

# SOGO: A Social Intelligent Negotiation Dialogue System

Ran Zhao  
Carnegie Mellon University  
Pittsburgh, PA  
rzhao1@cs.cmu.edu

Oscar J. Romero  
Carnegie Mellon University  
Pittsburgh, PA  
oscarr@cs.cmu.edu

Alex Rudnicky  
Carnegie Mellon University  
Pittsburgh, PA  
air@cs.cmu.edu

## ABSTRACT

In this paper, we propose a semi-automatic social intelligent negotiation dialogue system that interweaves task utterance with conversational strategies to engage human users in negotiation. Our two-phase system operates sequentially in a reasoning-and-generation loop: In the task phase, we leverage an off-the-shelf end-to-end dialogue model for negotiation to build a dialogue manager which decides the next system's task intention. Then, during the social phase, we employ a theory-driven, template-based natural language generator to realize the task intention as a genre of social conversational strategy. Subsequently, a set of conversational strategies are presented to a human expert who decides the final sentence to be uttered by the dialogue system. Compared to the baseline system, our proposed social intelligent dialogue system achieves a higher agreement rate and more "good deals" with humans while building interpersonal rapport.

## KEYWORDS

negotiation, dialogue system, conversational strategy, social reasoning, natural language generation

### ACM Reference format:

Ran Zhao, Oscar J. Romero, and Alex Rudnicky. 2018. SOGO: A Social Intelligent Negotiation Dialogue System. In *Proceedings of ACM IVA conference, Sydney, Australia, Nov 2018 (IVA'18)*, 9 pages. <https://doi.org/10.1145/nnnnnnn.nnnnnnn>

## 1 INTRODUCTION AND MOTIVATION

Whether we are deciding between a salad or fast food for lunch or asking a coworker to complete a project ahead of schedule in exchange for help at a later date, people negotiate every day. When we hold conflicting interests, we must negotiate to pursue our ultimate goals. Thus, negotiation is an act wherein participants with unique motives to cooperate and compete to maximize their own benefits. Virtual agents are a powerful tool for teaching negotiation skills and modeling negotiation in an agent shows great promise in a variety of domains [18]. Indeed, many current empirical studies are making progress in this area [14, 19, 23, 31].

Drawing on this work, we recognize that negotiation is both a challenging reasoning problem as well as linguistic problem. Although people are adept at navigating the trade-off between cooperation and competition, algorithms have yet to develop such reasoning and linguistic fluency. Therefore adversarial agents, such as AlphaGo, which aim to beat interlocutors in a zero-sum game, cannot work because they privilege competition, not cooperation. Similarly, a personal assistant like Alexa or Siri, which privileges cooperation, overlooks competition. Further, most studies regard negotiation as a sole reasoning or planning problem, like searching for optimized outcomes and thus aim to sharpen the agent's tactics. However, negotiation also relies on appropriate language to maintain relationships and optimize a plan. Prior research in human-human negotiation [22, 32] has shown that social factors such as trust and rapport underlie both challenges. In this study, we leverage different linguistic devices to build rapport between a human and a dialogue system, which thus fosters integrative agreement during negotiation.

Negotiators are encouraged to share crucial information and cooperate to reduce the risk of impasse because of a sense of *rappor*t, a feeling of connection and closeness with another [6, 32]. Also, Curhan et al., [8] shows that rapport helps to formalize a negotiator's intuition about objective outcomes and predict future objective value. Social scientist Spencer-Oatey [36, 37] explains the experience of establishing rapport according to three interrelating components: face, social rights and interactional goals. People use a variety of strategies to manage these three factors, which have been categorized into five major domains [37]: illocutionary domain, discourse domain, participation domain, stylistic domain and non-verbal domain. Most recent studies of rapport agents [17, 20] have investigated the importance of the non-verbal domain in human-agent interaction. However, we are interested in examining how the verbal channel produces rapport in the context of human-agent negotiation. Previously, Bronstein et al., [6] employed a hierarchical linear model to validate and quantify the contribution of the verbal channel in rapport management.

In general, we believe that introducing social conversation into complex negotiation communication will help establish and maintain rapport while facilitating negotiation process. To this end, we propose a two-phase computational model for negotiation: the task phase and the social phase. The task phase generates the next system task intention/move (e.g., to request a book). The social phase provides opportunities for social intentions/moves (e.g., self-disclosing a personal preference) realized by different conversational strategies. Conversational strategies are units of discourse that are larger than speech acts, which have been demonstrated to contribute to building, maintaining or even destroying interpersonal (or human-agent) bonds [35, 42]. The social phase is inspired by the work [2, 29], who exploits structures of casual conversation

---

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from [permissions@acm.org](mailto:permissions@acm.org).

IVA'18, Nov 2018, Sydney, Australia

© 2018 Association for Computing Machinery.

ACM ISBN 978-x-xxxx-xxxx-x/YY/MM...\$15.00

<https://doi.org/10.1145/nnnnnnn.nnnnnnn>

in human-human communication to improve the system's capabilities of generating task-related social moves that are more than idle chit-chat. The major function of task-related social moves is to prepare the interlocutor for the next negotiation move—rather than increase familiarity between interlocutors—by discussing some restricted common topics like the weather. For instance, the system discloses its personal interest in reading as the current social move (e.g., "I love reading"), which anticipates its next negotiation move of requesting the book (e.g. "Can I have the book?").

Our work does not seek to develop negotiation tactics but rather to use rapport management strategies to facilitate the negotiating process. Thus, we do not develop our own negotiation reasoning module but adopt and extend the end-to-end negotiation dialogue model from Facebook AI Research (FAIR) group as our task reasoning backbone [25]. FAIR has demonstrated the advancements of their neural dialogue models on the multi-issue bargaining task. Their task is to divide items from three item-types (books, hats, balls) between two participants. Since their neural dialogue model allows users to negotiate with the system through natural language, we are able to insert social moves into dialogues. Further, FAIR leveraged reinforcement learning to maximize the outcome of negotiation rather than imitating the actions of human users in supervised learning. Finally, FAIR replaced traditional likelihood decoding with a rollout algorithm, a kind of beam search conducted in the dialogue turn level, which allows the model to generate more effective negotiation strategies. Despite these advantages, FAIR suffers from low agreement with human users because of users' unwillingness to cooperate with such an uncompromising opponent. We believe that combining the negotiation tactics of the Facebook end-to-end model (FB) and our proposed rapport strategies will allow our social intelligent negotiation system (SOGO) to overcome these limitations and improve task performance while constructing an interpersonal bond with humans. To compare outputs from both systems, please refer to the sample dialogues in Table 1.

#### FB System

**Agent:** I'll take the book and balls if you want the book.

**User:** I need all the balls and the hat

**Agent:** You can have the hat if i can have the rest

**User:** I prefer all the balls and you can have the rest

**Agent:** deal

#### SOGO System

**User:** Since you really liked the ball last time, I am giving 2 to you. I get the hat and the books. Deal?

**Agent:** Thank you so much! 😊 I guess the book must be very important to you 😊. Shall I give you the the balls and 1 book?

**User:** You get the balls, and I get 2 books and the hat. Deal?

**Agent:** Why are you so good at this?! Seriously! ball and hat! Great!

😊 ..., can I have one ball and one hat please?

**User:** okay I guess.

**Agent:** Playing with you is making me a smarter chatbot. 😊

**Table 1: Sample human-agent negotiation dialogue of the FB System (top) and our SOGO System (bottom)**

In the next section we review prior studies related to negotiation and rapport agents. Then section 3 introduces our computational model of negotiation and section 4 provides an overview of the system architecture. To validate our model, we explain the set up

of our experiment in section 5. Finally, in section 6, we combine subjective metrics and objective metrics to evaluate the system's performance in negotiation and rapport building.

## 2 RELATED WORK

Below we describe related work on building a negotiation agent and a social intelligent agent.

### 2.1 Negotiation Agent

Modeling negotiation in an agent has become an increasing area of research. Different negotiation frameworks[13, 19, 30] have been established to serve as benchmarks for evaluating practical negotiation strategies. For instance, the generic negotiation framework Genius [19] facilitates research in bilateral multi-issue negotiation. Alternately, IAGO [30] allows a human to negotiate with a multi-modal virtual agent, whose facial expressions and nonverbal cues are accessible to its negotiating partner. Unlike the single-agent Genius and IAGO frameworks, DipGame [13] is a multi-agent system developed on the Diplomacy Game, where negotiation and relationships between players are essential for success. These frameworks provide testbeds for researchers to investigate different negotiation tactics and social strategies in human-agent interaction.

Prior research indicates that modeling negotiation in virtual humans requires the system to demonstrate cognitive skills for reasoning and social skills for communication, as well as a means to express its internal state [16]. Gratch and his colleagues introduced a Conflict Resolution Agent (CRA) that allows students to engage with virtual human role-players across multi-issue bargaining problems. They showed that participants perceived the same levels of satisfaction, cooperation and rapport when paired with a CRA as with a human. Furthermore, they confirmed that the CRA could help people improve their negotiation skills. To build on these promising findings, subsequent research has investigated ways to design an agent with better communication and negotiation skills. For instance, DeVault et al., [9] employed a wizard-controlled system to improve turn-taking skills, and Lucas et al., [26] investigated deceptive strategies to optimize one party's benefit while preserving the illusion of fairness towards the other party to improve negotiation policy. The analysis demonstrated that deceptive strategies increase feelings of satisfaction, trust and fairness of the offer from one's partner, which seems to be more credible and even preferred to honesty. Meanwhile, this work also ameliorates an agent's capability to understand human behavior by identifying multimodal signals of trustworthiness in human-human negotiation (i.e., different sets of behavior indicators for predicting objective trustworthiness and perceived trustworthiness). This finding will enable the system to infer human negotiation behaviors and display specific behavioral signals on demand to manipulate trustworthiness. Additionally, Johnson et al., [21] developed an autonomous agent that teaches people to negotiate through visualized feedbacks. Basically, the researchers encoded theoretical negotiation principles into several automatic quantifiable metrics that have been validated in pedagogical negotiations.

Current researchers have moved away from a traditional expert system to a reinforcement learning branch to design the algorithm for a negotiation agent. This new method promises to build a fully-automatic agent. The work on which our task phase is based [25]

introduces an end-to-end model that employs a reinforcement learning algorithm and that was shown to have learned both linguistic and reasoning skills without human annotations. However, this model is perceived to be too uncompromising and aggressive. Our work addresses this problem by building interpersonal rapport through conversational strategies, which could increase the human user's tolerance to one's agent partner. Alternately, Sun et al., [38] suggests designing a benevolent agent based on a reward-shaping method that diminishes rewards that make the agent feel less satisfied for consecutive rewards. We plan to integrate this method into our system in future work. In addition to the reward-shaping method, Lerer et al., [24] tried to solve social dilemmas by exploiting the tit-for-tat (TFT) strategy. Based on the definition of Markov TFT, the researchers utilized a reinforcement learning method to approximate the cooperative and defect policies as well as the switching policy. This work encourages the agent to become more adaptive to solve social dilemmas in many environments.

Although promising, this prior research is limited by a focus on modeling negotiation in a single interaction; they do not leverage the collected knowledge across the interaction to facilitate the process. Also, these studies solely consider and optimize the negotiation outcome but ignore the social outcome and its contributions to the task. Furthermore, most negotiation agents communicate with users through pre-defined symbols, not natural language, which detracts from the naturalness of the interaction and removes social cues from language.

## 2.2 Social Intelligent Agent

Many prior studies have addressed the issue of rapport between human users and virtual agents. An early paper [2] implemented a Real Estate Agent to build trust with users while selling houses. Based on sociolinguistics and social psychology, its computational architecture blended small talk and task talk in conversation, and demonstrated that social dialogue affects trust for extroverted users. Elsewhere, Coon et al., [7] created an algorithm to plan appropriate activities based on the user-agent relationship, ranging from stranger to companion. Each activity had its own required closeness level, and an activity planner optimized its activities to increase intimacy over time. Similarly, Bickmore et al., [3] categorized relationships into four intimacy levels and defined a set of dialogue acts for each level. Inspired by accommodation theory, the researchers modeled long-term user-agent relationships in the context of health counseling. Their proposed reactive algorithm selects the most appropriate dialogue act for the assessed intimacy level yet fails to increase intimacy. Gratch et al., [17] offers an alternative way of building rapport. They developed a sensitive listener agent that generates listening feedback behaviors to effectively coordinate with human users. In subsequent studies, Huang et al., [20] proposed a next generation rapport agent capable of enhancing mutual attention by predicting backchannels, increasing coordination by generating smooth turn-taking behaviors and maintaining positive communication through affective response and facial expressions. In addition to the nonverbal channels, Matasuyama et al., [28] also proposed several verbal strategies to facilitate rapport. They presented a socially-aware animated intelligent personal assistant that analyzes users' multimodal behaviors to estimate rapport levels,

reasoning and generating the conversational strategy instantiated by a set of observable verbal and nonverbal actions to build rapport.

Similar to research in negotiation agents, these work do not model the user across interactions and use the shared knowledge to build interpersonal bonds.

## 3 COMPUTATIONAL MODEL

In the following, we outline the details of our two-phase formalism in a computational model of negotiation. The task phase and the social phase are performed sequentially. Our task phase adopts the Facebook end-to-end negotiation dialogue model [25] which decides the system's next-task utterance. To make the system seem more human, in the social phase, our model displays all eligible strategies and realizes them into utterances that concord with former task utterances. These concordances are based on deep understanding of user's prior utterance and the system's next task utterance.

### 3.1 Task Phase

*3.1.1 End-to-end dialogue model.* Lewis et al., [25] utilize two-stage learning strategies by pre-training the model with supervised learning, then fine-tuning the parameter using reinforcement learning. In the following, we briefly discuss the advanced reinforcement learning with dialogue rollouts decoding based model from [25], which we deploy in our study. Each dialogue  $D$  is represented as a set of token  $x_t$  where the total number of tokens are  $T$ . Tokens are segmented by two special tokens WRITE and READ which indicates turn-taking between human and agent. The agent has an input goal and generates the outcome of the negotiation  $o$ . The model structure is a four GRU-based recurrent neural networks:  $GRU_g$  (Agent's goal encoder),  $GRU_w$  (dialogue token generator),  $GRU_o$  (forward output encoder),  $GRU_b$  (backward output encoder). In the first stage of supervised pre-training, given the word embedding  $W$ , Lewis et al., [25] firstly models the dependencies between language and input goals with the function (1):

$$p_{\theta}(x_t | x_0 \dots t-1, g) \quad (1)$$

Conditioning the input goals and generated dialogue, they predict negotiation outcomes with the function (2)

$$p_{\theta}(o | x_0 \dots T, g) \quad (2)$$

Thus, the objective function in the supervised learning stage can be represented as:

$$\mathcal{L}(\theta) = - \sum_{x,g} \sum_t l_o p_{\theta}(x_t | x_0 \dots t-1, g) - \sum_{x,g,o} l_o p_{\theta}(o | x_0 \dots T, g) \quad (3)$$

is a hyperparameter to balance token prediction loss and outcome prediction loss. Based on the negotiation outcome at the end of each dialogue, the agent receives a reward  $r(o)$ . In the second stage of reinforcement learning, given the discount factor  $\gamma$  and a running average of completed dialogue rewards  $\mu$ , the objective is to optimize the expected reward of each token generated by the agent ( $\mathbb{E}$ ) as follows:

$$L_{\theta}^{RL} = \mathbb{E}_{x_t \sim p_{\theta}(x_t | x_0 \dots t-1, g)} \left[ \sum_{x_t \dots T} \gamma^{T-t} (r(o) - \mu) \right] \quad (4)$$

In the decoding part, dialogue rollout algorithm [25] generates a small set of candidate utterances  $U = \{u_i | u_i = x_{n, n+k}\}$  and

chooses the utterance that maximizes the expected reward, with the following function:

$$u^* = \underset{u_i = x_{n, n+k}, u_i \in U}{\text{arg max}} (\mathbb{E}_{x_{(n+k+1 \dots T, o)} \sim p_\theta} [r(o)p_\theta(o|x_{0 \dots T})]) \quad (5)$$

Finally, both  $u^*$  and current turn user input utterances are sent to our social phase.

### 3.2 Social Phase

Our social phase transforms task utterances by introducing social moves. To effectively plan these moves, we need two types of information stored in our defined *conversational state*: (a) user/agent task intention/move and (b) user/agent model (e.g. personal preferences, dialogue history and context). Based on those information, our social language generator realizes the social moves with eligible conversational strategies. To obtain those information on the fly, we first apply a text classification algorithm to understand task intention from the utterance. Secondly, we use traditional information extraction method to construct a user/agent model. Finally, drawing on socio-psychology theories, we defined nine conversational strategies, which have a property of pre-conditions, from which we determine eligibility of specific conversational strategies given the conversational state. Meanwhile, we deploy emoticons to our generated sentences as indications of the illocutionary force in the textual utterances that they accompany.

**3.2.1 Intention Recognition.** In our study, understanding user and agent task intention/move is the premise for transforming the task utterance into conversational strategy. Based on our definition, each task intention/move consists of one speech act (e.g., Request, Offer) and one or more affiliated entity mentions (e.g., two books). We leverage vector-based text representation to build a speech act classifier and utilize a keyword matching algorithm to extract the entities mentioned in the sentence. Following these, the challenge might be multiple intentions in one utterance. For instance, "If you give me the ball, I will give you the book and two hats" refers to both Offer and Request. Thus, it is difficult for us to link the entity mentioned to its affiliated speech act. Our solution is to utilize the Stanford CoreNLP toolkit [27], which breaks the utterance into separate clauses with the smallest grammatical unit that can express a complete proposition, before training our speech act classifier. In this way, we guarantee that each clause includes only one intention. Both the human annotation and the trained classifier below are in the clause level.

**Speech Act Annotation** Based on empirical studies of human-agent negotiation dialogues [9, 16, 23], we discovered the five speech acts most closely related to rapport/face management and being widely used in negotiation shown in Table 2. The annotation work was conducted on the Amazon Turk platform. Six out of ten MTurkers passed the qualification test, i.e., completed previous tasks with more than 80% accuracy. 2,500 dialogues were annotated and used to train our speech act classifier, which served to annotate the rest of the corpus.

**Speech Act Classifier** We leverage the sentence classification library fastText [5] to train our supervised speech act classifier. fastText is essentially an extension of the word2vec model, which treats each word as a composition of character n-grams. We have set  $n$  as in the range between 3 and 6 (default value). We chose the

Speech Act	Precision	Recall	F1
Request	0.922	0.935	0.928
Reject	0.824	0.740	0.780
Accept	0.826	0.858	0.842
Elicit preference question	0.776	0.712	0.743
Offer	0.913	0.859	0.885

**Table 2: Performance of five speech act classifiers**

fastText toolkit for two reasons: (1) The Facebook negotiation corpus is domain-specific with a small vocabulary. Out-of-vocabulary presents a considerable problem. fastText generates better word embeddings for rare words and even those out-of-vocabulary words since it constructs the embeddings in the character-level. (2) fastText is memory-consuming: the number of n-grams in the character-level grows exponentially with the growth of corpus size. Since we have a small vocabulary in the corpus, this is not the case.

**3.2.2 User Model and Agent Model.** Both user model and agent model contain the dialogue history and context across the interactions that serve as long-term memory in human-agent interaction. Meanwhile, this memory offers the dialogue content for specific conversational strategies (e.g., reference to shared experience) that could index their built relationship. User model and agent model share most parts of the schema: (1) preferences, (2) historical game results (e.g. scores, deal items, game context), (3) speech act sequences, and (4) sentiment sequences. Besides, agent model also includes the conversational strategy sequences. In order to obtain the key information of user's preference in real-time, we developed a syntactic-based preference extractor. Specifically, in the clause breaking process, Stanford CoreNLP pipeline [27] generated a dependency tree of each clause as one of the intermediates. Thus, we wrote several subject-verb-object (SVO) templates to extract user preference on the dependency tree. In terms of the user's sentiment, we utilize the off-the-shelf rule-based sentiment classifier [15].

**3.2.3 Social Language Generator.** In this part, we adopt a theory-driven template-based approach to generate social moves. Zhao et al., [43] proposed a computation model of rapport that explains how humans in dyadic interactions build rapport over time through conversational strategies. Specifically, they find four major conversational strategies that positively impact rapport: Self-Disclosure (**SD**), revealing personal information to decrease social distance; Reference to Shared Experiences (**RSE**), which indexes common history; Praise (**PR**), which increases self-esteem in the listener and therefore raises interpersonal cohesion; and Violation of Social Norm (**VSN**), where general norms are purposely violated to accommodate the other's behavioral expectations. However, the authors studied peer tutoring-a scenario that elicits far fewer face-threatening speech acts, such as requests or rejections, than do negotiation dialogues. Thus, we are motivated to add speech act strategies which could boost politeness and appropriately address face-threatening speech acts. Concretely, based on [1, 4, 11], we include **Request**, **Reject**, **Gratitude**, **Greeting** and **Closing** strategies, each of which contains several sub-categories. For instance, head act is a core part of a request sequence. We tried to mitigate its face-threatening effect through different supportive moves: (1) *Preparator*: "I'd like to ask you something..." (2) *Grounder*: "I missed my book so much" (3) *Promise of reward*: "I will give you all the books in the next game." (4) *Imposition downgrader*: "Could you

Strategy	Sub-category	Realization
SD	Inner state	You know what, I really love reading.
RSE	Preference	Books are for you since you said you love reading last time
PR	Interaction	Negotiating with you is such fun.
VSN	Teasing	You messed up my thinking my friend....
Request	Grounder	Tomorrow is my creator's birthday and I do not get time to buy him a gift. Could you please give me the books for him?
Reject	Conditional	If you'd told me earlier, I could have given you the books.
Gratitude	Appreciation	You are a life savor!
Greeting	Friend	It is always a pleasure to play with you
Closing	SD Closing	Besides the game, I look forward to getting know you better

Table 3: Strategy Realizations

please give me the ball if you are not playing with it now?". We acquired several variations in sentence realization for each sub-category by hiring two native English writers. Table 3 shows some examples. Notice that some templates are designed for a specific negotiation entity (in red) and others are more general (in blue).

**Preconditions** Each strategy contains several pre-conditions that decide eligibility of usage given the current conversational state. In our model, pre-conditions are divided into categories, as shown in table 4.

Category	Pre-conditions
System and User Speech Act	Request, Reject, Greeting, Closing, Gratitude
Sentiment	positive and negative
Dialogue History	time of interaction, number of turns, historical game results
Entity	book-slot, hat-slot, ball-slot
System conv.strategies	SD, RSE, VSN, PR

Table 4: Preconditions

**Emoticons** Since visual access between participants in this study was limited, we substituted non-verbal cues with emoticons. Emoticons are generally accepted as non-verbal indicators of emotions that map directly onto facial expressions [34, 40], yet they also indicate the illocutionary force of an utterance [10]. They do not contribute to the propositional meaning of a sentence but construct a context in pragmatics for the text. For instance, using a smile emoticon when violating social conversational norms signals joking or teasing [44], which can significantly enhance interpersonal rapport between friends [33]. Following [12], who reveal that humans have six basic emotions, we provide six emoticon types: Happy, Sad, Fear, Anger, Surprise, and Disgust. Each type has two to three variants.

## 4 SYSTEM ARCHITECTURE

Below we detail the architecture of SOGO to operationalize our computational model, shown in Figure 1.

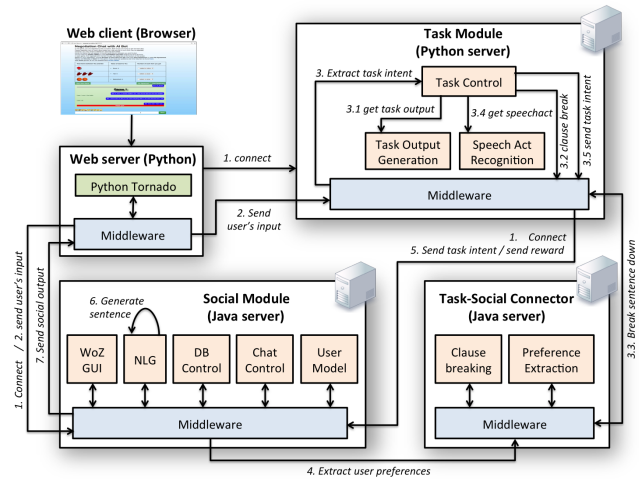


Figure 1: SOGO's Overall Architecture

**Overview** SOGO modules are deployed across four server nodes (two Java servers and two Python servers), and the system can be accessed by users through a web browser (Chrome and Safari are supported). The Web Server processes web client requests and delegates who should process each request; the Task Module is in charge of generating the next task intent according to the negotiation state and last user's input; the Task-Social Connector module runs different NLP Stanford APIs and bootstraps the transition from task to social phase; and finally the Social Module generates a social output according to the current conversational state, the given task output, and the chosen conversational strategy. All servers use a middleware layer that guarantees interoperability across multiple languages and operating systems, hides the underlying complexity of the environment, and masks the heterogeneity of networking technologies to facilitate programming of high-level features. The middleware layer provides multiple capabilities such as communication, message passing, concurrency, logging, service discovery, session management, and component pluggability. The Middleware layer uses ZMQ, a high-performance asynchronous socket-based messaging library for use in distributed and concurrent applications with minimal latency footprint.

**Pipeline** The Task-Social Connector server begins as a daemon service (a long running process) which listens for incoming requests. Then, the Social Module is launched and waits for other modules to connect. After a user initializes the interaction through the web browser and starts the negotiation game, the servers connect (the Web Server, the Task Module, the Social Module and the Social-Task Connector – step 1 in Figure 1). During the user's turn, his/her input (e.g., User: "I want two books and the ball") is sent to both the Task and Social Module (step 2), and the two processes run in parallel (steps 3 and 4). The Task Module generates a task output by using Facebook end-to-end model (FB) [25] (e.g., Agent: "I need one book and one hat, you can take the rest", step 3.1). The Task-Social Connector server splits the user's input and FB's output into two clauses (e.g., "I want two books" and "I want the ball", steps 3.2 and 3.3, respectively) and extracts the corresponding speech acts (e.g., <request, 2, books> and <request, 1, ball>, step 3.4). The composite task intent (i.e., the task output plus speech acts) is then sent to the

Dimension	Subjective Questions	t-test t(29)	$\mu_{SOGO}$	$\mu_{FB}$	Cohen's d
Coordination	1. I think that my agent and I were in sync with each other.	3.13**	5.10	3.77	0.81
	2. I felt uncomfortable and could not say everything that I wanted to say.	-3.05**	1.63	2.73	-0.79
	3. The interaction was frustrating.	-1.82	2.90	3.70	-0.47
Attentiveness	4. I felt that my agent was paying attention to what I was saying.	3.44**	5.23	3.53	0.89
	5. I was not really interested in what my agent was saying.	-0.24	1.63	1.70	-0.06
	6. My agent was respectful to me and considered to my concerns.	3.11**	5.43	4.10	0.80
Positivity	7. My agent was friendly to me.	3.90***	5.97	4.43	1.00
	8. I liked and felt warm toward my partner.	2.96**	5.10	3.80	0.76
	9. My agent cared about me.	3.36**	4.73	3.30	0.87
Face	10. Did you lose face (i.e., damage your sense of pride) in the negotiation?	0.00	1.76	1.76	0.0
Feeling about the negotiation	11. My agent was very uncooperative.	-1.36**	2.23	2.70	-0.35
	12. How satisfied are you with the balance between your own outcome and your agent's outcome(s)?	1.21**	5.07	4.57	0.31
	13. Did the negotiation build a good foundation for a future relationship with your agent?	3.83***	5.37	3.83	0.99
Perceived Rapport	14. I felt rapport between the agent and myself.	4.04**	5.10	3.50	1.04
Information Disclosure	15. I was willing to share information with my agent.	-0.77**	4.40	4.73	-0.20

**Table 5: Complete t-test statistical analysis of subjective questionnaire of rapport assessment by comparing SOGO system and Facebook end-to-end system. Effect size assessed via Cohen's d. Significance: \*\*\*:p <0.001, \*\*:p <0.01, \*:p <0.05**

Social Module (steps 3.5, and later, step 5). Meanwhile, the Social Module is extracting user preferences in order to update the user model (step 4) and displays user and system interaction in real-time on a WoZ (Wizard-Of-Oz) dashboard GUI. Once it receives the task intent from the Task Module (step 5), the Social Module executes a template-based Natural Language Generator (NLG) component which loads a set of pre-defined conversational strategies (using the DB Control). These strategies are then combined with a self-reflection mechanism based on a user's input parser and filtered using a rule-based system (step 6). Given the current dialogue state, a set of plausible social sentences are shown to the Wozer, who chooses one (e.g., Agent: "This book looks exactly like one my grandpa gave me, would you mind giving me that book and the hat that looks really nice on me? you can have the rest..."). The Social Module sends this output to the Web Server, which in turn displays it on user's browser (step 7).

## 5 EXPERIMENT SETUP AND DATA COLLECTION

We deployed the same multi-issue task as described in [25]. SOGO is semi-automatic: The functions described above are carried out automatically, but the Wozer decide which strategy to use when there are multiple available. In our experiment, we use Facebook end-to-end model as a baseline and compare it with our developed SOGO system. Thus, we recruited 60 English speakers on Amazon Mechanical Turk who were equally and randomly assigned to one of the conditions. To obtain high quality data, those workers were based in US or UK and had at least 95% approval rating with more than 5,000 previous HITs. Each participant played six games with the agent and completed a subjective questionnaire to reveal their feelings toward the game and interlocutor.

## 6 EVALUATION

In this study, we combine subjective and objective measures. Two-tailed independent sample t-tests were conducted on both the questionnaires and objective metrics to explore the difference of mean value of users' rating towards two systems. For all significant results ( $p < 0.05$ ), we also calculated effect size via Cohen's d to test for generalizability of results.

### 6.1 Subjective metrics

Based on items used in prior studies [8, 9, 16], we developed a 15-item self-report questionnaire that characterizes the interaction into dimensions of rapport, such as coordination, attentiveness, positivity, and so on; question 14 asked users to directly rate the overall feeling of rapport during the interaction. Responses were rated by each participant on a scale of 0 (Strongly Disagree) to 7 (Strongly Agree). Factor analysis proved only one factor for the 15 questionnaire items, which have high internal consistency with Cronbach's  $\alpha = 0.94$ . Table 5 shows the complete list of questions and results. We describe our findings of differences between two grounds on each dimension of rapport as follows:

**Coordination:** We observed that users felt more in synchronicity with the SOGO system ( $d=0.81$ ), as they could say almost everything that they wanted to say during the interaction ( $d=-0.79$ ). Next, we noticed users felt a little bit frustrated in both setups but showed no significant differences toward the two systems.

**Attentiveness:** Users reported that the SOGO system paid more attention to them ( $d=0.89$ ). Meanwhile, they also realized that the SOGO system was more respectful and considered their concerns ( $d=0.80$ ). Users stated that they were interested in listening to both systems. Thus, there is no significant difference among the two groups of participants, but the mean values of this question are all low ( $\mu_{SOGO}=1.63$ ,  $\mu_{FB}=1.70$ ).



**Positivity:** Users liked the SOGO system more and felt warm toward to their partner ( $d=0.76$ ). They experienced more of a sense of friendliness ( $d=1.00$ ) and caring from the SOGO system, as well ( $d=0.87$ ).

**Face:** Both groups reported a low degree of damaging their sense of pride but no significant difference was found across groups.

**Feeling about the negotiation:** We ameliorated the unpromising and uncooperative impression in users from the FB system to the SOGO system ( $d=-0.35$ ), though the rating of the SOGO system is still unsatisfied ( $\mu_{SOGO}=2.23$ ). Users felt more satisfied about the instrumental outcome in the SOGO system with low effective size ( $d=0.31$ ). Finally, we found that users regarded the whole negotiation process as a good foundation for a future relationship with the SOGO system. The attitude to these questions differ significantly in the two groups with a large effective size ( $d=0.99$ ).

**Perceived Rapport:** User perceive significantly higher rapport with SOGO system ( $\mu_{SOGO}=5.10$ ) comparing to FB system ( $\mu_{FB}=3.50$ ).

**Information Disclosure:** Users prefer to share more of their personal information with the FB system rather than the SOGO system but the results suffer from low effective size ( $d=-0.2$ ).

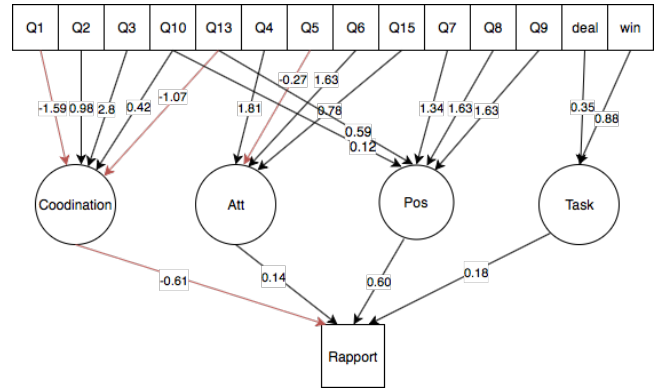
## 6.2 Objective metrics

Objective Metrics	t-test t(29)	$\mu_{SOGO}$	$\mu_{FB}$	Cohen's d
Win Times	2.59*	2.70	1.80	0.67
Deal Rate	7.74***	0.90	0.45	1.99
Average Dialogue Length	-1.50	6.80	7.59	-0.39
Average User Utterance Length	2.59*	7.17	5.48	$d=0.67$
Pareto Optimal	2.05**	96.67	80.00	0.53

**Table 6: Complete t-test statistical analysis of negotiation performance of SOGO system versus FB system. Effect size assessed via Cohen's d. Significance: \*\*\*:p <0.001, \*\*:p <0.01, \*:p <0.05**

In objective metrics, we firstly measured agent performance through three dimensions inherited from [25]: (1) Number of wins by the system (**Win Times**). Obviously, the SOGO system wins more often than the FB system, with a moderate effective size ( $d=0.67$ ). (2) Percentage of games that end up with an agreed-upon negotiation decision (**Deal Rate**). We know this is the major problem with the FB system in [25] since users even prefer not to agree rather than capitulate to an uncooperative system. The SOGO system significantly improves this agreement rate from 0.45 to 0.90 with high effective size ( $d=1.99$ ). (3) Percentage of Pareto optimal solutions for agreed deals (**Pareto Optimal**). The SOGO system performs better in this dimension because users prefer to adapt or even sacrifice themselves to agree with the system as a means of building rapport.

Next, since **Average Dialogue Length** and **Average User Utterance Length** are strong objective indicators of user engagement or interest [41], we include them here as well. We find that the average dialogue length for both systems are similar, however, users reply with more words in each utterance when they negotiate with the SOGO system ( $d=0.67$ ). Table 6 provides complete results.



**Figure 2: Structural Equation Model of Rapport**

Name	Value
RMSEA	0.094
SRMR	0.063
CFI	0.937
TLI	0.915

**Table 7: Model fit metrics. RMSEA = root mean square error of approximation; SRMR = Standardized Root Mean Square Residual; CFI = comparative fit index; TLI = tucker-lewis index**

## 6.3 Factor analysis of Rapport

Since [37, 42] have provided a clear computational model of rapport, we propose a four-factor model to explain the construct of negotiation rapport. Our goals are to validate this model and specify the variable loadings of each question to its corresponding factor of rapport. With respect to these goals, we conducted a confirmatory factor analysis, which assesses how well the proposed model captured the covariance between all variables in the model. In our case, subjective questions are observed variables, represented by square boxes; the subcomponents of rapport are the latent factors, drawn by circles. Overall, our proposed four-factor model is a relatively good fit based on the metrics listed in the table 7. In the top-level of the model, coordination and positivity correlate strongly to rapport in the context of negotiation. This result tells us that users value feelings of synchronicity and friendliness more than others. The negative sign of coordination loadings is due to the reversed-coded questions. Also, our findings confirm [8], namely that the outcome of the task (e.g., win or lose) does not greatly affect the building of a social bond. Unexpectedly, attentiveness seems like the least important factor of rapport, which is contradicts rapport theory [39]. One explanation of this phenomenon is that attentiveness has a large covariance with both coordination and positivity. Thus, its explanatory power toward the variance of rapport is reduced. In the bottom level, most questions have high variable loadings and figure 2 displays the complete results.

## 7 CONCLUSION AND FUTURE WORK

In this paper, we introduce and evaluate our social intelligent negotiation dialogue system, which can negotiate with people while building a social bond. These achievements come from our proposed two-phase computational model that blends social moves with task moves in an utterance. We leverage the off-the-shelf end-to-end dialogue model to decide the next task move and a theory-driven

template-based social language generator to introduce social skill to the system. The experiment demonstrates that our SOGO system behaves in sync with its user. People feel more comfortable and engaged during the interaction compared to the baseline FB system. Even though they stay in a semi-cooperative environment, people think our SOGO system is friendly and cares about them.

The design of our computational model and system architecture is supported by empirical work in social psychology, which helps identify areas in which the system can develop human-like qualities. Our work operationalizes these theories to practical human-agent interaction. Especially, in the social phase, we realize abstract theoretical findings into conversational strategies and speech act strategies. Our findings improve understanding of how to instantiate rapport in human-agent negotiation. Broadly speaking, we validate the discovery that strategies for building buddy relationships in human-human communication can be transferred to human-agent interaction.

In our future work, we plan to develop a fully-automatic SOGO system. Based on our collected human-agent dialogues, we will build a social reasoning module that behaves like [35] to rank and select eligible rapport strategies given a present conversational state. Meanwhile, we look forward to designing a rapport estimator that can capture the dynamics of rapport and provides an instant social reward in real-time. We also intend to create a dialogue model to optimize the sum of social outcomes as well as the negotiation result. Furthermore, we would like to extend our SOGO system to other domains such as counseling, job interview, elderly companion and etc. We expect the future challenges to be substantial, but rewarding, as we believe improving the human-like social capability of a system will greatly impact human-agent collaboration.

## REFERENCES

- [1] Leslie M Beebe and Tomoko Takahashi. 1989. Sociolinguistic variation in face-threatening speech acts. In *The dynamic interlanguage*. Springer, 199–218.
- [2] Timothy Bickmore and Justine Cassell. 2001. Relational agents: a model and implementation of building user trust. In *Proceedings of the SIGCHI conference on Human factors in computing systems*. ACM, 396–403.
- [3] Timothy Bickmore and Daniel Schulman. 2012. Empirical validation of an accommodation theory-based model of user-agent relationship. In *International Conference on Intelligent Virtual Agents*. Springer, 390–403.
- [4] Shoshana Blum-Kulka and Elite Olshtain. 1984. Requests and apologies: A cross-cultural study of speech act realization patterns (CCSARP). (1984), 196–213.
- [5] Piotr Bojanowski, Edouard Grave, Armand Joulin, and Tomas Mikolov. 2016. Enriching word vectors with subword information. *arXiv preprint arXiv:1607.04606* (2016).
- [6] Ilan Bronstein, Noa Nelson, Zohar Livnat, and Rachel Ben-Ari. 2012. Rapport in negotiation: The contribution of the verbal channel. *Journal of Conflict Resolution* 56, 6 (2012), 1089–1115.
- [7] William Coon, Charles Rich, and Candace L Sidner. 2013. Activity planning for long-term relationships. In *Intelligent Virtual Agents: 13th International Conference, IVA 2013, Edinburgh, UK, August 29-31, 2013, Proceedings*, Vol. 8108. Springer, 425.
- [8] Jared R Curhan, Hillary Anger Elfenbein, and Heng Xu. 2006. What do people value when they negotiate? Mapping the domain of subjective value in negotiation. *Journal of personality and social psychology* 91, 3 (2006), 493.
- [9] David DeVault, Johnathan Mell, and Jonathan Gratch. 2015. Toward natural turn-taking in a virtual human negotiation agent. In *AAAI Spring Symposium on Turn-taking and Coordination in Human-Machine Interaction*. AAAI Press, Stanford, CA.
- [10] Eli Dresner and Susan C Herring. 2010. Functions of the nonverbal in CMC: Emoticons and illocutionary force. *Communication theory* 20, 3 (2010), 249–268.
- [11] Miriam Eisenstein and Jean W Bodman. 1986. ăÏÿ very appreciateăÏÿ: expressions of gratitude by native and non-native speakers of American English. *Applied linguistics* 7, 2 (1986), 167–185.
- [12] Paul Ekman and Wallace V Friesen. 1986. A new pan-cultural facial expression of emotion. *Motivation and emotion* 10, 2 (1986), 159–168.
- [13] Angela Fabregues and Carles Sierra. 2011. Dipgame: a challenging negotiation testbed. *Engineering Applications of Artificial Intelligence* 24, 7 (2011), 1137–1146.
- [14] Peyman Faratin, Carles Sierra, and Nick R Jennings. 1998. Negotiation decision functions for autonomous agents. *Robotics and Autonomous Systems* 24, 3-4 (1998), 159–182.
- [15] CJ Hutto Eric Gilbert. 2014. Vader: A parsimonious rule-based model for sentiment analysis of social media text. In *Eighth International Conference on Weblogs and Social Media (ICWSM-14)*.
- [16] Jonathan Gratch, David DeVault, Gale M Lucas, and Stacy Marsella. 2015. Negotiation as a challenge problem for virtual humans. In *International Conference on Intelligent Virtual Agents*. Springer, 201–215.
- [17] Jonathan Gratch, Anna Okhmatovskaia, Francois Lamothe, Stacy Marsella, Mathieu Morales, Rick J van der Werf, and Louis-Philippe Morency. 2006. Virtual rapport. (2006), 14–27.
- [18] Robert H Guttman and Pattie Maes. 1998. Agent-mediated integrative negotiation for retail electronic commerce. In *International Workshop on Agent-Mediated Electronic Trading*. Springer, 70–90.
- [19] Koen Hindriks, Catholijn M Jonker, Sarit Kraus, Raz Lin, and Dmytro Tykhonov. 2009. Genius: negotiation environment for heterogeneous agents. In *Proceedings of the 8th International Conference on Autonomous Agents and Multiagent Systems-Volume 2*. International Foundation for Autonomous Agents and Multiagent Systems, 1397–1398.
- [20] Lixing Huang, Louis-Philippe Morency, and Jonathan Gratch. 2011. Virtual Rapport 2.0. (2011), 68–79.
- [21] Emmanuel Johnson, Jonathan Gratch, and David DeVault. 2017. Towards An Autonomous Agent that Provides Automated Feedback on Students' Negotiation Skills. In *Proceedings of the 16th Conference on Autonomous Agents and Multi-Agent Systems*. International Foundation for Autonomous Agents and Multiagent Systems, 410–418.
- [22] Dejun Tony Kong, Kurt T. Dirks, and Donald Ferrin. 2014. Interpersonal Trust within Negotiations: Meta-Analytic Evidence, Critical Contingencies, and Directions for Future Research. *57* (10 2014), 1235–1255.
- [23] Vasily Konovalov, Ron Artstein, Oren Melamud, and Ido Dagan. 2016. The Negotiation Corpus of Human-agent Negotiation Dialogues.. In *LREC*.
- [24] Adam Lerer and Alexander Peysakhovich. 2017. Maintaining cooperation in complex social dilemmas using deep reinforcement learning. *arXiv preprint arXiv:1707.01068* (2017).
- [25] Mike Lewis, Denis Yarats, Yann N Dauphin, Devi Parikh, and Dhruv Batra. 2017. Deal or No Deal? End-to-End Learning for Negotiation Dialogues. *ArXiv e-prints* (2017).
- [26] Gale Lucas, Giota Stratou, Shari Liebling, and Jonathan Gratch. 2016. Trust me: multimodal signals of trustworthiness. In *Proceedings of the 18th ACM International Conference on Multimodal Interaction*. ACM, 5–12.
- [27] Christopher D. Manning, Mihai Surdeanu, John Bauer, Jenny Finkel, Steven J. Bethard, and David McClosky. 2014. The Stanford CoreNLP Natural Language Processing Toolkit. In *Association for Computational Linguistics (ACL) System Demonstrations*. 55–60.
- [28] Yoichi Matsuyama, Arjun Bhardwaj, Ran Zhao, Oscar Romeo, Sushma Akoju, and Justine Cassell. 2016. Socially-aware animated intelligent personal assistant agent. In *Proceedings of the 17th Annual Meeting of the Special Interest Group on Discourse and Dialogue*. 224–227.
- [29] Nikita Mattar and Ipke Wachsmuth. 2012. Small talk is more than chit-chat. In *Annual Conference on Artificial Intelligence*. Springer, 119–130.
- [30] Johnathan Mell and Jonathan Gratch. 2016. IAGO: Interactive Arbitration Guide Online. In *Proceedings of the 2016 International Conference on Autonomous Agents & Multiagent Systems*. International Foundation for Autonomous Agents and Multiagent Systems, Singapore, 1510–1512.
- [31] Johnathan Mell and Jonathan Gratch. 2017. Grumpy & Pinocchio: Answering Human-Agent Negotiation Questions through Realistic Agent Design. In *Proceedings of the 16th Conference on AAMAS*. International Foundation for Autonomous Agents and Multiagent Systems, 401–409.
- [32] Janice Nadler. 2003. Rapport in conflict resolution. *Marq. L. Rev.* 87 (2003), 875.
- [33] Amy Ogan, Samantha Finkelstein, Erin Walker, Ryan Carlson, and Justine Cassell. 2012. Rudeness and rapport: Insults and learning gains in peer tutoring. In *International Conference on Intelligent Tutoring Systems*. Springer, 11–21.
- [34] Landra Rezabek and John Cochenour. 1998. Visual cues in computer-mediated communication: Supplementing text with emoticons. *Journal of Visual Literacy* 18, 2 (1998), 201–215.
- [35] Oscar J. Romero, Ran Zhao, and Justine Cassell. 2017. Cognitive-Inspired Conversational-Strategy Reasoner for Socially-Aware Agents. In *Proc. of the 26th Int. Joint Conf. on Artificial Intelligence, IJCAI-17*. 3807–3813.
- [36] Helen Spencer-Oatey. 2008. *Culturally speaking: Culture, communication and politeness theory*. Continuum International Publishing Group.
- [37] Helen Spencer-Oatey. 2008. Face, (im) politeness and rapport. (2008).
- [38] Fan-Yun Sun, Yen-Yu Chang, Yueh-Hua Wu, and Shou-De Lin. 2018. Designing Non-greedy Reinforcement Learning Agents with Diminishing Reward Shaping. (2018).



- [39] Linda Tickle-Degnen and Robert Rosenthal. 1990. The nature of rapport and its nonverbal correlates. *Psychological inquiry* 1, 4 (1990), 285–293.
- [40] Joseph B Walther and Kyle P D'Addario. 2001. The impacts of emoticons on message interpretation in computer-mediated communication. *Social science computer review* 19, 3 (2001), 324–347.
- [41] Zhou Yu, Alan W Black, and Alexander I Rudnicky. 2017. Learning conversational systems that interleave task and non-task content. *arXiv preprint arXiv:1703.00099* (2017).
- [42] Ran Zhao, Alexandros Papangelis, and Justine Cassell. 2014. Towards a dyadic computational model of rapport management for human-virtual agent interaction. In *Intelligent Virtual Agents*. Springer, 514–527.
- [43] Ran Zhao, Alexandros Papangelis, and Justine Cassell. 2014. Towards a dyadic computational model of rapport management for human-virtual agent interaction. In *International Conference on Intelligent Virtual Agents*. Springer, 514–527.
- [44] Ran Zhao, Tanmay Sinha, Alan W Black, and Justine Cassell. 2016. Socially-aware virtual agents: Automatically assessing dyadic rapport from temporal patterns of behavior. In *International conference on intelligent virtual agents*. Springer.