

How Rude are You?: Evaluating Politeness and Affect in Interaction

Swati Gupta¹, Marilyn A. Walker¹, Daniela M. Romano¹

¹Department of Computer Science, University of Sheffield, Regent Court, 211 Portobello street,
Sheffield, UK, S1 4DP

{s.gupta, m.walker, d.romano}@dcs.shef.ac.uk

Abstract. Recent research on conversational agents emphasises the need to build affective conversational systems with social intelligence. Politeness is an integral part of socially appropriate and affective conversational behaviour, e.g. consider the difference in the pragmatic effect of realizing the same communicative goal with either “Get me a glass of water mate!” or “I wonder if I could possibly have some water please?” This paper presents POLLY (Politeness for Language Learning), a system which combines a spoken language generator with an artificial intelligence planner to model Brown and Levinson’s theory of politeness in collaborative task-oriented dialogue, with the ultimate goal of providing a fun and stimulating environment for learning English as a second language. An evaluation of politeness perceptions of POLLY’s output shows that: (1) perceptions are generally consistent with Brown and Levinson’s predictions for choice of form and for discourse situation, i.e. utterances to strangers need to be much more polite than those to friends; (2) our indirect strategies which should be the politest forms, are seen as the rudest; and (3) English and Indian native speakers of English have different perceptions of politeness.

1 Introduction

Recent research suggests that computers are perceived as social actors, who must exhibit social intelligence and awareness, rather than merely as computational machines that perform tasks assigned to them by the user [1,7,11,17,18,20,26]. This social role awareness involves the ability to behave in a socially correct manner, where an integral part of this behaviour is conversation, the ability to communicate appropriately, according to the situation and the feelings of the interlocutors. For example, consider the difference in the pragmatic effect of realizing the same communicative goal with either “Get me a glass of water mate!” or “I wonder if I could possibly have some water please?”.

According to theories in sociolinguistics, choices of these different forms are driven by sociological norms among human speakers [6,9] *inter alia*, and work on computational models for conversational agents has recently begun to build on these sociolinguistic theories. Walker et al. [26] were the first to utilize and implement Brown & Levinson’s [6] theory of politeness, henceforth B&L, in conversational agents, in order to provide interesting variations of character and personality in an interactive narrative application. Other work has explored building affective conversational systems that are considerate of the emotions of the user and which exhibit appropriate emotions [2,17], and work has shown that the expression and recognition of personality is strongly linked to positive and negative affect [14,15]. But politeness is an integral part of affective conversational behaviour, as impoliteness exhibits negative feelings towards the hearer, and may hurt the hearer’s feelings or make the hearer angry.

This paper presents POLLY (Politeness for Language Learning), a system which combines a spoken language generator with an AI Planner to model B&L’s theory of politeness in task-oriented dialogue. The value of politeness strategies based on B&L has been demonstrated in several conversational applications, e.g. tutorial dialogue [11,12,18], animated presentation teams [1,21], and real estate sales [7]. Recent research also shows that human tutors employ politeness strategies while interacting with students and that pedagogical agents that use polite language provide affective scaffolding to the instructors and contribute to the learners’ motivational state to help them learn difficult concepts [11,19,27]. André et al. [3] use politeness strategies to mitigate face threats resulting from dialogue acts and investigate how the user’s affective response to the system can be improved. Roman et al. [22] found that politeness plays a role in dialogue summarization since human summarizations tend to report politeness as a result of their point of view (which interlocutor they are asked to empathize with or if they are supposed to act as an observer). This was most evident for reporting impolite behaviour. This bias could be for saving the face of the

interlocutor they empathize with, which is directly related to the notion of self-esteem. Morand & Ocker [16] suggest how politeness contributes to the study of role relations in computer-mediated communication. They point out that in task-oriented speech acts, emotion work appears in the form of politeness and the degree, type and tactics of politeness provide important cues regarding actors' relational orientations towards each other. Reeves and Nass [20] also observed that users are polite to the computers because they are considerate about the computer's feelings, and are not likely to speak in a manner that might hurt its feelings. Thus, previous work suggests that computers should reciprocate with humans by being polite, as social behaviours are not accomplished in isolation from the responses to them, and sociological norms dictate that humans expect reciprocity.

Here, we explore the effect of politeness in natural conversation by evaluating the use of different politeness strategies in task-oriented dialogues in a collaborative task domain of cooking, where subjects are asked to collaborate with another person to make a recipe [10,24,25]. We show that: (1) politeness perceptions of POLLY's output are generally consistent with B&L's predictions for choice of form and for discourse situation, i.e. utterances to strangers need to be more polite than those to friends; (2) our indirect strategies which should be the politest forms, are seen as the rudest; and (3) English and Indian speakers of English have different perceptions of politeness. Section 1 describes POLLY's architecture and functionality. Section 2 describes an evaluation of users' perceptions of automatically generated task-oriented polite language and Section 3 presents the experimental results. Section 4 sums up and compares our results with previous work.

2 POLLY's Architecture and Theoretical Basis

POLLY consists of two parts: an AI Planner based on GraphPlan [5] and a Spoken Language Generator (SLG), as illustrated in Figure 1. GraphPlan is a classic STRIPS-style planner which, given a goal, e.g. cook pasta, produces a plan of the steps involved in doing so. POLLY then allocates the plan steps to two agents as a shared collaborative plan to achieve the cooking task, with goals to communicate about the plan via speech acts (SAs) needed to accomplish the plan collaboratively, such as Requests, Offers, Informs, Acceptances and Rejections [10,24,25].

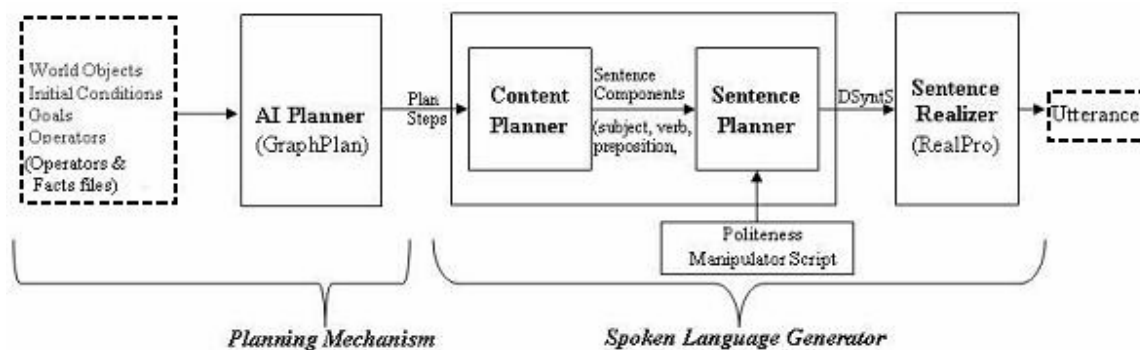


Fig. 1. POLLY's Architecture

The SLG then generates variations of the dialogue based on B&L's theory of politeness that realizes this collaborative plan, as in [1,26]. This is explained in more detail below and an example dialogue is shown in Table 2. When this dialogue is embedded in our virtual reality environment [23], the human English language learner will be able to play the part of one of the agents in order to practice politeness in a real-time immersive environment.

2.1 Brown and Levinson's Theory

B&L's theory states that speakers in conversation are rational actors who attempt to realize their speech acts (SAs) to avoid threats to one another's **face**, which consists of two components. **Positive face** is the desire that at least some of the speaker's and hearer's goals and desires are shared by other speakers.

Negative face is the want of a person that his action be unimpeded by others. Utterances that threaten the conversants' face are called Face Threatening Acts (FTAs). B&L predict a universal of language usage that

Table 1. The individual B&L strategies used for Request and Inform speech acts

B&L	Request Speech Act		Inform Speech Act	
	Strategy Forms	Strategy Names	Strategy Forms	Strategy Names
Direct	Do X.	RD1Imperative	X	ID1DirectAssert
	Do X please.	RD2ImperativePlz	-	-
	You must do X.	RD3ImperativeInsist	-	-
	You could do X.	RD4AsModAbility	-	-
Approval	Could you please do X mate?	RAp1QModAbility	Do you know that X?	IAp1QKnowledge
	If you don't mind you can do X.	RAp2AsModAbility	Do you know that X mate?	IAp2QueryKNowledgeAddress
	Would it be possible for you to do X?	RAp3AsPossible	-	-
	I'm sure you won't mind doing X.	RAp4AsOptimism	-	-
Autonomy	Could you possibly do X for me?	RAu1QModAbility	It seems that X.	IAu2AsAppear
	I know I'm asking you for a big favour but could you please do X?	RAu2ApologizeQModAbility	I am wondering if you know that X.	IAu1AsConfuse
	I'm wondering whether it would be possible for you to do X.	RAu3AsConfusePossibility	-	-
	Would you not like to do X?	RAu1QOptimism	-	-
Indirect	X is not done yet.	RI1AsNegation	-	-
	X should have been done.	RI2AsModRight	-	-
	Someone should have done X.	RI3AsModRightAbsSub	-	-
	Someone has not done X yet.	RI4AsNegationAbsSub	-	-
	<i>Where X is a task request. For example 'You could chop the onions,' or 'Would it be possible for you to clean the spill on the floor?'</i>	<i>These strategies are applied to the various tasks requests X.</i>	<i>Where X is an inform event, like 'Do you know that the milk is spoilt mate?' or 'I'm wondering if you know that you have burnt the pasta.'</i>	<i>These strategies are applied to the various inform events X.</i>

the choice of linguistic form can be determined by the predicted Threat Θ as a sum of 3 variables:

1. P: power that the hearer has over the speaker;
2. D: social distance between speaker & hearer;
3. R: a ranking of imposition of the speech act.

Linguistic strategy choice is made according to the value of the Threat Θ . We follow Walker et al.'s [26] four part classification of strategy choice. The **Direct strategy** is used when Θ is low and executes the SA in the most direct, clear and unambiguous way. It is usually carried out either in urgent situations like "Please Help!", or where the face threat is small as in informing the hearer "I have chopped the vegetables" or if the speaker has power over the hearer, "Did you finish your homework today?". The **Approval strategy** (Positive Politeness) is used for the next level of threat Θ - this strategy is oriented towards the need for the hearer to maintain a positive self-image. Positive politeness is primarily based on how the speaker approaches the hearer, by treating him as a friend, a person whose wants and personality traits are liked, and by using friendly markers "Could you please chop the vegetables mate?" or stating optimism "I'm sure you won't mind washing the dishes!" The **Autonomy Strategy** (Negative Politeness) is used for greater face threats, when the speaker may be imposing on the hearer, intruding on their space or violating their freedom of action. These face threats can be mitigated by apologizing, "I know I'm asking you for a

big favour but could you please wash the dishes?” or by minimizing imposition, “I just want to ask you if you could close the door.” The **Indirect Strategy** (Off Record) is the politest strategy and is therefore used when Θ is greatest. It depends on speaking in an indirect way, with more than one attributable intention so that the speaker removes himself from any imposition. For example by using metaphor and irony, rhetorical questions, understatement, or hints such as “Its cold in here,” which implies a request to close the door, or being vague like “Someone should have cleaned the table.” Table 1 lists the B&L strategies used in the evaluation experiment in Section 3.

2.2 Planning Mechanism

Planning is the process of generating a sequence of actions that can achieve a pre-specified goal. In particular, our planner generates the sequence of actions that are to be performed for cooking pasta. The information needed to create dialogic utterances for the agents in a dialogue are extracted from the plan. The planner output can be in any form depending upon the planner used, but a mapping between the components of the plan and the lexicalized entries for the syntactic structure of the utterance realizing that plan component is required. This mapping is typically referred to as a generation dictionary. Given a generation dictionary (a mapping) the dialogue generation component of the system can be used with any planner.

The AI planner GraphPlan [5] has been used for POLLY. GraphPlan applies the Planning Graph Analysis approach to a STRIPS-like planning domain where the operators have preconditions, add effects, and delete effects which are actually conjuncts of propositions that have parameters that can be instantiated to objects in the world. Thus a planning problem consists of a STRIPS-like domain, a set of objects, a set of propositions or initial conditions and a set of problem goals required to be true at the end of a plan. Graphplan takes the objects, initial conditions, goals and operator definitions as input (see Figures 1, 2 and 3) and creates a plan for cooking pasta. The facts file shown in Figure 3 defines the objects of the world, the initial conditions that need to be true and the final goals that have to be achieved. The operators file shown in Figure 2 contains the operator definitions with their parameters, preconditions and effects where the parameters are initialised with the objects of the world as defined in the facts file.

```
(operator chop
  (params (<v> V) (<k> K))
  (preconds (available <v>)
            (available <k>))
  (effects (chopped <v>)
           (del available <v>)
           (del not-chopped <v>)
           (not-placed pan1 burner1)))

(operator cook
  (params (<v> V) (<p2> P2))
  (preconds (chopped <v>)
            (ingredients-added other-ingredients <v>))
  (effects (cooked <v>)
           (del chopped <v>)))
```

Fig. 2. Excerpt of the operators file

```
(vegetables V)
(knife K)
(pasta P1)
(water W)
....
(preconds
  (available vegetables)
  (available knife)
  (available pasta)
  (available water)
  .....)
(effects
  (ready pasta))
```

Fig. 3. Excerpt of the facts file

An excerpt from the output plan is:

- Step 1: place_pan_burner
- Step 2: turn-on_burner
- Step 3: boil_pasta_pan
- Step 4: chop_vegetables_knife
- Step 5: place_pan_burner
- Step 6: add_oil_pan

In our plan operators, lexicalization is directly encoded in the plan, so the generation dictionary is not needed, i.e. lexical entries such as place, pan, burner, etc are directly picked up from the plan steps by the language generator.

2.3 SLG (Spoken Language Generation)

The SLG is based on a standard architecture [8] with three components: Content planning, utterance planning and surface realization. See Figure 1. The politeness strategies are implemented through a combination of content selection and utterance planning. The linguistic realizer RealPro is used for realization of the resulting utterance plan [13], and the content planning and utterance planning components produce outputs that can be transformed into RealPro input, which we discuss first.

The **Surface Realizer** RealPro takes a dependency structure called a Deep-Syntactic Structure (DSyntS) as input and realizes it as a string. DSyntS are unordered *trees* with labelled nodes and arcs where the nodes are lexicalized. Only meaning bearing lexemes are represented and not function words. An example of a DSyntS for the utterance “I have chopped the vegetables.” is given below. The attributes to all the nodes are explicitly specified, such as tense, or article. The two nodes are specified with relations I and II, where I is the subject and II is the object.

```
"chop" [ lexeme: "chop" class: "verb" taxis: "perf" tense: "pres" ]
(
  I "<PRONOUN>" [ lexeme: "<PRONOUN>" number: "sg" person: "1st" rel: "I" ]
  II "vegetable" [ lexeme: "vegetable" article: "def" class: "common_noun" number: "pl" rel: "II" ]
)
```

<pre>"cook" [lexeme: "cook" class: "verb" tense: "pres" mood: "imp"] (II "vegetable" [lexeme: "vegetable" article: "def" class: "common_noun" number: "pl" rel: "II"]) This would be realized simply as "Cook the vegetables." It is transformed to create utterances which vary in politeness according to B&L.</pre> <p style="text-align: center;">Base DSyntS</p>	<pre>"wonder" [lexeme: "wonder" class: "verb" tense: "pres"] (I "<PRONOUN>" [lexeme: "<PRONOUN>" number: "sg" person: "1st" rel: "I"] II "be" [lexeme: "be" class: "verb" mood: "cond" tense: "pres" rel: "II"] (II "possible" [lexeme: "possible" rel: "II"] I "it" [lexeme: "it" rel: "I"] III "cook" [lexeme: "cook" class: "verb" mood: "inf-to" question: "+" rel: "III"] (II "vegetable" [lexeme: "vegetable" article: "def" class: "common_noun" number: "pl" rel: "II"] I "<PRONOUN>" [lexeme: "<PRONOUN>" number: "sg" person: "2nd" rel: "I"]) ATTR "whether" [lexeme: "whether" rel: "ATTR"] ATTR "perhaps" [lexeme: "perhaps" class: "adverb" rel: "ATTR"])) Realized as "I wonder whether it would perhaps be possible for you to cook the vegetables."</pre> <p style="text-align: center;">Base DSyntS manipulated to create another polite DSyntS (the RAu7AsConfusePossibility strategy)</p>
<pre>"cook" [lexeme: "cook" class: "verb" mood: "cond" question: "+"] (I "<PRONOUN>" [lexeme: "<PRONOUN>" number: "sg" person: "2nd" rel: "I"] II "vegetable" [lexeme: "vegetable" article: "def" class: "common_noun" number: "pl" rel: "II"] ATTR "like_to" [lexeme: "like_to" class: "adverb" rel: "ATTR"] ATTR "mate" [lexeme: "buddy" rel: "ATTR"]) Realized as "Would you like to cook the vegetables mate?"</pre> <p style="text-align: center;">Base DSyntS manipulated to create a polite DSyntS (the RAu9QOptimism strategy)</p>	

Fig. 4. Transformation from base DSyntS to RAu9QOptimism and RAu7AsConfusePossibility strategies for the CookVeg task

The **Content Planner** interfaces to the AI Planner, selecting content from the preconditions, steps and effects of the plan. According to B&L, direct strategies are selected from the steps of the plan, while realizations of preconditions and negating the effects of actions are techniques for implementing indirect strategies. For instance, in case of the first direct request strategy RDIImperative (stands for Request SA, Imperative direct strategy) in Table 1, which is realised as ‘Do X’, task X is selected from the steps of the plan and since it is a request SA and imperative strategy, it is realized simply as ‘Do X’. Similarly, in case of the first indirect strategy RIIAsNegation (Request SA, Assert Negation Indirect strategy) which is

realized as ‘X is not done yet’, the content is selected by the negation of effects of the action of doing X. The content planner extracts the components of the utterances to be created, from the plan and assigns them their respective categories, for example, lexeme *get/add* under category verb, or *knife/oil* under direct object, and sends them as input to the Sentence Planner.

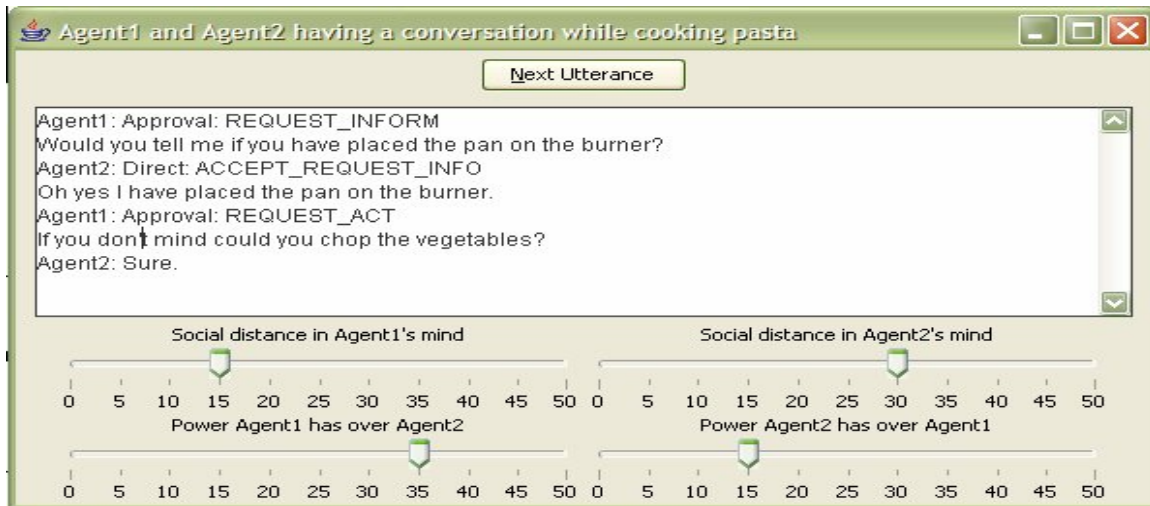


Fig. 5. A screenshot of the textual interface

The **Sentence Planner** then converts the utterance components to the lexemes of DSyntS nodes to create basic DsyntS for simple utterances [4], which are then transformed to create variations as per B&L’s politeness strategies. At the moment our interface is text based, but our plan is to embed it in Sheffield’s virtual reality environment. A screenshot of our textual interface is in Figure 5. The Sentence Planner creates SAs of two kinds: Initiating SAs such as request, inform, suggest, and offer and Response SAs such as inform SA and acceptance and rejection of requests or suggestions. To generate a conversation, first the initiating SAs are created followed by response SAs. The subject is implicitly assumed to be first person singular (I) in case of offer, inform, accept and reject, second person singular (you) in request_act and request_inform and first person plural (we) in case of suggest and accept_suggest. Each SA has multiple variants for realizing its politeness strategies, some of which are shown in Table 1.

For realizing these B&L strategies, a number of transformations on the basic DSyntS were implemented that were hypothesized to vary the politeness of a utterance. These politeness transformations are divided into four categories: **Address** form which means a friendly manner of addressing someone like ‘mate’. **Abstracting the subject** by saying ‘someone should have washed the dishes’ instead of addressing the hearer directly. **Softeners** like ‘if you don’t mind,’ ‘if you know,’ ‘please’ and ‘possibly’. **Additives** consisted of *Apologizing* like admitting impingement as in “I know I’m asking you for a big favour”, using *must* “You must take out the trash” and explicitly stating that you are asking a favour as in “Could you chop the onions for me?” For example if we want variations for a Request_act SA in which one agent requests the other to cook vegetables, the Content Planner sends the verb (cook) and the direct object (vegetable) to the Sentence Planner which then creates a base DsyntS. Figure 4 shows the RAu9QOptimism transformation for the CookVeg task (which stands for Request act speech act, Query optimism autonomy strategy for the task cook vegetables). In addition, in the second row of Table 1, the Sentence Planner transforms the selected content by adding ‘please’ for the second direct request strategy RD2ImperativePlz, and in the third row, it adds ‘must’ to the RD3ImperativeInsist. Under indirect strategy in Table 1, the strategy of abstracting the subject by saying ‘someone’ instead of addressing the hearer directly is shown as RI4AsNegationAbsSub. An example run of a dialogue generated by the system for two agents cooking pasta is given in Table 2.

Table 2. An example run of the system for two agents cooking pasta with vegetables

Agent	Utterance	SA and Politeness Strategy
Agent1	Could you tell me if you have placed the pan on the burner?	Approval: REQUEST_INFORM
Agent2	<i>Oh yes, I have placed the pan on the burner.</i>	Direct: ACCEPT_REQUEST_INFO
Agent1	Have you turned-on the burner mate?	Approval: REQUEST_INFORM
Agent2	<i>I am not sure.</i>	Direct: REJECT_REQUEST_INFO
Agent2	<i>Could I boil the pasta in the pan for you?</i>	Autonomy: OFFER
Agent1	Alright if it is not a problem.	Autonomy: ACCEPT_OFFER
Agent2	<i>Do you know that I have chopped the vegetables with the knife?</i>	Approval: INFORM
Agent1	Ok.	Direct: ACCEPT_INFORM
Agent2	<i>Do you know that I have added the oil to the pan my friend?</i>	Approval: INFORM
Agent1	Yeah.	Direct: ACCEPT_INFORM
Agent1	I have added the vegetables to the pan.	Direct: INFORM
Agent2	<i>Alright.</i>	Direct: ACCEPT_INFORM
Agent1	Could I add the other-ingredients to the vegetables?	Approval: OFFER
Agent2	<i>That is nice of you but no please do not bother yourself.</i>	Approval: REJECT_OFFER
Agent2	<i>I am wondering whether you would like to cook the vegetables in the pan.</i>	Autonomy: REQUEST_ACT
Agent1	Please do not mind but I can not do that.	Autonomy:REJECT_REQUEST_ACT

3 Experimental Method

We conducted an experiment to study the perception of politeness by subjects in different discourse contexts, with subjects who were native speakers of English, but from two different cultural backgrounds: 11 were British and 15 were Indians, most of them students of mixed gender with an age between 20 to 30 years. Subjects were administered a web-based questionnaire and presented with a series of collaborative cooking tasks. They were asked to rate various utterances automatically generated by POLLY as though the utterance *had been said to them by their partner* in the process of cooking a recipe together. The subjects were asked to rate how polite they perceived their partner to be, on a five point Likert-like scale: Excessively Overpolite, Very Polite, Just Right, Mildly Rude or Excessively Rude. All of the tasks were selected to have relatively high R (ranking of imposition) as per B&L's theory. Requests were to 'chop the onions', 'wash the dishes', 'take out the rubbish' and 'clean the spill on the floor.' The events for the propositional content of the Inform SAs were "You have burnt the pasta", "The milk is spoilt", "You have broken the dish" and "The oven is not working". The subjects rated a total of 84 utterances spread over these eight different tasks as shown in Table 3. There was also a text box for subjects to write optional comments.

There were five experimental variables: (1) Speech act type (request vs. inform); (2) B&L politeness strategy (direct, approval, autonomy, indirect); (3) discourse context (friend vs. stranger); (4) linguistic form of the realization of the B&L strategy; (5) cultural background (Indian vs. British). The politeness strategies were selected from strategies given by B&L for each level of politeness, and are shown in Table 1. We did not manipulate the power variable of B&L.

For each task, subjects were told that the discourse situation was either cooking with a **Friend**, or with a **Stranger**. This was in order to implement B&L's D variable representing social distance. A friend has a much lower social distance than a stranger, thus Θ should be much greater for strangers than friends.

The speech acts tested were: **Request** and **Inform**. The ranking of imposition R for speech acts has Requests with higher R than Inform, so Θ should be greater for requests, implying the use of a more polite B&L strategy.

For the Request speech act, each subject judged 32 example utterances, 16 for each situation, Friend vs. Stranger. There were 4 examples of each B&L strategy, direct, approval, autonomy, indirect. The B&L strategies for requests are given in Table 1.

For the Inform speech act, subjects judged 10 example utterances for each situation, friend and stranger, with 5 B&L strategies, used to inform the hearer of some potentially face-threatening event. Of the five, there was one direct, two approval and two autonomy utterances. No Indirect strategies were used for Inform SAs because those given by B&L of hints, being vague, jokes, tautologies are not implemented in our system. The B&L strategies for Informs are also in Table 1. The distribution of the utterances used in the experiment is in Table 3.

Table 3: Distribution of the 84 utterances used in the experiment

Speech Act	Situation	Tasks	B&L Strategies				Total	
			Direct	Approval	Autonomy	Indirect		
Request	Friend	chop onions	4	4	4	4	16	64
		clean spill on floor	4	4	4	4	16	
	Stranger	wash dishes	4	4	4	4	16	
		take out rubbish	4	4	4	4	16	
Inform	Friend	oven not working	1	2	2	0	5	20
		burnt the pasta	1	2	2	0	5	
	Stranger	milk is spoilt	1	2	2	0	5	
		broken the dish	1	2	2	0	5	

4 Results and Observations

We calculated an ANOVA with B&L category, situation (friend/stranger), speech act, syntactic form, politeness formula and the nationality of subjects as the independent variables and the ratings of the perception of politeness by the subjects as the dependent variable. Results are in Tables 4, 5, and 6 and are discussed below.

B&L strategies Effect: The four B&L strategies (Direct, Approval, Autonomy and Indirect) had a significant effect on the interpretation of politeness ($df=3$, $F=407.4$, $p<0.001$). See Table 6. The overall politeness ratings from least polite to most polite were Indirect, Direct, Approval and then Autonomy strategy. See the graph in Figure 6. It must be noted that as opposed to our findings, B&L posit that the indirect strategy is the politest. This may be so because the indirect realizations that our generator produces from the AI planner are the effect-not-achieved forms like the indirect request strategies (RI1AsNegation, RI2AsModRight, RI3AsModRightAbSub and RI4AsNegationAbsSub) as shown in Table 1, which sound like a complaint or sarcasm. Other Indirect strategies given by B&L like giving hints, being vague, sarcasm or jokes are situation dependent and require general knowledge that we have not yet implemented.

Situation Effect (Friend/Stranger): Figure 7 and Table 4 show that politeness ratings averaged over the four B&L's strategies for friend, which is 3 or 'just right', is more than that for a stranger, which is a little above 2 or 'mildly rude'. Table 6 shows that utterances spoken by a friend are rated as more polite than those spoken by a stranger, for all four B&L strategies ($df=1$, $F=123.6$, $p<0.001$). This shows the effect of B&L's social distance variable, i.e. when social distance is large, a politer utterance is appropriate but if we use an utterance that assumes too much social distance, , the utterance is regarded as too polite.

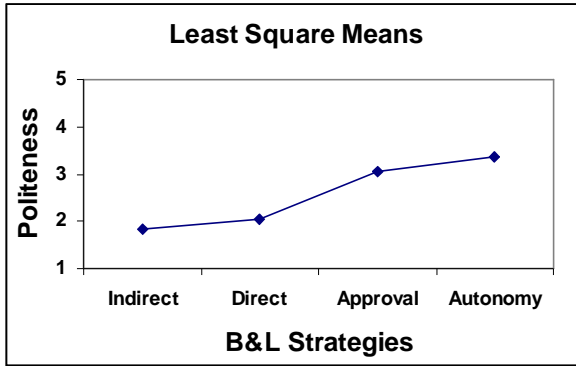


Fig. 6. The effect of B&L Strategies on politeness. 1 = Excessively Rude, 5 = Overly Polite

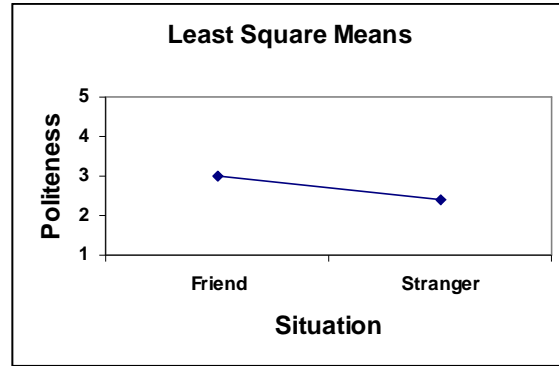


Fig. 7. The effect of situation on politeness. 1 = Excessively Rude, 5 = Overly Polite

SA Effect (Request/ Inform): Inform SA was rated as more polite than Request SA ($df=1, F=61.4, p<0.001$) as shown in Figure 8. This is in line with what B&L say about face threat, that speech acts such as Requests carry more face threat than speech acts such as Inform as they impede upon the hearer’s freedom of action and thus require more face redressing or politeness strategies.

Utterance Form Effect: We divided the utterances into four categories, used for B&L strategy realizations, as per their syntactic forms. *Queries* are those that interrogate the listener, like the Approval strategy RAp1QModAbility for requests, “Could you please wash the dishes mate?” *Assertions* in case of a request SA refer to utterances that make a request by asserting something, such as asserting that the precondition holds or asserting the ability of the hearer as in the Approval strategy RAp2AsModAbility, “If you don’t mind you can chop the onions.” For the inform SA, Assertions refer to polite declaratives that use some politeness formulas or additives with autonomy and approval strategies. On the other hand *Direct Assertions* refer to utterances that directly assert something without using much politeness and are used to realize the direct form of the Inform SA, like the ID1DirectAssert strategy, “You have burnt the pasta.” Lastly, *Imperatives* are those utterances that directly command the user to perform some action, like the RD3ImperativeInsist strategy, “You must clean the spill on the floor”.

In case of requests, Queries were rated as the politest followed by Assertions and then Imperatives ($df=2, F=279.4, p<0.001$). In case of the Inform SA, Assertions are considered to be the most polite, followed by Queries and then Direct Assertions ($df=2, F=36.0, p<0.001$). The overall order of politeness ratings is in Figure 9 with politeness scores in Tables 4 and 5.

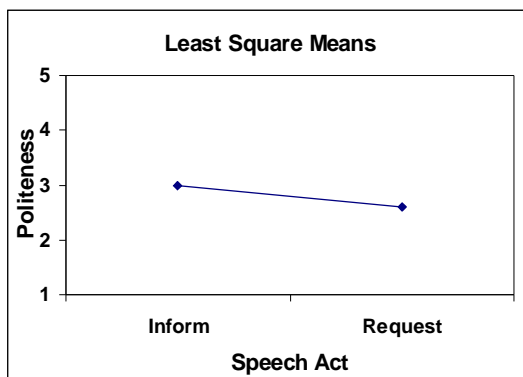


Fig. 8. The effect of speech act on politeness. 1 = Excessively Rude, 5 = Overly Polite




Fig. 9. The effect of utterance form on politeness. 1 = Excessively Rude, 5 = Overly Polite

Nationality Effect: We found that the politeness interpretations of Indian and British subjects were significantly different. Indians perceived the utterances as overall much more polite than the British did, as shown in Figure 10. This was most evident when the partner was a Friend ($df=1$, $F=6.0$, $p<0.01$), and for Requests ($df=1$, $F=6.37$, $p<0.01$) whereas perceptions were almost equal for strangers. This demonstrates a cultural effect, namely that Indian native speakers of English are more informal in their communication, especially when they are talking to a friend. Although the overall degrees of politeness of the four B&L strategies was rated higher by Indians, the order of the ranking of the strategies was the same for both Indians and British (indirect being the least polite, followed by approval, autonomy and direct) which shows that broad universality is still preserved.




Fig. 10. The effect of nationality on politeness. 1 = Excessively Rude, 5 = Overly Polite

Politeness Formula Effect: We observed that utterances with the address form 'mate', which is more specific to the British culture, were rated more polite than those without it ($df=1$, $F= 49.8$, $p<0.001$). Abstracting the subject (used in indirect strategy) made the utterance less polite ($df=1$, $F=125.0$, $p<0.001$) and adding Softeners notably increased politeness ($df=4$, $F=104.0$, $p<0.001$). In case of additives, apologies were rated to be most polite, followed by those that explicitly asked for a favour and utterances that used an insisting adverb such as 'must' were least polite ($df=3$, $F=185.6$, $p<0.001$). Average ratings of politeness formulas are in Table 5.

Table 4. Mean values of situation and utterance forms in relation to the speech acts

Situation		Request Inform	
		Friend	2.8
	Stranger	2.3	2.8
Utterance Form	Imperative	1.9	NA
	Assertion	2.4	3.2
	Queries	3.3	3.0
	Direct Assertion	NA	2.4

Table 5. Overall mean values of the utterance forms and politeness formulas

		Overall Score
Utterance Form	Imperative	1.8
	Assertion	2.5
	Queries	3.2
	Direct Assertions	2.4
Politeness Formula	AddressForm	3.1
	AbstractSubject	2.0
	Softeners	3.3
	Additives	3.0

Table 6. Mean values of the politeness ratings of SAs and situations for B&L's strategies and their overall mean score

Speech Act	Situation	Direct	Approval	Autonomy	Indirect	Overall
		Request	2.0	3.0	3.4	1.8
	Inform	2.4	3.0	3.2	NA	3.0
	Friend	2.3	3.3	3.6	2.0	3.0
	Stranger	1.8	2.8	3.1	1.7	2.4

5 Discussion and Conclusion

We presented an implementation of a system, called POLLY, that combines a general AI planner with a spoken language generator, for generating polite language as per the theory of Brown and Levinson, and demonstrated how to extract language from a plan to automatically generate task-oriented conversations [24,25]. One of the strengths of B&L's theory is the assumption that conversational agents are rational actors, who make explicit use of planning representations, with content selection for discussing actions

based on the steps, preconditions and effects of the plan. The general approach is very similar to that of Walker et al's [26] application of B&L to conversational agents for interactive narrative, but while they used a planner representation, they did not integrate a planner and their approach was not evaluated. Here, we have presented an experiment which shows that the B&L strategies have a significant effect on humans' perception of politeness. The utterances evaluated by our subjects were produced by POLLY and there was no human moderator, whereas the evaluation experiment of Cassell & Bickmore was Wizard-of-Oz [7]. As far as cultural differences are concerned, our experiment showed strong differences in the perception of politeness by Indian and British native speakers of English in the case of SAs with B&L's high ranking of imposition such as requests, and in situations where B&L's social distance variable was lower, such as when talking to a friend. In contrast, Johnson et al. [12] showed that the perceptions of politeness of American and German speakers in the domain of tutorial dialogues were identical. André et al. [1] proposed the idea of animated presentation teams for presenting information to the user but they investigated only personality and not politeness and their NLG was template based. Our generator is to be applied in the domain of teaching English as a Second Language (ESL). Previously, Porayska-Pomsta [18] applied B&L's theory in the tutorial domain for modelling teacher's corrective responses, with a generator based on case-based reasoning, selecting utterances from human-human dialogues rather than building a generator based on B&L. Johnson et al. [11] also had a similar approach for generating socially appropriate tutorial dialogue, with a template-based NLG component, for a language training system that provides training in a foreign language and culture. Their goals are similar to ours, but their language courses have a strong task-based focus, on skills needed to cope with specific situations such as introducing yourself, obtaining directions and arranging meetings. Rehm and Andre [21] have shown that the interpretation of politeness strategies is affected by the gestures used in an embodied conversational agent, and we have not yet carried out experiments on an embodied agent.

One apparent limitation of our work to date, is that an examination of the utterances in the example run in Table 2 suggests that the language is somewhat inappropriate. This is because this system presents a design for a plug-in that can be used in different applications and for demonstration purpose we have kept the utterance generation random. Because of this random nature, there arise two problems. One, because of the random retrieval of utterances from the set of DSyntS for each B&L category, some of the sequences seem inconsistent in a particular context. Each set of utterances under a specific category includes the various realizations given by B&L and these realizations may not always be appropriate in every context; a limitation of B&L's theory is that its equation it does not take the context into account. Second, random selection may result in a repetition of the type of strategy realization.

In future work, we aim to modify the language generator to make it more robust, explore B&L's power variable, as well as the interplay of the three sociological variables and integrate POLLY into a virtual reality environment for learning politeness while learning English as a second language.

References

1. André, E., Rist, T., Mulken, S.v., Klesen, M., & Baldes, S. The automated design of believable dialogues for animated presentation teams. In *Embodied Conversational Agents* (pp. 220–255). MIT Press (2000)
2. André, E., Dybkjær, E., Minker, W., Heisterkamp, P. (Eds.) *Affective Dialogue Systems, Tutorial and Research Workshop, ADS 2004, Proceedings. Lecture Notes in Computer Science 3068* Springer (2004 a)
3. André, E., Rehm, M., Minker, W., & Buhler, D. Endowing spoken language dialogue systems with emotional intelligence. In André, E., L. Dybkjær, W. Minker and P. Heisterkamp (Ed.), *Affective Dialogue Systems, Springer Verlag* (2004 b). (pp. 178--187)
4. Berk, Lynn M. *English syntax: from word to discourse*, Oxford University Press (1999)
5. Blum, A., Furst, M. Fast Planning Through Planning Graph Analysis. *Artificial Intelligence* (1997) 90:281-300.
6. Brown, P. & Levinson, S. *Politeness: Some Universals in Language Usage*. Cambridge u.a.: Cambridge University Press (1987)
7. Cassell, J. Bickmore, T. W. Negotiated Collusion: Modeling Social Language and its Relationship Effects in Intelligent Agents. *User Model. User-Adapt.Interact.* (2003) 13(1-2):89-132.
8. Dale, R. and Reiter, E. *Building Natural Language Generation Systems*. Studies in Natural Language Processing. Cambridge University Press (1995)
9. Goffman, E. *The presentation of self in everyday life*, Doubleday Press, 1959.

10. Grosz, B. J. and Sidner, C. L. Plans for discourse. In P. R. Cohen, J. L. Morgan, and M. E. Pollack, editors, *Intentions in Communication*, pages 417- 444. MIT Press, Cambridge, MA (1990)
11. Johnson, L.W. and Rizzo, P. and Bosma, W.E. and Ghijsen, M. and van Welbergen, H. Generating socially appropriate tutorial dialog. In: *ISCA Workshop on Affective Dialogue Systems*, (2004) pp. 254-264.
12. Johnson, L., Mayer, R., André, E., & Rehm, M. Cross-cultural evaluation of politeness in tactics for pedagogical agents. *Proc. of the 12th Int. Conf. on Artificial Intelligence in Education* (2005)
13. Lavoie, B., and Rambow, O. RealPro – a fast, portable sentence realizer. In *Proceedings of the Conference on Applied Natural Language Processing (ANLP'97)*, Washington DC (1997)
14. Mairesse, F. and Walker, M.A. Words Mark the Nerds: Computational Models of Personality Recognition through Language. In *Proceedings of the 28th Annual Conference of the Cognitive Science Society (CogSci 2006)*.
15. Mairesse F. and Walker M.A. PERSONAGE: Personality Generation for Dialogue. In *Proceedings of the 45th Annual Meeting of the Association for Computational Linguistics. ACL* (2007).
16. Morand, D. A., and Ocker, R. J. Politeness theory and Computer-Mediated communication: A Sociolinguistic Approach to Analyzing Relational Messages. *Proceeding of the 36th Hawaii International Conference on System Sciences, HICSS* (2003)
17. Piwek, P. An Annotated Bibliography of Affective Natural Language Generation. Technical Report ITRI-02-02, University of Brighton (2002)
18. Porayska-Pomsta, K. Influence of Situational Context on Language Production: Modelling Teachers' Corrective Responses. PhD Thesis. School of Informatics, University of Edinburgh (2003)
19. Porayska-Pomsta, K. and Pain, H. Providing Cognitive and Affective Scaffolding through Teaching Strategies, *Proceedings of the 7th International Conference on Intelligent Tutoring Systems (ITS)*, (2004)
20. Reeves, B. and Nass, C. *The Media Equation*. University of Chicago Press (1996)
21. Rehm, M. and Andre, E. Informing the Design of Agents by Corpus Analysis. *Conversational Informatics*, T. Nishida Ed. (2007)
22. Roman, N., Piwek, P., Carvalho, A. Politeness and Bias in Dialogue Summarization: Two Exploratory Studies. In: J. Shanahan, Y. Qu, J. Wiebe (Eds.), *Computing Attitude and Affect in Text: Theory and Applications*, The Information Retrieval Series, Vol. 20, Dordrecht: Springer (2006)
23. Romano, D.M. Virtual Reality Therapy. *Developmental Medicine & Child Neurology Journal*. (2005) Sep; 47(9):580
24. Sidner, C. L. An artificial discourse language for collaborative negotiation. In *Proc. 12th National Conf. on AI* (1994) pages 814–819,
25. Walker, M.A. The effect of resource limits and task complexity on collaborative planning in dialogue. *Artificial Intelligence Journal*, 85:1-2, 1996.
26. Walker, M., Cahn, J. and Whittaker, S. J. Improving linguistic style: Social and affective bases for agent personality. In *Proc. Autonomous Agents'97* (1997) 96–105. ACM Press
27. Wang, N., Johnson, W.L., Rizzo, P., Shaw,E., & Mayer, R. Experimental evaluation of polite interaction tactics for pedagogical agents. *Proceedings of IUI '05*. (2005)