri, V. N. 2009. Nonverbal communication theories. *Encyclopedia of Communication Theory* 690–694.
- Lam, K. S., and Aman, M. G. 2007. The repetitive behavior scale-revised: independent validation in individuals with autism spectrum disorders. *Journal of autism and developmental disorders* 37(5):855–866.
- Laugeson, E. A., and Frankel, F. 2011. *Social skills for teenagers with developmental and autism spectrum disorders: The PEERS treatment manual*. Routledge.
- Marsella, S. C.; Pynadath, D. V.; and Read, S. J. 2004. Psychsim: Agent-based modeling of social interactions and influence. In *Proceedings of the international conference on cognitive modeling*, volume 36, 243–248.
- Martens, C. 2015. Ceptre: A language for modeling generative interactive systems. In *Eleventh Artificial Intelligence and Interactive Digital Entertainment Conference*.
- Morrison, H., and Martens, C. 2017. A generative model of group conversation. In *Proceedings of the 12th International Conference on the Foundations of Digital Games*, 66. ACM.
- Myles, B. S. 2003. *Social skills training for children and adolescents with Asperger syndrome and social-communication problems*. Autism Asperger Publishing Company.
- Ogan, A.; Finkelstein, S. L.; Walker, E.; Carlson, R.; and Cassell, J. 2012. Rudeness and rapport: Insults and learning gains in peer tutoring. In *ITS*, 11–21. Springer.
- Patrick, N. J. 2008. *Social skills for teenagers and adults with Asperger syndrome: a practical guide to day-to-day life*. Jessica Kingsley Publishers.
- Ryan, J. O.; Mateas, M.; and Wardrip-Fruin, N. 2016. A lightweight videogame dialogue manager. In *DiGRA/FDG*.
- Tartaro, A., and Cassell, J. 2007. Using virtual peer technology as an intervention for children with autism. *Towards universal usability: designing computer interfaces for diverse user populations*. Chichester: John Wiley 231:62.
- Wendler, D. 2014. *Improve your social skills*.
- Zegarac, V. 1998. What is “phatic communication”? *Pragmatics AND Beyond New Series* 327–362.