

Improvising Linguistic Style: Social and Affective Bases for Agent Personality

Marilyn A. Walker
AT&T Labs Research
600 Mountain Ave.
Murray Hill, NJ 07974
walker@research.att.com

Janet E. Cahn
MIT Media Lab
20 Ames St.
Cambridge, MA 02139
cahn@media.mit.edu

Stephen J. Whittaker
AT&T Labs Research *†
600 Mountain Ave.
Murray Hill, NJ 07974
whittaker@research.att.com

Abstract

This paper introduces Linguistic Style Improvisation, a theory and set of algorithms for improvisation of spoken utterances by artificial agents, with applications to interactive story and dialogue systems. We argue that linguistic style is a key aspect of character, and show how speech act representations common in AI can provide abstract representations from which computer characters can improvise. We show that the mechanisms proposed introduce the possibility of socially oriented agents, meet the requirements that lifelike characters be believable, and satisfy particular criteria for improvisation proposed by Hayes-Roth.

Introduction

Just because you are a character doesn't mean that you have character. Wolf to Raquel in *Pulp Fiction*, Q. Tarantino.

Linguistic Style Improvisation (henceforth LSI) concerns the choices that speakers make about the SEMANTIC CONTENT, SYNTACTIC FORM and ACOUSTICAL REALIZATION of their spoken utterances. This paper argues that linguistic style is a key aspect of an agent's character. We present a novel theory of, and algorithms for, Linguistic Style Improvisation by computer characters.

As an example of how linguistic style can convey character, consider Victor Laszlo's request for two

This research was supported in part by Hewlett Packard Laboratories U.K., by the University of Pennsylvania, by Mitsubishi Electric Research Laboratories, by the MIT Media Laboratory, and by AT&T Research Laboratories.

Permission to copy without fee all or part of this material is granted provided that the copies are not made or distributed for direct commercial advantage, the ACM copyright notice and the title of the publication and its date appear, and notice is given that copying is by permission of ACM. To copy otherwise, or to republish, requires a fee and/or specific permission. Agents '97 Conference Proceedings, copyright 1997 ACM.

cointreaux in 1, from the *Casablanca* screenplay in Figure 1. In the film, this request is delivered in pleasant tones.

(Laszlo and Ilsa enter Rick's Cafe)
Headwaiter: Yes, M'sieur?
Laszlo: I reserved a table. Victor Laszlo.
Waiter: Yes, M'sieur Laszlo. Right this way.
(Laszlo and Ilsa follow the waiter to a table)
Laszlo: Two cointreaux, please.
Waiter: Yes, M'sieur.
Laszlo: (to Ilsa) I saw no one of Ugarte's description.
Ilsa: Victor, I feel somehow we shouldn't stay here.

Figure 1: Excerpt from the *Casablanca* script.

- (1) a. Two cointreaux, please.

However, consider the alternative stylistic realizations in 2 for requesting two cointreaux:

- (2) a. Bring us two cointreaux, right away.
b. You must bring us two cointreaux.
c. We don't have two cointreaux, yet.
d. You wouldn't want to bring us two cointreaux, would you?

Clearly, speakers make stylistic choices when they realize their communicative intentions, and their realizations express their character and personality. And, based on these stylistic realizations, listeners draw inferences about the character and the personality of the speaker. Thus, algorithms for LSI are important for any domain in which agents speak, such as characters for interactive drama systems, multimodal interface agents and spoken dialogue agents (CPB⁺94; LB95; RWS⁺94; MDBP94; HRB94; Kam95).

Our work on LSI draws from two theoretical bases: computational work on SPEECH ACTS (All79; Coh78;

header:	REQUEST-ACT(speaker, hearer, action)
precondition:	WANT(speaker,action) CANDO(hearer,action)
decomposition-1:	surface-request(speaker,hearer,action)
decomposition-2:	surface-request(speaker,hearer, INFORMIF (hearer,speaker, CANDO(hearer,action)))
decomposition-3:	surface-inform(speaker,hearer, $\neg$ (CANDO(speaker,action)))
decomposition-4:	surface-inform(speaker,hearer, WANT(speaker,action))
effects:	WANT(hearer,action) KNOW(hearer, WANT(speaker, action))
constraint:	AGENT (action,hearer)

Figure 2: Definition of the REQUEST-ACT plan operator from Litman and Allen, 1990

Lit85), and social anthropology and linguistics research on social interaction(Gof83; BL87). The *Speech Acts* section introduces the components of speech act theory that we draw on; the *Social Interaction and Linguistic Style* section discusses in detail Brown and Levinson’s theory of linguistic social interaction. We argue that these two theories in combination yield a rich generative source of different characterizations for artificial agents. The *Computing Linguistic Style* section then explains how these theories provide the basis for generating the improvisations such as those in 2, above. The *Implementing Emotional Dispositions* section discusses how we augment these improvisations by selecting for the speaker an emotional disposition and its attendant acoustical correlates(Cah90). The *Examples* section illustrates how the theory is implemented in the domain of interactive story and dialogue. Finally we discuss how LSI extends and differs from other recent approaches to both interactive drama and text generation and propose useful extensions to our current work.

Speech Acts

Speech acts were first proposed as a small set of communicative intentions such as REQUEST or INFORM that underly all utterance production(Sea75). In any language based application, interactive dialogue can be represented as sequences of speech acts by multiple characters. Therefore, LSI uses speech acts as the abstract representation for utterances, and plans as the basis for improvisation — each spoken utterance is represented as an instantiation of a plan operator and these instantiations are interleaved with descriptions of physical acts in a real or simulated world.

The inventory of speech acts is defined by the application. Ours consists of the initiating acts of INFORM, OFFER and two types of REQUEST: REQUEST-INFO and REQUEST-ACT. We also use three types of response speech acts for acceptance and rejection,

corresponding to each major type of initiating act: ACCEPT-INFORM, ACCEPT-OFFER AND ACCEPT-REQUEST; and REJECT-INFORM, REJECT-OFFER and REJECT-REQUEST.

Each speech act definition includes (a) specifying the conditions under which a speaker performing the speech act could be successful at achieving a communicative intention, and (b) specifying the effects on the hearer if the speaker is successful. Earlier computational work proposed that speech acts should be implemented in a standard AI planning system as plan operators that include the act’s DECOMPOSITION, PRECONDITIONS and EFFECTS, thereby enabling computer agents to plan utterances in the same way that they plan physical acts(All79; Coh78; Lit85). An example plan-based representation of a REQUEST-ACT (for example, Laszlo’s request in 1a) based on Litman and Allen’s work, is given in Figure 2(LA90).

A critical basis of our improvisation algorithms is speech act theory’s distinction between the underlying intention of a speech act, and the surface forms of the utterance that can REALIZE the speech act. This distinction is seen in Figure 2: the REQUEST-ACT speech act specifies an underlying intention (the desired effect) of the speaker getting the hearer to do (or want to do) a particular action; while the four decompositions specify the different ways that the underlying speech act can be realized by *surface* speech acts, that is, by particular sentential forms such as declarative sentences or questions. For example, the sentential equivalents of decompositions 1 to 4 in Figure 2 might be those in 3a to 3d respectively, where *action* represents an action description:

- (3) a. Do *action*.
- b. Can you do *action*?
- c. I can’t do *action*.

d. I want *action*.

Our algorithms for improvisation, to be discussed in the *Computing Linguistic Style* section, are mechanisms for deciding how to realize a given underlying intention as a particular surface form. While previous work on dialogue generation has focused on informational motivations and effects (MP93), we focus here on the impact of social and affective parameters on the selection of utterance form and content.

Social Interaction and Linguistic Style

Whenever agents realize a particular speech act, they make choices about the linguistic style with which that act is realized. Our main idea is that all these choices have a major effect on our perception of an agent's character and personality. Given the goal of achieving a particular communicative intention in a given social setting, an agent must choose among all the possible variations in SEMANTIC CONTENT, SYNTACTIC FORM and ACOUSTICAL REALIZATION. We call these choices a STRATEGY for realizing a particular communicative intention.

The generative account we present is derived from Brown and Levinson's theory of social interaction (BL87) in which they identify a number of different variables and give examples of how different values for the variables produce different communicative outcomes. In LSI, we take their framework, refine its specification where necessary, and specify the computational mechanisms required to implement it.¹

Maintaining public face An important basis of the theory is that all agents have and know each other to have:

1. FACE: An agent's public self image, which consists of the desire for:
 - (a) AUTONOMY: Freedom of action and freedom from imposition by other agents;
 - (b) APPROVAL: A positive consistent self-image or personality that is appreciated and approved of by other agents;
2. Capabilities for RATIONAL REASONING such as means-end reasoning, deliberation, and plan recognition.

Social variables and face Given the desire to maintain their own and others' face, and beliefs about their own and others' rationality, the agents' algorithm

¹Due to space constraints, we are unable here to present a full exegesis of their theory, the interested reader is referred to (BL87).

for choosing a strategy for realizing a particular speech act relies on evaluating three socially determined variables:

1. D(S,H): the SOCIAL DISTANCE between the speaker and hearer.
2. P(H,S): the POWER that the hearer has over the speaker.
3. R_α : a RANKING OF IMPOSITION for the act α under discussion.

Human agents use personal experience, background knowledge, and cultural norms to determine the values for these variables. For example, SOCIAL DISTANCE often depends on how well S and H know one another, but also on social class and status. POWER comes from many sources, but often arises from the ability of S to control access to goods that H wants, such as money.

The RANKING OF IMPOSITION relies on the fact that all agents' basic desires include the desire for autonomy and approval. Thus particular speech act types can be ranked as higher impositions simply by how they relate to agents' basic desires.

Speech acts that can function as a threat to H's desire for autonomy include those that predicate some future act of H, as well as speech acts that predicate some future act of S toward H, such as offers, which put pressure on H to accept or reject them. This means that the act types of REQUEST-INFORM, REQUEST-ACT and OFFER threaten H's desire for autonomy. The INFORM speech act also threatens H's desire for autonomy on the basis that it is an attempt by S to affect H's mental state.

Speech acts that threaten H's desire for approval include all rejections, including the act types REJECT-INFORM, REJECT-OFFER and REJECT-REQUEST.²

Given our inventory of speech acts, and the range of the variables D and P, we instantiate the theory with the ranking of imposition R_α based on the speech act type, as shown in Figure 3 below.³

Linguistic style strategies and social variables

As social and rational actors, S and H attempt to avoid threats to one another's face. Given values for social

² Other speech acts not in our inventory, such as criticisms and complaints, also threaten H's desire for approval (BL87).

³The values we use here serve to illustrate the model and range of phenomena. The actual values of the ranking of imposition need to be empirically determined with respect to the culture being modeled. We also discuss in our concluding section how R_α should be a function of both speech act type and propositional content, rather than purely speech act type as we do here.

Speech Act	R_α
accept-request	5
accept-inform	5
accept-offer	10
inform	15
request-info	20
offer	25
reject-offer	30
reject-inform	35
reject-request	40
request-act	45

Figure 3: A ranking R_α on imposition of various types of speech acts with values from 1 to 50.

distance $D(S,H)$, power $P(H,S)$ and ranking of imposition R_α , the agent S estimates the THREAT Θ to H of performing the speech act α by simply summing these variables as in equation 4:

$$(4) \quad \Theta = D(S,H) + P(H,S) + R_\alpha$$

Once a value for Θ has been calculated, the agent uses it to choose among one of the following four strategies for executing a speech act:⁴

- (5) a. DIRECT: Do the act directly.
- b. APPROVAL-ORIENTED: Orient the realization of the act to H 's desire for approval;
- c. AUTONOMY-ORIENTED: Orient the realization of the act to H 's desire for autonomy;
- d. OFF-RECORD: Do the act off record by hinting, and/or by ensuring that the interpretation of the utterance is ambiguous.

The lowest values of Θ lead to the DIRECT strategy and higher values lead to the OFF-RECORD strategy. In LSI, the range for each of the social variables D , P and R_α is between 0 and 50. Therefore, the Θ sum will range from 0 to 150. DIRECT strategies correspond to Θ values through 50, APPROVAL-ORIENTED strategies to Θ values from 51 to 80, AUTONOMY-ORIENTED strategies for Θ values from 81 to 120 and OFF-RECORD strategies for Θ values from 121 to 150.⁵

Each strategy can be realized by a wide range of sub-strategies, whose SEMANTIC CONTENT is selected from the plan-based representation for a speech act and whose SYNTACTIC FORM is selected from a library

⁴Brown and Levinson include a strategy of not executing the speech act at all because the face threat is too great.

⁵Again these values are estimates selected for illustrative purposes.

of syntactic forms. And since there are many ways to realize each strategy, realizations within particular ranges are heuristically assigned to the upper or lower end of the scale, or assigned to the same values of the scale to support random variation.

Emotion as an element of linguistic style

Varying the affect of the spoken realization is a critical aspect of linguistic style. Although Brown and Levinson state that expressions of strong emotion threaten both S and H 's desires for approval and autonomy, they do not further specify the relation between strategies for selecting SEMANTIC CONTENT and SYNTACTIC FORM, and those for selecting the ACOUSTICAL REALIZATIONS in the utterance which most directly express emotions. In order to explore this interaction, we adopt a very simple view of emotional expression: emotional disposition is an orthogonal dimension to social variables, and each character is simply assigned an emotional disposition at the start.

Computing Linguistic Style

Because LSI is defined on the basis of speech act types alone, what we have described so far is domain independent. However, the content of each speech act is domain specific. For example, in Figure 2, domain specific content is represented by the **action** variable in the definition of REQUEST-ACT. Similarly, the domain specific content in an INFORM speech act is represented by a **proposition** variable. Thus to test LSI, specific domains must be represented in terms of the actions and propositions of that domain. For example, Figure 6 represents the domain specific action of serving two cointreaux.

We have tested LSI on speech acts derived from two domains: a task-oriented dialogue in which two agents discuss furnishing a two room house (Wal96a), and a segment of the *Casablanca* script shown in Figure 1 (Wal96b; WABM95).⁶

As shown in Figure 4, LSI takes an input a sequence of speech acts representing a dialogue, and a SOCIAL STRUCTURE which consists of a value between 0 and 50, for both social distance D and for power P , for each pair of agents in the dialogue. Then, for each speech act in the script or the dialogue, the speaker

⁶The task oriented dialogue representation is generated off-line by a planner, while the *Casablanca* script speech act representation is constructed by hand. In both cases, we use the generator FUF(Elh92) to generate surface forms. Because FUF does not operate directly on predicate logic representations used in plans, we therefore augment these with manually generated FUF equivalents. Future implementations will include a transducer that generates FUF forms automatically from plan representations.




Figure 4: Overview of LSI Algorithm

header:	SERVE(waiter, customer, two-cointreaux)
precondition:	HAS(restaurant, two-cointreaux)
decomposition:	BRING(waiter, customer, two-cointreaux)
effects:	HAS(customer, two-cointreaux)

Figure 5: A possible plan in the restaurant domain for serving two cointreaux

determines the social distance D between him/herself and the hearer, the power P that the hearer has over him/her, and the value on R_α for the speech act type as in Figure 3. Then by equation 4, the speaker calculates the value of Θ , and uses this to select one of the strategies given in 5 above.

We will now demonstrate how the algorithm operates, by showing how different linguistic strategies result from different social structures. In each case we will use the example from *Casablanca*, in which Laszlo orders two cointreaux from Emil, and assume that the algorithm operates on the representations in Figures 2 and 6.⁷ Since there are many more realizations of the strategies than can be discussed here, interested readers are referred to (BL87).

Direct strategies

Direct strategies result from social structures in which both social distance D and power P are small. In the case of our two cointreaux example, imagine that Laszlo and Emil are old friends, and that Emil, as the waiter, has no power over Laszlo. This could be modeled in our framework with a social structure in which

⁷Actually we will derive some of the decompositions in Litman's definition by rule(AP80; GL71).

the social distance D between Emil and Laszlo is 4 and the power P that Emil has over Laszlo is 0. According to Figure 3, the R_α for REQUEST-ACT is 45. Using equation 4 and the values for P , D and R_α , the value for Θ is 49, leading Laszlo to select a direct form strategy for realizing his request.

The realizations for all DIRECT forms, irrespective of speech act type, are based on the SEMANTIC CONTENT of the decomposition step of the speech act. Each speech act type has an associated *default* SYNTACTIC FORM.

For example, in the case of REQUEST-ACTS we assume that the default syntactic form is an imperative.⁸ Thus the simplest strategy for realizing a direct form is the REALIZE-DIRECT-STRATEGY: Realize the content of the decomposition step with its associated default syntactic form. For a request such as *Two cointreaux, please*, this would result in an utterance such as:

(6) Bring us two cointreaux.

Direct realizations can also be ordered within the

⁸For speech acts such as INFORM, the default syntactic form is a declarative sentence, and for speech acts which are subtypes of ACCEPT or REJECT, the default forms are *Okay*, *Yes* or *No*, respectively.

range of 0 to 50 so that lower values correspond to styles that convey that H has no power (P is low). One way to make a REQUEST-ACT is the POWER-DIRECT-STRATEGY: Add *you must* or *right away* to the direct form. This is illustrated in 7 and 8:

(7) Bring us two cointreaux right away.

(8) You must bring us two cointreaux.

Approval oriented strategies

Approval oriented strategies result from social structures in which there are minor differences in both power P and social distance D between the interactants, so that these factors play a weak role in strategy selection. Strategies for orienting the realization of a speech act to the hearer's desire for approval include intensifying interest or attention to H, implying that S and H are cooperators who have the same perspective or desires, or conveying that S and H are part of the same social group or are friends.

One way to convey that S and H have the same desires when making a request is the OPTIMISM-APPROVAL-STRATEGY: S expresses optimism that H will want to do what S wants H to do. This strategy results from selecting the semantic content to be realized from the *want hearer action* effect of the request-act (as in Figure 2)⁹, and realizing this semantic content with a declarative sentence that includes a tag question. This strategy results in surface forms such as:

(9) You'd like to bring us two cointreaux, wouldn't you?

One way to imply that S and H are in the same social group and that S believes that the relative P between himself and H is small is the GROUP-APPROVAL-STRATEGY: Use in-group address forms such as *buddy*, *mate*, *honey*, *doll*, *my man*, depending on the group. For a request, this is implemented by concatenating an in-group address form, *my man*, to the direct realization of the speech act, resulting in surface forms such as:

(10) Hey Emil, my man, bring us two cointreaux.

For ACCEPT-OFFER or ACCEPT-REQUEST speech acts, approval oriented forms are those that explicitly assert the WANT effect of the offer or request speech act, such as:

(11) I'd be glad to.

⁹A similar strategy of assuming that the effect already holds can also be used for INFORM speech acts.

and

(12) With pleasure.

For rejections, approval oriented forms are those by which H affirms a social relationship with S such as:

(13) I'm sorry, I can't. Normally I'd love to.

Autonomy oriented strategies

Autonomy oriented strategies result from social structures in which there are significant differences between the two agents in either power P or social distance D. Under these circumstances S will choose strategies that make minimal assumptions about H's wants and desires, leaving H the option not to do the act, and disassociate S from possible infringement of H's autonomy.

Note that the effect field in Figure 2 encodes information about H's wants and desires. Thus, one rule is to be pessimistic about H's desires. This can be achieved by selecting semantic content from this effect field with the NEGATE-EFFECT-AUTONOMY-STRATEGY: State that the want effect doesn't hold. This produces a form such as:

(14) You wouldn't want to bring us two cointreaux would you?

In addition, note that the precondition field in Figure 2 encodes information about H's abilities. One way of leaving H the option not to do the act is for S to produce a query with this precondition as the semantic content, leaving H the option of saying that s/he is unable to do the act. This is the QUERY-ABILITY-AUTONOMY-STRATEGY, which results in forms such as:

(15) Can you bring us two cointreaux?

One way of disassociating S and H from an autonomy infringement is to produce an indirect form of a request with the ASSERT-WANT-PRECONDITION-AUTONOMY-STRATEGY: State that the want precondition holds. This results in forms such as:

(16) We'd like two cointreaux.

Another strategy for avoiding an autonomy infringement is the IMPERSONALIZE-ACTOR-AUTONOMY-STRATEGY: Impersonalize who actually performs the requested act. This results in proposals with no actor specified. It is also possible to produce proposals in which the act itself is unspecified, by selecting the semantic content for the request from the effect field of the domain act. For example, in Figure 6, the effect is that the customer has two cointreaux. Using this field as the semantic content results in surface forms such as:

(17) Let us have two cointreaux.

INFORM speech acts also have realizations that are autonomy oriented. An INFORM speech act can impinge on H's autonomy concerning what s/he wants to believe. One way to orient to H's autonomy is to soften the strength of an assertion by HEDGING it (PFB80). For example, consider Laszlo's utterance of *I reserved a table*. This can be hedged by simply embedding the declarative sentence, which is produced from the decomposition step of the plan for an INFORM, with hedging phrases such as *I feel*, *I believe*, *It seems*, *As you may know*, *I think*, *I heard*, or adding other hedges such as *somehow*, *sort of*, *kind of* to the verb phrase. This strategy is encapsulated in 18 and produces forms such as 19:

(18) HEDGE-INFORM-STRATEGY: Augment any inform statement with either a pre-sentential or a verbal hedge.

(19) I believe I reserved a table.

An example of hedging in the original script (Figure 1) is Ilsa's assertion:

(20) Victor, I feel somehow we shouldn't stay here.

Hedging the strength of the assertion can also function as an approval oriented strategy since it is a simple way to avoid disagreement.

Off record strategies

Off record strategies result from social situations in which there are significant values for social distance D or major discrepancies in power P between two agents, or from an act that is a large imposition on H. Tactics for going off record are difficult to implement because strategies for doing so involve indirect inference paths that are difficult to model computationally. There are, however, several simple ways to make a request off record by constructing hints from plan-based representations. One strategy is the ASSERT-NEGATION-DOMAIN-EFFECT-STRATEGY, in which S asserts that the effect of the domain plan does not hold, as in:

(21) We don't have two cointreaux yet.

Another strategy is the ASSERT-DOMAIN-PRECONDITION-HOLDS-STRATEGY: Assert that the precondition of the domain plan holds. For example, Laszlo's utterance of *I reserved a table* is a statement that the domain precondition for being shown to a table holds. Thus the original realization in the script is an off record form.

Another strategy is the ABSTRACT-AGENT-AND-NEGATE-EFFECT-STRATEGY: Select the semantic content as the decomposition of the domain plan, abstract the agent role, and negate the assertion of the decomposition. This leads to an implicature (Hir85). The result is shown below:

(22) Someone hasn't brought us two cointreaux.

In the current implementation of LSI, autonomy oriented forms are sometimes substituted for off record forms in order to provide more variability when characters choose to go off record.

Implementing Emotional Dispositions

Once a character's emotional disposition has been set, all of that character's utterances are synthesized with the acoustical correlates of that emotion. We implement this by drawing on Cahn's theory of expressing affect in synthesized speech (Cah90), and use a version of her Affect Editor program developed expressly for interactive theater and simulated conversation.

The Affect Editor computes instructions for a speech synthesizer (so far, the DECtalk3 and 4.1) so that it produces emotional and expressive synthesized speech. The output is a set of synthesizer instructions; the input is a combination of text and acoustical parameter values. The parameters (seventeen in all) control the presence in the speech signal of various aspects of pitch, timing, voice quality and phoneme quality.

Because some of the acoustical properties are moderated by linguistic properties of the text, the words in the text must be annotated for part of speech, focus information (expressed as a likelihood of receiving intonational stress, that is, as the inverse of the accessibility of items in memory), and then the text itself marked with all possible phrase boundaries according to syntax and grammatical role.

The acoustical parameters have numerical values. Their adjustment around zero — representing neutral affect — allows various shadings of emotional expression, for example, from calm to sad to completely dejected, or from enthusiasm to harsh anger. Our current LSI implementations make use of parameter value sets for seven emotional dispositions: Angry, Annoyed, Disgusted, Distracted, Gruff, Pleasant and Sad.

Example Runs of Linguistic Style Improvisation

To demonstrate the effect of LSI, we apply it to the first five lines of the *Casablanca* script in Figure 1, where agent A is Laszlo and agent B is the waiter. We provide an underlying abstract representation for this excerpt in terms of speech acts as specified in Figure 6.

(Laszlo and Ilsa enter Rick's Cafe)	
Headwaiter: Yes, M'sieur?	(OFFER)
Laszlo: I reserved a table. Victor Laszlo.	(REQUEST-ACT)
Waiter: Yes, M'sieur Laszlo. Right this way.	(ACCEPT-REQUEST)
(Laszlo and Ilsa follow the waiter to a table)	
Laszlo: Two cointreaux, please.	(REQUEST-ACT)
Waiter: Yes, M'sieur.	(ACCEPT-REQUEST)

Figure 6: Assumed Speech Acts for an excerpt from the *Casablanca* script.

We use extreme power and social distance parameter settings in the examples to demonstrate the range of variation that is possible.

A direct/angry speaker with an approval-oriented/pleasant hearer In a social structure in which A's emotional disposition is angry, and B's is pleasant, modeled by setting $D(A,B) = 0$, $P(B,A) = 0$, $D(B,A) = 30$, and $P(A,B) = 30$, A will choose direct strategies and an angry delivery, and B will choose approval oriented strategies, delivered in pleasant tones. The result of this social structure applied to the *Casablanca* excerpt is:

(23) W: Could I help you?

L: You must take us to a table. I am Victor Laszlo.

W: It's a pleasure.

L: Bring us two cointreaux, right away

W: I'd be glad to.

An autonomy-oriented/distraught speaker with a direct/pleasant hearer In a social structure where A's emotional disposition is distraught, and B's is pleasant, modeled by setting $D(A,B) = 40$, $P(B,A) = 40$, $D(B,A) = 0$, and $P(A,B) = 0$, A will choose autonomy oriented strategies and a distraught delivery. and B will choose the lower end of direct strategies and a pleasant delivery. The effect of this social structure on the *Casablanca* excerpt is:

(24) W: I will help you

L: Can you take us to a table? As you may know, I am Victor Laszlo

W: Yes, if you insist.

L: You wouldn't want to bring us two cointreaux, would you?

W: Yes, if I must.

The values that produce 24 portray Laszlo as a wimp, for several reasons. First, Laszlo, who is the customer, is orienting to the waiter's autonomy. Second, the distraught delivery is very high pitched and tentative. Finally, the fact that the waiter is rude highlights their differences in linguistic style.

Related Work

There are two areas of related work: recent work on interactive drama systems—in particular, Hayes-Roth's work on improvisation by computer characters; and the longer running body of work on natural language generation.

Interactive drama systems In empirical studies of human reactions to lifelike computer characters, Nass *et al.* (NST95) show that linguistic style leads to specific inferences about character. However, they rely on pre-scripted linguistic forms to demonstrate its effects and no generative mechanism is supplied. Other work in this area, for example, that of Maes *et al.* (MDBP94) and Loyall and Bates (LB95) has focused on the behavior on non-speaking animals, so that linguistic style has not been considered. Where characters do speak, their utterances are in the main pre-scripted (BL93), or generation does not focus on variations in linguistic style (CPB⁺94).

Hayes-Roth's work on improvisation does allow for linguistic variation, but this arises by selection from a finite set of forms, and again no generative mechanism is given (HRB94; HRBS95). However this work provides a useful set of requirements for improvisation mechanisms of computer characters (HRBS95), which our mechanisms for LSI satisfy:

1. *Interesting variability* in a character's interpretation of a given direction on different occasions;
2. *Random variability* in the way a character performs a specific behavior on different occasions;
3. *Idiosyncrasies* in the behaviors of different characters;

4. *Plausible motivations* for character's behavior;
5. *Recognizable emotions* associated with character's behaviors and interactions.

The dialogues in 23 and 24 demonstrate that social structure variables produce *interesting variability*, *random variability*, and *idiosyncrasies*. In addition, because Brown and Levinson's theory is based on empirical observation of human interaction in many cultures, a theory of LSI based on it satisfies Hayes-Roth's last two criteria. Since the theory captures linguistic universals, human users should be able to ascribe plausible motivations and recognize the emotions associated with a character's behavior. Especially, the motivations the theory ascribes are not only descriptive and explanatory, but *predictive* and *generative*.

Text generation Previous work on natural language generation has addressed the problems of how surface forms can be generated from underlying speech acts (MP93; Coh78; Dal88), *inter alia*. However in the main, the variables that determine linguistic choice have all been task-related. The generation research has therefore addressed the role of linguistic choice in indicating information structure; foregrounding and backgrounding information; reducing cognitive overload, and the impact of these factors on inducing change in the hearer's beliefs. This task oriented perspective ignores other aspects of choice and interaction, namely, agents' motivations, and socially appropriate responses and behavior.

One exception is the work of Hovy (Hov93), who does consider the effect of social factors on generation. However, Hovy is concerned with generating news stories (text) which, in speech act terms are sequences of INFORM speech acts. In contrast, our work focuses on the generation of conversation, which requires a much wider range of speech acts. Furthermore, the news story genre affords fewer opportunities for social factors to affect generation given the anonymity of the generic text reader.

Discussion and Future Work

In this paper, we have argued that linguistic style is an under-researched aspect of character, and presented a theory of, and algorithms for, Linguistic Style Improvisation by computer characters. This work expands the set of parameters that have been investigated in research on natural language generation of conversational speech.

Possible interesting extensions to our work would be to introduce social feedback into our model, allowing linguistic actions to directly affect the SOCIAL STRUCTURE in the course of an interaction. We hope to ex-

plore a reciprocal feedback loop to social structure, in which, for example, one agent's linguistic friendliness results in another agent adjusting their beliefs about social distance, and hence changing the second agent's future linguistic strategies. This should result in interpretable and interesting changes in the way two agents treat one another over the course of a social interaction. We also hope to examine in more detail the relationship of acoustical expression of emotions to choices about linguistic semantic content and syntactic form.

Another possible extension concerns a more complex function for calculating the RANKING OF IMPOSITION R_α . The problem is that R_α should be a function of *both* the speech act type, and the type of the action *in the domain*. For example, a REQUEST-ACT that H pass the salt is less of an imposition than a REQUEST-ACT that H give S five dollars. We conjecture that a function for R_α could be based on inputs α and a domain act δ , if the speech act planner could access information about the effort involved with the execution of the domain act δ .

In sum, we have shown how LSI can be applied to computer characters in both interactive fiction and task-oriented dialogue simulation. In future work, we hope to investigate applying the same mechanisms to characters for personal assistants for spoken language interfaces (BL93; Kam95; YLM95). We believe that the combination of dimensions we have focused on provides a motivated and artistically interesting basis for making choices about linguistic style, that these choices are closely related to human perceptions of character and personality, and that they provide a rich generative source of linguistic behaviors for life-like computer characters.

Acknowledgments

Thanks to Gene Ball, Justine Cassell, Larry Friedlander, Dawn Griesbach, Don Martinelli, Phil Stenton, and Loren Terveen for interesting discussions about linguistic style, to Charles Callaway and Michael Elhadad for consultation on using FUF, to Obed Torres for implementation work on VIVA, and to Wendy Plesniak for artistic inspiration.

References

- [All79] James F. Allen. A Plan-Based Approach to Speech Act Recognition. Technical report, University of Toronto, 1979.
- [AP80] James F. Allen and C. Raymond Perrault. Analyzing Intention in Utterances. *Artificial Intelligence*, 15:143–178, 1980.

- [BL87] Penelope Brown and Steve Levinson. *Politeness: Some Universals in Language Usage*. Cambridge University Press, 1987.
- [BL93] J. Eugene Ball and Daniel T. Ling. Natural Language Processing for a Conversational Assistant. Technical Report MSR-TR-93-13, Microsoft, 1993.
- [Cah90] Janet E. Cahn. The Generation of Affect in Synthesized Speech. *Journal of the American Voice I/O Society*, 8:1–19, July 1990.
- [Coh78] Philip R. Cohen. On Knowing What to Say: Planning Speech Acts. Technical Report 118, University of Toronto, Department of Computer Science, 1978.
- [CPB⁺94] J. Cassell, C. Pelechaud, N. Badler, M. Steedman, B. Achorn, T. Becket, B. Douville, S. Prevost, and M. Stone. ANIMATED CONVERSATION: Rule-Based Generation of Facial Expression Gesture and Spoken Intonation for Multiple Conversational Agents. In *Proceedings*, pages 413–420. SIG-Graph, 1994.
- [Dal88] R. Dale. *Generating Referring Expressions in a Domain of Objects and Processes*. PhD thesis, University of Edinburgh, 1988.
- [Elh92] Michael Elhadad. *Using Argumentation to Control Lexical Choice: a Functional Unification Implementation*. PhD thesis, Columbia University, 1992.
- [GL71] D. Gordon and G. Lakoff. Conversational Postulates. In *Papers from the 7th Regional Meeting of the Chicago Linguistic Society*, pages 63–84. CLS, 1971.
- [Gof83] Erving Goffman. *Forms of Talk*. University of Pennsylvania Press, Philadelphia, Pennsylvania, U.S.A., 1983.
- [Hir85] Julia Hirschberg. *A Theory of Scalar Implicature*. PhD thesis, University of Pennsylvania, 1985.
- [Hov93] Eduard H. Hovy. Automated Discourse Generation Using Discourse Structure Relations. *Artificial Intelligence Journal*, 63:341–385, 1993.
- [HRB94] Barbara Hayes-Roth and Lee Brownston. Multiagent Collaboration in Directed Improvisation. Technical Report KSL 94-69, Knowledge Systems Laboratory, Stanford University, 1994.
- [HRBS95] Barbara Hayes-Roth, Lee Brownston, and Erik Sincoff. Directed Improvisation by Computer Characters. Technical Report KSL 95-04, Knowledge Systems Laboratory, Stanford University, 1995.
- [Kam95] Candace Kamm. User Interfaces for Voice Applications. In David Roe and Jay Wilpon, editors, *Voice Communication between Humans and Machines*, pages 422–442. National Academy Press, 1995.
- [LA90] Diane Litman and James Allen. Recognizing and Relating Discourse Intentions and Task-Oriented Plans. In Philip R. Cohen and Jerry Morgan and Martha E. Pollack, editor, *Intentions in Communication*. MIT Press, 1990.
- [LB95] A. Bryan Loyall and Joseph Bates. Behavior-Based Language Generation for Believable Agents. Technical Report CMU-CS-95-139, School of Computer Science, Carnegie Mellon University, 1995.
- [Lit85] Diane Litman. Plan Recognition and Discourse Analysis: An Integrated Approach For Understanding Dialogues. Technical Report 170, University of Rochester, 1985.
- [MDBP94] Pattie Maes, Trevor Darrell, Bruce Blumberg, and Sandy Pentland. Interacting with Animated Autonomous Agents. In *AAAI Spring Symposium on Believable Agents*. AAAI, 1994.
- [MP93] Johanna D. Moore and Cécile L. Paris. Planning Text for Advisory Dialogues: Capturing Intentional and Rhetorical Information. *Computational Linguistics*, 19(4), 1993.
- [NST95] C. Nass, J. Steuer, and E. Tauber. Computers are Social Actors. In *Conference on Computer Human Interaction*, pages 72–78. CHI, 1995.
- [PFB80] Ellen Prince, Joel Frader, and Charles Bosk. On Hedging in Physician-Physician Discourse. In *AAAL Symposium on Applied Linguistics in Medicine*. AAAL, 1980.
- [RWS⁺94] Charles Rich, Dick Waters, Carol Strohacker, Michal Roth, and Yves Schabes. Demonstration of an Interactive Multimedia Environment. *Computer*, 27(12):15–22, 1994.
- [Sea75] John R. Searle. Indirect Speech Acts. In *Syntax and Semantics III: Speech Acts*, pages 59–82. Academic Press, New York, 1975.
- [WABM95] Dick Waters, David Anderson, John Barbus, and Joe Marks. The SPLINE Scalable Platform for Interactive Environments. Technical report, Mitsubishi Electric Research Labs, 1995.
- [Wal96a] Marilyn A. Walker. The Effect of Resource Limits and Task Complexity on Collaborative Planning in Dialogue. *Artificial Intelligence Journal*, 85(1–2), 1996.
- [Wal96b] Marilyn A. Walker. The VIVA Virtual Theater. Technical Report (forthcoming), Mitsubishi Electric Research Labs, 1996.

[YLM95] Nicole Yankelovich, Gina-Anne Levow, and Matt Marx. Designing Speech Acts: Issues in Speech User Interfaces. In *Conference on Human Factors in Computing Systems*. CHI, 1995.