

Grapevine: A Gossip Generation System

Foadad Khosmood
University of California, Santa Cruz
1156 High Street
Santa Cruz, CA 95064
+1-831-460-1004
foaad@soe.ucsc.edu

Marilyn Walker
University of California, Santa Cruz
1156 High Street
Santa Cruz, CA 95064
+1-831-459-1058
maw@soe.ucsc.edu

ABSTRACT

Generating believable and contextual dialogue among non-player-characters (NPC) remains one of the major challenges in interactive entertainment. Dialogue scenes in virtual environments are crucial to narrative progression and user believability, yet they continue to demand heavy authorial burden. In this paper, we describe our project Grapevine, a system for generating gossip-style conversation. We model the gossip conversation with a series of speech-acts controlled by a dialogue manager. We model characters with traits derived from the Big Five theory of personality. Grapevine also maintains an independent belief matrix, allowing for modeling of phenomena such as dishonesty, misunderstanding and bias. The dialogue manager decisions are a function of both narrative progression and personality traits. Surface text realization is achieved using RealPro, an off-the-shelf realizer (Lavoie & Rambow, 1997) and stylistically enhanced with the PERSONAGE generator (Mairesse & Walker, 2007). We demonstrate the current performance of the system with sample output of a three character series of gossip dialogues and discuss results of our 50 person validation survey.

Categories and Subject Descriptors

Design, Human Factors

General Terms

Algorithms, Design, Human Factors

Keywords

dialogue, gossip, natural language generation, Grapevine

1. INTRODUCTION

Spontaneous conversations ordinarily contain many examples of gossip -- that is, references to, and social evaluations of, the personal qualities, behaviors, and affairs of a person who is not

present (Eder & Enke, 1991; Gottman & Mettetal, 1986; Parker Teasley, Meissner, & McClellan, 1994). The “evaluative” nature of gossip means that it is a primary means for expressing and discussing aspects of character and shared value systems. The “behind-the-back” aspect of gossip is one of its most distinctive features. Studies suggest that children find gossip with their friends highly entertaining and use gossip as a basis for building solidarity against others (Gottman, 1983; Gottman & Mettetal, 1986; Parker & Gottman, 1989; Rysman, 1977). Thus gossip functions as a form of self expression for the individual, and functions in relationships and groups to promote solidarity and intimacy, indicate social status, and establish norms for appropriate behavior (Eder & Enke, 1991; Fine & Rosnow, 1978; Gluckman, 1967; Suls, 1977; Abrahams, 1970; Parker & Seal, 1996). More speculatively, gossip has also been described as a useful evolutionary tool and a “hard wired” survival skill important for managing reputations in large distributed groups (McAndrew, 2008; Dunbar, 1998).

Given that some form of gossip is present in almost any social setting, there has been little work in computational modeling of gossip. Although rumors and related language generation models have been featured in digital games in the past, we are not aware of any previous research specifically on the generation of gossip in the context of interactive entertainments. There are interactive chat bots such as Gossip Galore (Cheng, 2009) that can converse with a human about gossip related topics using knowledge bases they have gathered from the web. But generating believable and contextual dialogue among non-player-characters (NPCs) for a real life phenomenon such as gossip is still a challenge in interactive entertainment.

Dialogue scenes in virtual environments can be crucial to narrative progression and user believability, and they often provide key elements of the back-story. Yet they continue to demand detailed handcrafting by authors, which is both difficult and expensive to create, and perhaps more importantly, to modify. In this paper, we describe our project Grapevine, a system for specifically generating gossip-style conversation of the kind one might find around the proverbial work place water cooler or break room. Figure 1 shows two examples of Grapevine conversations, generated using character specific scripts enhanced with PERSONAGE stylistic transformations that express personality linguistically.

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. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

FDG 2010, June 19-21, Monterey, CA, USA

Copyright 2010 ACM 978-1-60558-937-4/10/06... \$10.00

PAT: Greetings and salutations to you Jo!
 JO:
 PAT: Tell us Ms. Jo, what is up with you?
 JO: Mhmm... I mean, I am not rather bad.
 JO: Ok, I mean, what is rather happening?
 PAT: Actually, there hasn't been a thing mate, you know, alright?
 JO: Nice try! I mean, something must be mmhm....
 PAT: Yes. There is one thing, you know buddy!
 PAT: Oh gosh it's about Alex and Jacki! I should keep going pal?
 JO: OK.
 PAT: I hear.... Alex enjoys Jacki pal, you know!
 JO: Yes. Alex loves Jacki you sure?
 PAT: OK. It's your turn pal, you know!
 JO: Yes. Mhmm... I mean, there is one thing.
 JO: Ok, i see, it's about Kevin. I mean, I should say it?
 PAT: Ok.
 JO: Rumors say.... Mhmm... err... Kevin was fired.
 PAT: Yes. wow!

PAT: Welcome my friend. Hello Chris!
 CHRIS: Hellos to you too Pat!
 PAT: Do say Mister Chris, how goes it?
 CHRIS: I am darn fine buddy!
 CHRIS: Yeah, what is really going on?
 PAT: Actually, there hasn't been a thing, you know!
 CHRIS: Interesting. Actually, something must be buddy!
 PAT: Yes. Actually, there is one thing, you know buddy, you see?
 PAT: Oh gosh yeah, it's about Jess and the boss, you know, alright? I should say?
 CHRIS: Yea.
 PAT: I have heard that.... Jess is just having with the Boss an affair, you know!
 CHRIS: Yes. yes, that explains a lot...

Figure 1. Grapevine-generated conversations, with characteristic scripts enhanced with PERSONAGE.

Our approach builds on our prior work in affective generation which attempts to replicate some of the social functions of conversation rather than treat dialogue as simply a mechanism for information exchange. In this work, we argue that evaluative language is a primary locus for self-expression, a type of speech event which clearly manifests the speaker's personality (Mairesse & Walker, 2007; 2008). Here we investigate whether automatically produced variants based on personality parameterization can enhance the authenticity of gossip-style conversation. Note that the parameters implemented in Grapevine allow us to generate many hundreds of variations on a single conversation. For example, in the production of gossip, some of the key variations are *slant* as indicated by lexical choice, determination of agency in events, and the use of hedging and other pragmatic markers.

We have also been influenced by several existing story telling machines, particularly Wide Ruled (Skorupski et. al., 2007) which is capable of accepting arbitrary characters and states and making story telling decisions based on them. We envision one use for Grapevine as a more concentrated expert module with many assets that can be plugged into an authorial assistance system like Wide Ruled or be utilized by a higher level drama manager.

In order to model gossip with reasonable accuracy we have identified some elementary requirements: psychological character traits for the conversants, a belief system capable of recording character state of knowledge and perceived knowledge, a historical account of what has already transpired in various conversations, a surface text realization system to translate between speech act desires and utterances and finally, a dialogue manager to direct the conversation. These components are discussed in Section 2. In Section 3 we report on an initial user evaluation of the conversations that Grapevine produces. Finally in Section 4 we present conclusions and describe our future work.

2. MODELING GOSSIP

Our aim is to build a model of in Grapevine that will generate gossip-style dialogues. Several other concepts are instrumental to the concept of gossip and must be modeled as well. These include character, beliefs, dialogue manager, surface text realization and personality based enhancements. In this section we will review these other concepts.

2.1 Architecture and Overview

Grapevine is built in Python using the Natural Language Tool Kit (NLTK). Stylistic enhancements are achieved using PERSONAGE (Mairesse & Walker, 2007) and RealPro (Lavoie, & Rambow, 1997) is used for surface text realization.

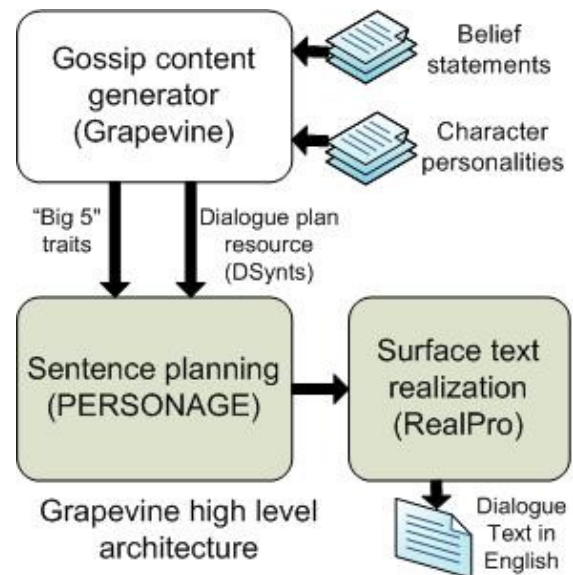


Figure 2. Grapevine high level architecture and interfaces

The dialogue manager is implemented as a state machine and controls the whole sequence of the gossip conversation. The conversation is modeled as a series of exchanges, each exchange consisting of two or more speech-acts, and each exchange functioning to complete one speaker/hearer gossip joint intention (Clark & Shaffer, 1989; Grosz & Sidner, 1990). We model characters with Big Five Personality traits and link these to their

propensity for gossip as well as the type of things they gossip about, and their reaction to gossip provided by others. Thus, the personality of a character affects its behavior at both the dialogue level and at the utterance level. Grapevine also maintains an independent belief matrix, allowing for modeling of phenomena such as dishonesty, misunderstanding and bias. Personality traits affect dialogue state transitions, sentence planning and surface text realization. The PERSONAGE language generator functions at the utterance level; it takes as input the personalities of the conversants, a content pool from the dialog manager represented in terms of simple syntactic structures, and the current belief states of the conversants. It is capable of producing many different variants of an utterance reflecting the personality of the speaker, and their propensity to, and attitude towards gossip.

2.2 Character

A large part of both the motivation and the act of engaging in gossip is tied to the subject's psychological traits and character. The "Big Five" theory of personality traits (Goldberg, 1993) provides a number of supertraits/subtraits that we hypothesize are relevant for gossip, and which we have operationalized in Grapevine. We focus on subtraits of the "Big Five" traits of Extraversion, Agreeableness and Conscientiousness because these traits have been found to be most easily manifested in language behavior (Mairesse & Walker, 2008). Table 1 specifies some of the subtraits at the high and low ends of the Big Five traits of Extraversion, Agreeableness and Conscientiousness.

Table 1. Example adjectives associated with "Big Five" traits

trait	high	low
Extraversion	warm, gregarious, assertive, talkative, optimistic, spontaneous, active, sociable	shy, quiet, reserved, passive, solitary, moody, joyless
Agreeableness	trustworthy, considerate, friendly, generous, helpful, altruistic	unfriendly, selfish, suspicious, uncooperative, malicious
Conscientiousness	competent, disciplined, dutiful, achievement striving, deliberate, careful, orderly	narrow-minded, conservative, ignorant, simple

Using the example adjectives and applying them to our gossip scenario, we implement six subtraits in Grapevine.

Each subtrait represents an axis between two polar opposite adjectives like "talkative" and "quiet" (see Extraversion high and low adjectives in Table 1). Subtraits are implemented using standard scales of 1..7 as used in research on personality psychology; we focus on the extremes and a neutral value for each trait in this study. The neutral designation gives us several advantages. First, it allows for partial specification of a character's subtraits. Thus, we can, if desired concentrate just on a few psychological traits. This corresponds to stylization of literary, cinematic and computer generated characters, since what usually makes a character memorable to the audience are only a few observable or dominant traits, even though the character may possess many more. Table 2 outlines the subtraits used in

Grapevine, associations with the "Big Five" traits and the effects of these parameters on system processes of state transitions and surface text generation.

Table 2. Grapevine Character Attributes

Subtraits	Big-5 Trait	Effects
talkative_quiet	Extraversion	state transitions surface text
gossip_reserved	Conscientiousness, Agreeableness	state transitions
sensitive_ insensitive *	Agreeableness	surface text
liar_honest	Conscientiousness	state transitions
trusting_skeptic	Agreeableness	state transitions surface text
formal_casual	Extraversion	surface text

* we consider "insensitive" synonymous with "rude"

2.3 Beliefs

We use the term "belief" to model the information known to each character. Such information is subject to re-examination or being gossiped about depending on the situation. For example one of the beliefs being discussed in the Grapevine generated exchange in Figure 1 is "Jess is having an affair with the Boss".

Grapevine supports beliefs of n "order" where n is 0, 1 or 2 and refers to the number of people involved in the belief. The aforementioned "Jess is having an affair with the Boss", for example, is a belief of second order while "The company is going bankrupt" involves no people and is of 0 order.

Beliefs are read by the system before any character selection or dialogue management. Each belief is specified by its event (like "having an affair"), order, participants ("Jess", "The Boss"), and a belief matrix. Table 3 describes inputs used to create the belief matrix for "Jess is having an affair with the Boss":

Table 3. Input for Grapevine belief creation

input parameter	Example	notes
Event	"having an affair"	can be anything
participants list	"Jess", "The Boss"	participants need not be Grapevine characters
knowers list	"Pat", "Jo"	list of Grapevine characters that know the event to be true

The belief matrix is a two dimensional table of 0.0-1.0 bound floating point values that correspond to each character's knowledge (belief) status about the event, as well as beliefs about other character's knowledge about the same event. Table 4 is an illustration of a belief matrix for the event "Jess is having an affair with the Boss". Pat, Chris and Jo are registered characters in Grapevine.

Table 4. A Grapevine belief matrix for “Jess is having an affair with the Boss”

	Pat	Jo	Chris
Pat	1.0	0.0	0.0
Jo	1.0	1.0	0.25
Chris	0.0	0.0	1.0

The intersection of a character with itself is the degree of belief held by the character about the event. The intersection of character A (column) and character B (row) contains the degree of certainty of knowledge that A has about B knowing about the event. The entries in the table should read, for example, as follows:

- Pat believes that Jess is having an affair with the Boss with 1.0 certainty.
- Jo believes that Pat knows Jess is having an affair with the Boss with 1.0 certainty.
- Jo believes that Chris knows Jess is having an affair with the Boss with 0.25 certainty.

Non-binary belief values such as the Jo/Chris intersection above allow for modeling of lying, dishonesty, skepticism and confronting contrary information in general. At present Grapevine simply considers any value above 0.9 as “true” and any value below “0.1” as false. Probabilistic truth functions will be considered for future work.

Under certain circumstances when a character comes into contact with contradictory information, an algorithm is used to reconcile previous knowledge with the new form. The result is a floating point belief value that Grapevine stores in belief matrices.

2.4 Dialogue manager

The Dialogue manager (DM) is responsible for traversing through the dialogue states, generating and storing utterances, as well as making belief modifications. The typical gossip conversation is modeled into two phases: Greet phase and reveal phase.

Every state except three special states (GreetPrep, EndGreet and End) also corresponds to a particular dialogue act and generates an utterance by one character. GreetPrep also produces text, but this text is just a brief announcement of the characters being present in the dialogue that follows, rather than a dialogic contribution. The states are further explained in Table 5. At almost every dialogue state, three functions are performed by the dialogue strategy active at that state:

1. Realize the utterance based on character personality traits and their utterance’s semantic representation (logical form), and store that utterance in the *history* object. We use the off the shelf surface realizer, RealPro (Lavoie & Rambow, 1997) to achieve the final surface text realization.
2. Update belief matrices, if any, based on any information transactions that might have occurred in the state. (Example: character b receives the gossip message and will now update the corresponding belief matrix based on its level of skepticism.)
3. Determine the next state based on character traits and pre-defined state transitions.

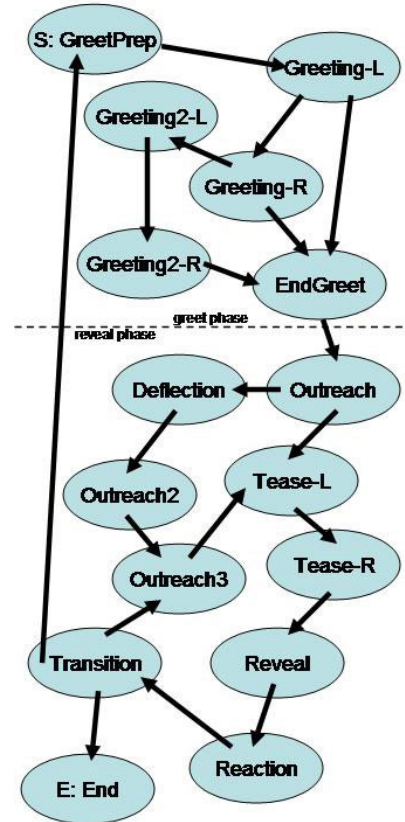


Figure 3. Grapevine’s dialogue state machine

Three special states serve different functions focused on dialogue management, rather than language generation. The InitialSetting state ensures that a valid gossip interaction can occur between the first two characters on the character list. By default the dialogue always occurs between the first two characters on the Grapevine character list, but other character may be substituted in after each transition. Grapevine iterates through a list of user-specified beliefs and evaluates each against the two characters for suitability of use for the gossip exchange. If no valid gossip session is possible, then the dialogue manager attempts to substitute another character into the top spot and try again. If all options have been exhausted, Grapevine terminates.

When an appropriate belief and a suitable pair of characters are selected, the next task by the DM is to determine which character should initiate the greeting. Selecting the initiator of the greeting phase is independent of the Revealing phase of the dialogue. In order to portray the best greeting experience, the DM selects the initial speaker based on the subtraits talkative_quiet and sensitive_insensitive.

The second special state, EndGreet evaluates the two chosen characters as in the InitialSetting state, except that instead of selecting the best greeting, it selects for the best gossip experience by examining all the beliefs known by either of the two characters and determining which character should be the information provider and which character should be the information receiver. The DM tries to find beliefs known by one character and

unknown to the other with greatest distance in knowledge gap possible. Character positions (a) and (b) are re-assigned accordingly. The character soliciting the information, (a), will begin the revealing phase in the Outreach state.

Table 5. Grapevine states

State	Description	Utterance logical form
INITIAL	Select 2 suitable npc's	setting(a,b)
GREET-L	Initial hello by npc "a"	greetl(a,b)
GREET-R	Reply greeting	greetr(b,a)
GREET2-L	2 nd . Greeting line such as "how are you?"	greet2l(a,b)
GREET2-R	Reply to secondary greeting	Greet2r(b,a)
ENDGREET	Re-evaluates and possibly re-assigns A,B npc's based on who has better secret = B	-- no utterance
OUTREACH	General attempt at fishing for information	outreach(a,b)
DEFLECTION	Initial deflection of attempt	deflection(b,a)
OUTREACH2	Second, more emphasized attempt to get some info	outreach2(a,b)
OUTREACH3	B starts accommodating	outreach3(b,a)
TEASE-L	B reveals some of the secret, asks for confirmation	teasel(b,a)
TEASE-R	A confirms, info requested	teaser(a,b)
REVEAL	B reveals full secret	reveal(b,a)
REACTION	A reacts according to personality and beliefs	reaction(a,b)
TRANSITION	DM determines, based on hand-crafted probability distribution, to: 1. continue conversation with roles reversed, or 2. swap in a third character and start from beginning, or 3. determine no more suitable gossip exists and move to END	-- no utterance
END	History is printed, program ends	histDump()

Lastly, the End state serves as the conclusion of the dialogue which may involve any number of two-person gossip chats. At this stage, the accumulated chat history is printed to the screen for the user. For PERSONAGE stylistic enhancing, output is in DSyntS XML and involves one more step before final realization.

Figure 4 specifies the dialogue act and the associated utterances for a dialogue generated by Grapevine.

```
[0][INITIAL SETTING] => Pat and Jo are in the room.
[1][GREET-L] => Pat: Welcome my friend. Hello Jo!
[2][GREET-R] => Jo: ...
[3][GREET2-L] => Pat: So enlighten me , Jo, how goes it?
[4][GREET2-R] => Jo: I'm great
[5][OUTREACH] => Jo: Anything happen?
```

```
[6][DEFLECTION] => Pat: Oh not much, not much
[7][OUTREACH2] => Jo: Oh PALLEASE! are you sure?
[8][OUTREACH3] => Pat: I heard something, actually...
[9][TEASE-L] => Pat: This involves Alex and Jacki. ready?
[10][TEASE-R] => Jo: Please go on.
[11][REVEAL] => Pat: I found out Alex loves Jacki.
[12][REACTION] => Jo: You don't say!
[13][TRANSITION] => Pat: How about you now?
[14][OUTREACH3] => Jo: Ok, fine!
[15][TEASEL] => Jo: This concerns Kevin. Interested?
[16][TEASER] => Pat: Don't stop!
[17][REVEAL] => Jo: I was told Kevin has been fired.
[18][REACTION] => Pat: Thanks for the info
```

Figure 4. State-Utterance Correspondance for a gossip session

We can see in the first few lines that Pat is a lot more talkative than Jo. He uses more words than he has to, and he repeats the same speech acts, for example: "not much not much". This is the reason Grapevine chose to start the greeting phase with Pat, rather than Jo. Jo is dismissive of the elaborate greetings. This is a manifestation of subtraits "talkative", "sensitive" and "formal" for Pat and "quiet", "rude" and "casual" for Jo.

At the EndGreet state, Grapevine has performed a role reversal and now Jo begins to solicit gossip from Pat. In response to Pat's initial deflection, Jo is again acting impatient, rude and insensitive, demanding an answer "Oh PALLEASE!"

After the revelation of the first event, "Alex loves Jacki", Grapevine's DM selects the role reversal action and Pat asks Jo to return the favor and share another piece of gossip. This is the "Transition state" to "Outreach3" arrow in Figure 3 above.

2.5 Generator

Each state plans one or two sentences based on the function of the state and the character subtraits. These sentences are in the form of a Deep Syntactic Structure (DSyntS) tree. The DSyntS form is based on the Meaning-Text Theory (MTT) of Igor Mel'cuk and colleagues and serves as the input for the sentence realizer RealPro (Lavoie & Rambow 1997). A number of factors make DSyntS an ideal representation for our work. First, unlike phrase structure representations, DSyntS are dependency representations. There are no non-terminal nodes such as VP. DSyntS allow non grammatical representations. The "deep syntactic" part of the DSyntS means that only meaning bearing lexemes are represented and not function words. These factors allow for a cleaner separation of meaning structure and style and therefore cleaner manipulation of style which we depend on in Grapevine.

The Grapevine DM produces a DSyntS for every sentence generated. This process already generates variations based on personality and narrative progression as well as probabilistic phrase selection, all of which is encoded in the DM. The combined output functions as a dynamic text plan for the surface text realization process. Figure 5 shows a DSyntS representation of the sentence "Jess is having an affair with the Boss." These DSyntS are encoded in XML in Grapevine.

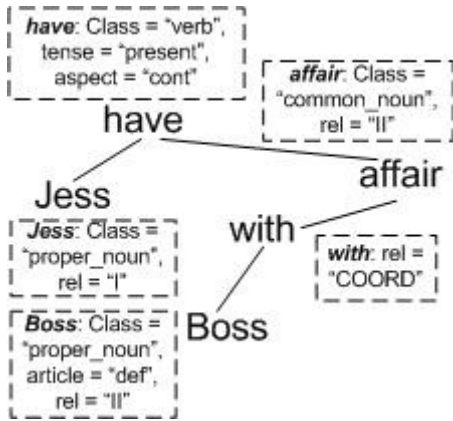


Figure 5. DSyntS representation of “Jess is having an affair with the Boss”

2.6 Personality-based enhancements

We say that personality manifests through language both in terms of *what* is being said, and *how* it is being said. While the Grapevine DM produces some personality based variations by the virtue of its content selection role (the “what”), starting from Grapevine 2.0, further personality based enhancements are achieved at the utterance level (the “how”). For example, Grapevine DM reflects that a quiet/shy character would not take secondary greeting branches in Figure 3. But the utterances that are ultimately generated for such a character need to themselves reflect the personality. As of the 2.0 version, we employ PERSONAGE to generate such personality enhancements resulting in even more variations in the dialogue.

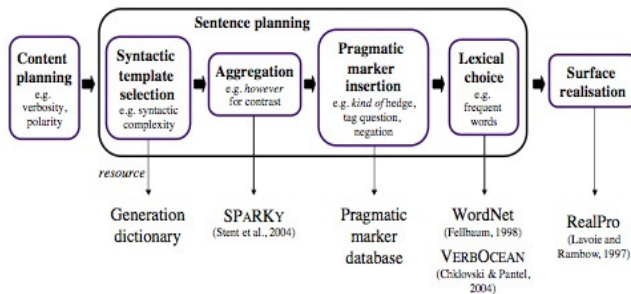


Figure 6. The architecture of the PERSONAGE generator.

PERSONAGE is a system capable of manipulating DSyntS to induce personality-based parameters before final surface realization. These enhancements are based on Big Five personality traits. Grapevine personalities are a set of subtraits of the Big Five, as shown in Table 2. The architecture of PERSONAGE is shown in Figure 6. The first component of the architecture is the content planner which specifies the structure of the information to be conveyed. The resulting content plan tree is then processed by the sentence planner, which selects syntactic templates for expressing individual propositions, and aggregates them to produce the utterance’s full syntactic structure. The pragmatic marker insertion component then modifies the syntactic structure locally to produce various pragmatic effects, depending on the markers’ insertion constraints. The lexical choice component selects the most appropriate lexeme for each content

word, given the lexical selection parameters. Finally, the RealPro surface realizer (Lavoie and Rambow 1997) converts the final syntactic structure into a string by applying surface grammatical rules, such as morphological inflection and function word insertion. Table 6 contains some examples of PERSONAGE enhancements made using personality information for the characters.

Table 6. Examples of original and enhanced utterances

Original	Personality-Enhanced
PAT: This is about Jess and the Boss. Should I continue?	PAT: Oh gosh yeah, it's about Jess and the Boss, you know, alright? I should continue?
JO: I'm OK.	JO: Mhm... i see, I am rather OK.
CHRIS: I'm OK.	CHRIS: Yeah, actually, I am darn OK, am I not?

3. EXPERIMENTAL EVALUATION

3.1 Survey design

To evaluate Grapevine, we conduct a 50 person survey on the campus of the University of California at Santa Cruz. Our aim is to evaluate two distinct aspects of the system. First, we want to know if the events being gossiped about are obvious, that is, can normal readers tell what the subject of the Gossip is? This point would indicate that Grapevine has succeeded in generating something that people associate with gossip and furthermore, people can identify the subjects of the gossip.

Second, we want to know if character traits exhibited in the dialogue are detectable by subjects. For example, can people tell that Pat is supposed to be talkative? This experiment was carried out with Grapevine 1.0.

Subjects are asked to judge a sample conversation from Grapevine and answer a survey. Grapevine currently produces two-person conversation sessions until there are no more valid beliefs eligible to be gossiped about. The resulting output is three to five conversations given our settings for the survey which typically fill about 1.5 pages. Two separate output runs are produced each involving the same exact set of events and characters, but the list of “knowers” (see Table 3) are changed to generate more variation in the output and the pairings. Four beliefs are inserted into the system (Table 7). The first three survey questions examine each character’s personality subtraits. The questions are:

How would you describe Pat/Jo/Chris ?

Four pairs of choices are provided as well as “none of the above”:

talkative, quiet, sensitive, rude, formal, casual, trusting, skeptic, none of the above

The subject is instructed to circle at least one of the choices. The fourth question asks the subject to write down the four gossip events mentioned in the text.

3.2 Results

Fifty subjects completed the survey correctly. We first analyzed the event question responses as shown in Table 7. Table 7 shows subjects can easily identify the gossip events being revealed in the dialogue.

Table 7. Gossip event identification results

gossip event	# correct	% correct
Alex loves Jackie	47	94%
Ingrid works as an exotic dancer	47	94%
Jess is having an affair with the Boss	49	98%
Kevin has been fired	49	98%

Next, we examine the subtraits survey questions. Each question instructs subjects to circle one or more subtraits per character. Each character asserted four subtraits in the conversation. Many subjects selected only one subtrait (presumably the most dominant) and others selected two or three. A random process chance of selecting each of the correct subtraits is 1 in 8 or 12%. We use this as the baseline to evaluate the per-character subtrait identification results.

Table 8. Subtrait identification results for character Pat

subtrait	# correct	% correct
formal	8	16%
sensitive	1	2%
talkative	43	86%
trusting	12	24%

Table 9. Subtrait identification results for character Jo

subtrait	# correct	% correct
casual	21	42%
quiet	8	16%
rude / insensitive	9	18%
skeptic	13	26%

Table 10. Subtrait identification results for character Chris

subtrait	# correct	% correct
casual	22	44%
sensitive	1	2%
talkative	25	50%
trusting	5	10%

The results in Table 7 indicates that the Grapevine generated language was clear enough for subjects to know the content of the gossip. However, the results in Tables 8 through 10 on identifying character subtraits indicates there is possibly room for improvement in our generation engine, as well as in the way that we elicit subject’s judgments. In future we will explore using the TIPI scale for personality perception judgments (Mairesse & Walker, 2008). While the percentages of correct answers are in general much lower, they are still significant. All but three

instances, or 9 out of 12 trait questions, are above the 12% baseline discussed before.

What we find immediately in these numbers is a wide gap between some -- one might call “dominant” -- subtraits, compared to others in the same character. Pat’s dominant subtrait, for example is “talkative”. This was easily identified by the subjects, while hardly anyone identified one of his other traits, “sensitive”. Obviously there is linguistic indeterminism in communicating these personality attributes both at the generation and at the understanding levels. Sensitive utterances might be mistaken for talkative ones. There might be dependency between the two subtraits. It might be the case that “sensitive speech” is objectively more verbose.

The words themselves are ambiguous and we made no effort to define them for the subjects. The word “sensitive” was probably the clearest case of ambiguity. We implemented sensitive, as a positive quality meaning “agreeable” or sympathetic, while we believe many people thought of it as “touchy” or “jumpy” which was the opposite of our intention. Admittedly, a subtrait like “talkative” is easier to identify with visual clues than for example “trusting”.

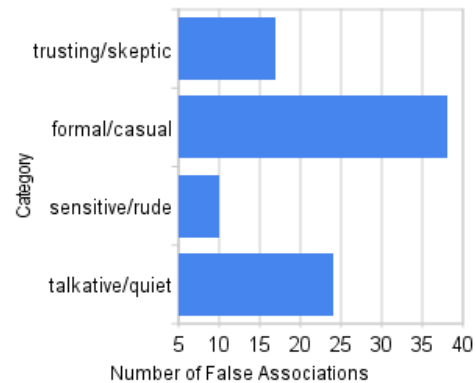


Figure 7. False association statistics

We instruct the subjects to circle one or more subtraits, not giving a specific number and not even identifying those subtraits that we have implemented as opposites (like “sensitive” versus “rude”). This may have been the reason for a large number of surveys with only one or two selections per character.

Figure 7 displays the number of false associations. A false association occurs when the subject circles a subtrait that the character does not have. Each circle results in one false association. Thus several false associations are possible per survey. The results are categorized by subtrait pairs. By far the largest number of false associations is about formal/casual distinction. Specifically, many times a formal character is mistaken for a casual one. Of the 38 false associations in the formal/casual category 28 of them are Pat being labeled “casual” when he was formal. In fact only 8 people correctly identified Pat as formal meaning the “formalness” of Pat was not communicated through the dialogue language. The responses were more correlated with casual rather than formal. This was the only case where a significant number chose an opposite subtrait.

It is rather difficult to interpret this latter phenomenon. One can point to the fact that “formality” is a relative term and no clear bar was given to the subjects by which to evaluate it. Across the

board, there was clear tendency not to associate any character with “formal.” Another possibility was discussed before. The dominant trait for Pat is talkative and the signs pointing to the talkative subtrait may have overwhelmed those hinting at a “formal” subtrait. At the end, subjects may have concluded that whoever is chatty is probably not formal.

4. CONCLUSIONS AND FUTURE WORK

We present Grapevine as a state-driven gossip generation tool making use of PERSONAGE personality-based style enhancements and using RealPro as the surface text realizer. Our intention is to study gossip dialogue generation and provide the tool as plug-in to larger planning-based drama managers in interactive entertainment systems or authorial assistance tools such as Wide Ruled.

We consider the biggest area for improvement the sentence planning architecture. We would like to further integrate available databases such as WordNet or ConceptNet for the purpose of increasing the variability of our sentences and sentence structures. We also would like to work on developing a clean API, perhaps an online resource that can be easily incorporated into larger systems and provide a gossip dialogue generation module that can be easily integrated into interactive story systems.

Future work will include more user surveys to validate newer features of Grapevine. Although the facilities exist, we did not implement any belief matrix functionality dealing with lies or dishonesty which we would like to experiment with in the future. Lastly, the Gossip model itself is being enhanced to contain a third phase whereby conversants continue discussion past the revealing phase and mention other facts related to the gossip subject leading perhaps to a probabilistic speculation model to be activated for some personalities.

5. REFERENCES

- [1] Abrahams, R. D. 1970. “A performance-centered approach to gossip.” *Man*, 5, 290–301.
- [2] Cheng, Xiwen et. al. (2009, April). “Gossip Galore”, Proceedings of the EACL 2009 Demonstrations Session, pp. 13-16, Athens, Greece.
- [3] Clark, H. H., and Schaefer, E. F. 1989. “Contributing to discourse.” *Cognitive Science*, 13, 259-294.
- [4] Dunbar, Robin. 1998. *Grooming, Gossip, and the Evolution of Language*. Harvard University Press. ISBN 0674363361.
- [5] Eder, D., & Enke, J. L. 1991. “The structure of gossip: Opportunities and constraints on collective expression among adolescents”. *American Sociological Review*, 56, 494–508.
- [6] Gluckman, M. 1967. “Psychological, sociological, and anthropological explanations of witchcraft and gossip: A clarification”. *Man*, 2, 20–34.
- [7] Goldberg, L. R. 1993. “The structure of phenotypic personality traits”. *American Psychologist*, 48, 26-34.
- [8] Gottman, John and Mettetal, Gwendolyn. 1986. “Speculations About Social and Affective Development: Friendship and Acquaintanceship through Adolescence.” Pp 192-237 in *Conversations of Friends: Speculations on Affective Development*, edited by J. Gottman and J. Parker. New York: Cambridge University Press.
- [9] Grosz, B. J. and Sidner, C. L. 1990. “Plans for discourse”. In Cohen, P. R., Morgan, J., and Pollack, M. E., editors, *Intentions in Communication*, pages 417-444. The MIT Press: Cambridge, MA.
- [10] Leaper, Campbell and Heathre Holliday. 1995. “Gossip in same-gender and cross-gender friends’ conversations”, *Personal Relations*, 2 (1995), pp. 237-246., Cambridge University Press.
- [11] Lavoie, B., and Rambow, O. RealPro – a fast, portable sentence realizer. 1997. In *Proceedings of the Conference on Applied Natural Language Processing (ANLP’97)*, Washington DC
- [12] Mairesse, François and Marilyn A. Walker. 2007. PERSONAGE: Personality generation for dialogue. In Proc. of the 45th Annual Meeting of the Association for Computational Linguistics (ACL), 496–503.
- [13] Mairesse, François, Walker, Marilyn, Mehl, Matthias and Moore, Roger. 2007. “Using Linguistic Cues for the Automatic Recognition of Personality in Conversation and Text”. *Journal of Artificial Intelligence Research (JAIR)*, 30, pages 457-500.
- [14] Mairesse, François and Walker, Marilyn. (2008, March). “A Personality-based Framework for Utterance Generation in Dialogue Applications”. In *Proceedings of the AAAI Spring Symposium on Emotion, Personality, and Social Behavior*, Palo Alto
- [15] Mel’cuk, Igor. 1988. *Dependency Syntax: Theory and Practice*, State University of New York Press.
- [16] McAndrew, Frank T. “The Science of Gossip: Why We Can’t Stop Ourselves”, *Scientific American*, 2008.
- [17] Parker, J.G., Teasley, S.D., Meissner, R.M., & McClellan, T.L. (1994, June). “The social construction of preadolescent gossip: Content, gender, and interpersonal constraints”, annual symposium of the Jean Piaget Society, Chicago.
- [18] Parker, Jeffrey G., Gottman, John M. 1989. “Social and emotional development in a relational context: Friendship interaction from early childhood to adolescence.” in *Peer relationships in child development*. (pp. 95-131). Oxford, England: John Wiley & Sons. xiv, 434 pp.
- [19] Rysman, A. How the “gossip” became a woman. 1977 *Journal of Communication*, 26, 176-180.
- [20] Skorupski, J, Jayapalan, L., Marquez, S. and Mateas, M. 2007 “Wide Ruled: A friendly interface to author-goal based story generation”, *Lecture Notes in Computer Sciences*, 4871:26.
- [21] Suls, J. M. 1977. Gossip as social comparison. *Journal of Communication*, 27, pp. 164–168