Archetype-Driven Character Dialogue Generation for Interactive Narrative

Jonathan P. Rowe, Eun Young Ha, James C. Lester

Department of Computer Science
North Carolina State University
Raleigh, NC 27695, USA
{jprowe, eha, lester}@ncsu.edu

Abstract. Recent years have seen a growing interest in creating virtual agents to populate the cast of characters for interactive narrative. A key challenge posed by interactive characters for narrative environments is devising expressive dialogue generators. To be effective, character dialogue generators must be able to simultaneously take into account multiple sources of information that bear on dialogue, including character attributes, plot development, and communicative goals. Building on the narrative theory of character archetypes, we propose an archetype-driven character dialogue generator that uses a probabilistic unification framework to generate dialogue motivated by character personality and narrative history to achieve communicative goals. The generator's behavior is illustrated with character dialogue generation in a narrative-centered learning environment, CRYSTAL ISLAND.

Keywords: Agents in narrative

1 Introduction

Devising robust, believable virtual agents is a central problem in interactive narrative. Because characters are instrumental in defining and advancing plots, as well as in creating compelling experiences for audiences [1], creating virtual agents that play the roles of these characters is a central issue in interactive narrative generation. Recent years have seen great strides in virtual agents for interactive narrative in education [2], [3], training [4], [5], [6], and entertainment [7], [8], [9]. This work has largely centered on creating adaptive, narrative-centered interactions that afford significant degrees of user control and autonomy, with the promise of coherent and engaging experiences that satisfy the myriad criteria of narrative utility [10], [11].

Character dialogue is often the driving force in moving a plot forward. Characters drive narrative by fostering empathy and conflict within a story. Life-like characters are revealed to the audience through a combination of dialogue and events within the story world. To support believable character-character and character-player interactions, a computational model of character dialogue generation for interactive narrative must satisfy three requirements. First, it must generate character-appropriate dialogue. Because each character's dialogue must "follow clearly and

validly from the character that uses it" [1], dialogue must be appropriate for character personalities and their associated motivations, preferences, and constraints. Second, it must consider the narrative context in which the dialogue will be delivered. As it generates dialogue, it must take into account narrative history to respond to the possible interactions initiated by the user and other characters. Third, it must perform robustly: it must be able to cope with combinations of goals, character personality attributes, and narrative history that cannot be anticipated prior to runtime.

To address these requirements, we propose an archetype-driven character dialogue generator. The dialogue generator builds on *character archetypes* [12], which are narrative-theoretic blueprints of well-established sets of traits for a particular character and role, such as fears, goals, motivations, and personality characteristics. Character archetypes exist outside of any one particular narrative scenario; they are adopted because of their power to define consistent sets of character traits that are both familiar and believable to audiences. The archetype-driven model of character dialogue generation employs probabilistic unification grammars that enable it to simultaneously consider multiple sources of information (character archetypes, narrative history, and communicative goals) to dynamically generate character-appropriate dialogue that achieves specific communicative goals for specific plot contexts. The model of dialogue generation has been implemented in a character dialogue generator for characters inhabiting CRYSTAL ISLAND, a narrative-centered learning environment for the domain of middle school microbiology.

This paper is structured as follows. Section 2 discusses related work in virtual agents for narrative environments and dialogue generation, which is followed by an introduction to the CRYSTAL ISLAND narrative-centered learning environment in Section 3. Section 4 describes the archetype-driven model of character dialogue generation; it presents an interactive narrative architecture that houses the character dialogue generator, the archetype-based character representation, and the probabilistic unification formalism used by the character dialogue generator. The dialogue generator's behavior is illustrated with a scenario from CRYSTAL ISLAND in Section 5. Concluding remarks and directions for future work follow in Section 6.

2 Related Work

A significant body of work has investigated intelligent virtual characters for interactive narrative environments. Among the earliest and most influential work is that by Carnegie Mellon's Oz group [7], which explored a range of issues in creating believable virtual characters for interactive drama, including extensions to their behavior specification language to support natural language generation [13]. The interactive drama Façade proposed a beat-based structuring of narrative content and character behaviors to create dramatic, adaptive vignettes [9]. Façade made significant advancements in character believability and player control, although dialogue was not dynamically generated.

Beyond incorporating rich emotional models for virtual characters [2], [4], work on character representation has informed the approach to character dialogue generation introduced here. Rizzo et al. [14] propose a goal-based model of character personality

and social behavior for use in a virtual environment. One of the objectives of this approach was to ensure that "agents endowed with different personalities [could] perform the same high-level behaviors in different ways" [14]. The approach proposed here shares this objective. Mosher and Magerko [15] introduce character representations that use personality templates and "stereotypes," an approach paralleling the use of character archetypes here. The abstract character representations that they use are psychologically motivated and emphasize psychological traits whereas the proposed archetype representation has literary underpinnings and focuses on character preferences and goals.

Work on character dialogue generation in interactive narrative has explored a number of representational choices. Cavazza and Charles [8] propose a unified representation for narrative and communicative acts that takes advantage of a hierarchical task network (HTN) representation for narrative planning in the I-Storytelling environment. In contrast, THESPIAN uses a decision-theoretic approach to reasoning about character goals to control socially normative dialogue behavior [6]. THESPIAN's approach accommodates probabilistic preferences and weights. Other notable work on character dialogue includes T2D, which maps Rhetorical Structure Theory structures to DialogueNet structures to generate robust and extensible character-character dialogue from a monological text [16].

The natural language generation community has investigated generation techniques that consider the speaker's personality. Walker et al. [17] demonstrate that linguistic style is a significant influence on communication, and that listeners infer a speaker's personality from the speaker's linguistic style. Their Linguistic Style Improvisation (LSI) theory analyzes the speaker's choices of semantic content, syntactic form, and acoustic features according to inter-personal relationships. While LSI focuses on social dimensions of linguistic style, another line of research has explored personality-informed variations in computer-generated dialogue. Isard et al. [18] represent linguistic personality with n-gram language models. Their language models were trained on a corpus that is labeled with the dimensions of big-five personality models. Mairesse and Walker [19] map personality-dependent linguistic traits to their dialogue generator's parameters to control various aspects of dialogue generation, such as content selection and structure, syntactic template selection, aggregation, pragmatics, and lexical selection. While both lines of work represent significant advances, they do not consider interactive narrative factors bearing on character dialogue generation.

3   CRYSTAL ISLAND Narrative-Centered Learning Environment

The archetype-driven model of character dialogue generation is being investigated in CRYSTAL ISLAND (Figure 1), a narrative-centered learning environment for the domain of microbiology for middle school students [20]. CRYSTAL ISLAND features a science mystery set on a recently discovered volcanic island. Students play the role of the protagonist, Alex, who is attempting to discover the identity and source of an unidentified infectious disease plaguing a newly established research station. The story opens by introducing the student to the island and members of the research team

for which the protagonist's father serves as the lead scientist. Several of the team's members have fallen gravely ill, including Alex's father. Tensions have run high on the island, and one of the team members, Ford, suddenly accuses another, Quentin, of having poisoned the other researchers. It is the student's task to discover the outbreak's cause and source, and either incriminate or exonerate Quentin.

The virtual world of CRYSTAL ISLAND, the semi-autonomous characters that inhabit it, and the user interface were implemented with Valve Software's Source™ engine, the 3D game platform for Half-Life 2. The Source engine currently provides much of the low-level (reactive) character behavior control. Narrative generation, pedagogical planning, and character behavior management are the subject of ongoing work.
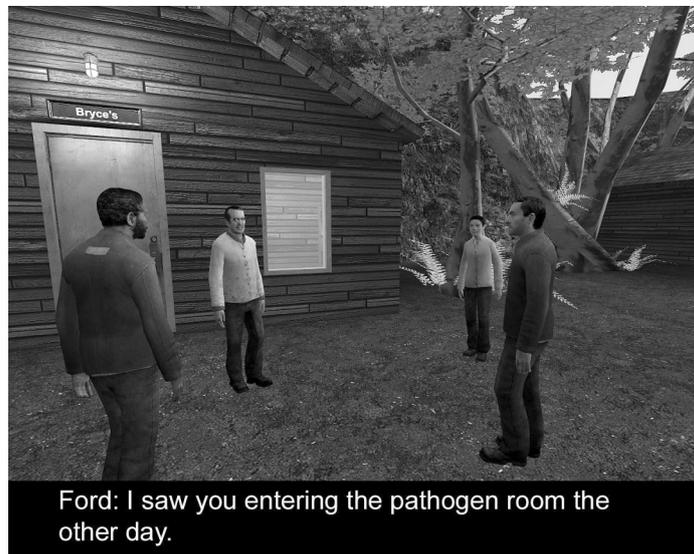

Fig. 1. CRYSTAL ISLAND narrative-centered learning environment

4 Archetype-Driven Character Dialogue Generation

Figure 2 shows an overview of the components of an interactive narrative architecture. Prior to run time, interactive narrative authors use the Character Composer's authoring facilities to instantiate characters from archetypes to populate the story world. The resultant character compositions furnish the raw materials for the Narrative Generator, which considers character attributes and the Narrative History as it dynamically constructs narrative plans. The Character Dialogue Generator is passed communicative goals by the Narrative Generator. Communicative goals consist of character objectives achievable by performing certain speech acts. To achieve communicative goals, the Character Dialogue Generator considers character

compositions and the Narrative History as it generates the dialogue. The Character Behavior Controller coordinates character dialogue and action in the world model.
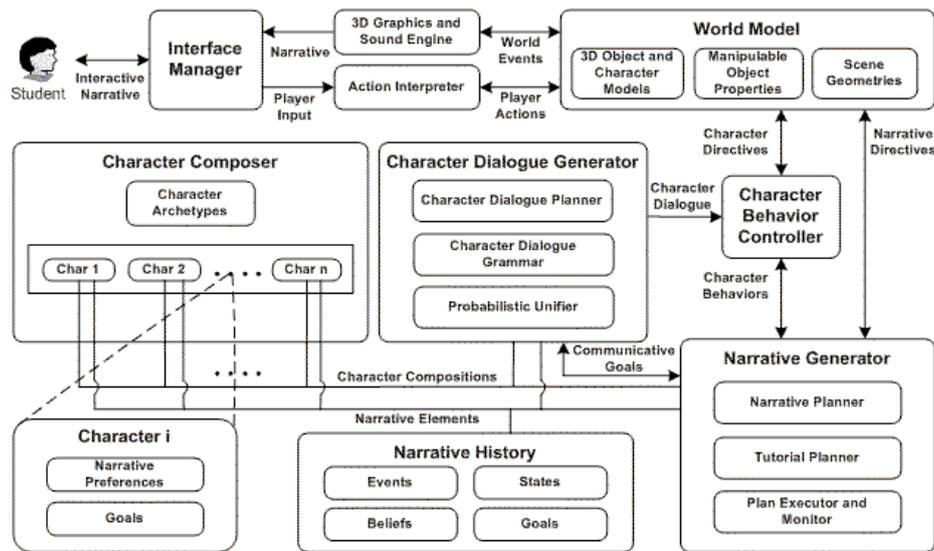

Fig. 2. Interactive narrative architecture

The four components that drive character dialogue are the Character Composer, the Narrative Generator, the Narrative History, and the Character Dialogue Generator:

- **Character Composer:** A semi-automated module that manages story-independent and story-dependent representations for defining virtual characters, the character composer includes a library of character archetypes that define abstract character traits and attributes. In essence, the library is a cast of canonical character descriptions. Authors interact with the Character Composer to instantiate abstract archetypes into fully defined, story-dependent characters that will inhabit the interactive narrative environment. During execution, the instantiated characters contribute to the dialogue generation processes. Character representations are discussed in detail below (Section 4.1).

- **Narrative Generator:** Chief among the functionalities of interactive narrative is the capacity to adapt plot direction and character responses to a player's actions in the environment. The Narrative Generator's execution and monitoring facilities serve as the interface between narrative planning and the world model with which players interact. The Plan Executor reconciles narrative plans and character definitions with the current narrative state to devise communicative goals and behaviors for the narrative's characters. These goals are emitted as behavioral directives to the Character Behavior Controller or as communicative goals to the Character Dialogue Generator. As these goals are pursued within the world model, the Narrative Generator monitors goal status in order to update the

Narrative History or invoke narrative re-planning. For narrative-centered learning environments, the narrative generator also includes a tutorial planner.

- **Narrative History:** The Narrative History maintains a set of both authored and generated narrative events, facts, character beliefs, and goals comprising the story-to-date. The initial state for the Narrative History encodes the narrative back-story predating any plot progress that a student has observed or participated in. As students progress through the narrative, additional elements are added to the Narrative History. Narrative elements inform (1) the Character Composer when it instantiates virtual characters from abstract archetypes, (2) the Narrative Planner by contributing to narrative plan formulation, and (3) the Character Dialogue Generator to guide the dialogue generation search process.

- **Character Dialogue Generator:** Given communicative goals by the narrative planner, the character dialogue generator's activities are informed by character archetypes and narrative history. The character dialogue generator is discussed in detail below (Section 4.2).

4.1 Archetype-Based Character Representation

To capture the regularities exhibited by commonly occurring categories of characters, character representations employ an archetype model. Archetypes are a powerful structuring tool used by authors to create believable but distinctive characters in their stories [12]. The prevalence of character archetypes can be observed throughout literature and other narrative media. For example, Shakespeare's Hamlet is a classic example of the *Recluse* archetype, whereas Hamlet's mother Gertrude exemplifies the *Matriarch* archetype [12]. For interactive narrative, constructing computational models of character archetypes offers the potential to create more portable, story-independent character specifications than traditional story-dependent approaches. Much as they do for authors, character archetypes offer interactive narrative attractive interactional properties such as character identifiability and believability.

To support the representations required by the Character Dialogue Generator's probabilistic unification grammar framework (as described in Section 4.2), character archetypes are defined using a weighted preference schema. Character archetypes are currently heuristically defined, and are composed of the following features:

- **Narrative Preferences:** For a given archetype, narrative preferences map normalized probabilities to subsets of the elements contained in the Narrative History. These mappings quantify the importance of narrative elements, which consist of narrative events, facts, and beliefs, to a particular character archetype. In this paper, events are occurrences that modify the narrative state in some manner, facts are true assertions about characters or the world state that do not explicitly reference events, and beliefs encompass events and facts whose actual truth values are not certain. Two primary types of narrative preferences are used here: categorical preferences and abstraction preferences. *Categorical preferences* distinguish among different types of narrative elements, quantifying whether a character would value concrete facts and events from the story, or de-emphasize them in favor of beliefs. *Abstraction preferences* are story-independent descriptions used to distinguish subsets of story-specific narrative elements that

are important to a particular archetype. Because archetypes are generally story-independent, it is necessary to define narrative abstractions that can label and encompass elements from a particular story. For example, characters that embody the *Abuser* archetype (Table 1) are often the target of some ego-damaging event, catalyzing them to seek revenge in response [12]. To encode this regularity, an interactive narrative "author" can create a narrative abstraction called *Revenge-Catalyst* prior to runtime. This abstraction can be specified to aggregate a subset of narrative elements once a story is authored and entered into the Narrative History. Because these elements are important to *Abuser* characters, its archetype specification will associate a high weight with the *Revenge-Catalyst* narrative abstraction and emphasize the associated elements during runtime.

- **Goals:** Character goals are explicit objectives pursued by a character, and are assigned normalized probabilities in a manner similar to Narrative Preferences. Character archetypes can also value or de-value goals independent of a story. For example, the *Traitor* archetype typically seeks to advance his or her social or professional status. This goal is independent of any story in which the *Traitor* participates. After characters are instantiated, their goals inform the Narrative Generator's narrative planning.

The approach described offers two principle benefits. First, the use of a probabilistic, preference-driven formalism provides a degree of flexibility in character's decision-making processes that supports robust performance in character dialogue generation. Second, the use of archetypes permits the partial definition of story-independent characters. These definitions can subsequently be used to instantiate story-specific characters that automatically adopt the features defined in an associated archetype. Archetypes introduce the opportunity to partially automate the character creation process once a sufficient library of character archetypes has been encoded. The current choice of features for encoding archetypes, and the generality of our character representations, are the subjects of ongoing investigation.

To illustrate, consider two character archetypes derived from [12], *Traitor*, and *Abuser* (Table 1). The *Traitor* is an archetype that typifies a cool, rational and organized antagonistic character whose priorities lie with their job or their own professional advancement. These characters often see business and life as a game to be won, and avoid chaotic or strongly emotional situations. Examples of this archetype from literature and film include *Wall Street*'s Gordon Gekko, and *A Christmas Carol*'s Ebenezer Scrooge. In its corresponding computational representation, one can observe this archetype's preference for objective, rational information through its high weighting of facts and events. Additionally, the choice of goals pertinent to the *Traitor* archetype emphasizes professional objectives such as the success of the research expedition (e.g. *See-Organization-Succeed*) and career advancement (e.g. *Advance-Status*), rather than personal or emotional concerns. The *Traitor* also values the *Incriminate-Evidence* abstract narrative element, weighting incriminating information that may eventually be useful for competitive advantage.

In contrast, the *Abuser* archetype typifies antagonistic characters that are strongly guided by their emotions, often exhibiting intensity, passion, and aggression. This archetype strongly values the ways that others view him, and it may carry a delicate ego. Examples of the *Abuser* include Shakespeare's Othello and *The Great Gatsby*'s Tom Buchanan. We encode this archetype by more strongly weighting beliefs over

events and facts to de-emphasize the rational, objective side of its personality. The archetype's tendency toward aggression and ego is characterized by its high preference for the *Embarassment-Sequence* abstract narrative element, a sequence referencing an ego-damaging event, as well as a heavy weight placed on the related *Avenge-Transgression* goal.

Table 1. Example partial definitions of the *Traitor* and *Abuser* character archetypes.

Archetype	*Traitor*		*Abuser*	
Narrative Preferences:				
Element Categories	Events	0.5	Events	0.3
	Facts	0.4	Facts	0.2
	Beliefs	0.1	Beliefs	0.5
Element Abstractions	Embarass-Sequence	0.2	Embarass-Sequence	0.5
	Conflict-Catalyst	0.3	Conflict-Catalyst	0.3
	Incriminate-Evidence	0.5	Incriminate-Evidence	0.2
Goals:	Advance-Status	0.4	Advance-Status	0.2
	See-Org-Succeed	0.3	See-Org-Succeed	0.1
	Bring-Order	0.3	Avenge-Transgression	0.7

4.2 Probabilistic Unification Grammars for Character Dialogue Generation

The principle objective of the Character Dialogue Generator is to create natural language utterances that satisfy a communicative goal provided by the narrative generator, that are appropriate for the specified character archetype, and that take the narrative history into account as the syntax and semantics of the dialogue are planned. Turn-taking conversational behavior is handled separately by the Narrative Generator, and is beyond the scope of this paper. The Character Dialogue Generator considers three sources of information: communicative goals from the Narrative Generator, archetype information from the Character Composer, and narrative elements from the Narrative History. First, the communicative goal specifies the speech act that should be performed at a given dialogue turn. Second, the character archetype provides generator preferences over the types of topics reflected by the archetype. Third, the narrative history provides a repository of current (and previous) topics.

In addition, the Character Dialogue Generator considers three additional sources of information and defines them as parts of the dialogue generation grammar. These are the preferences of communicative goals both over topics and over syntactic templates, and the preferences of character archetypes over syntactic templates. Syntactic characteristics of templates are represented using three features associated with template strings: *mood* (indicative, interrogative), *modality* (explicit, implicit), and *viewpoint* (1st singular, 1st plural, 3rd, other). Explicit vs. implicit modality determines whether the topic is explicitly expressed in the utterance or not. These features have been employed to express different linguistic traits representing either different types of characters [19] or different social relationships between characters [8]. Preferences over alternate values for each feature are represented using normalized probabilities, as in the case of topic preferences.
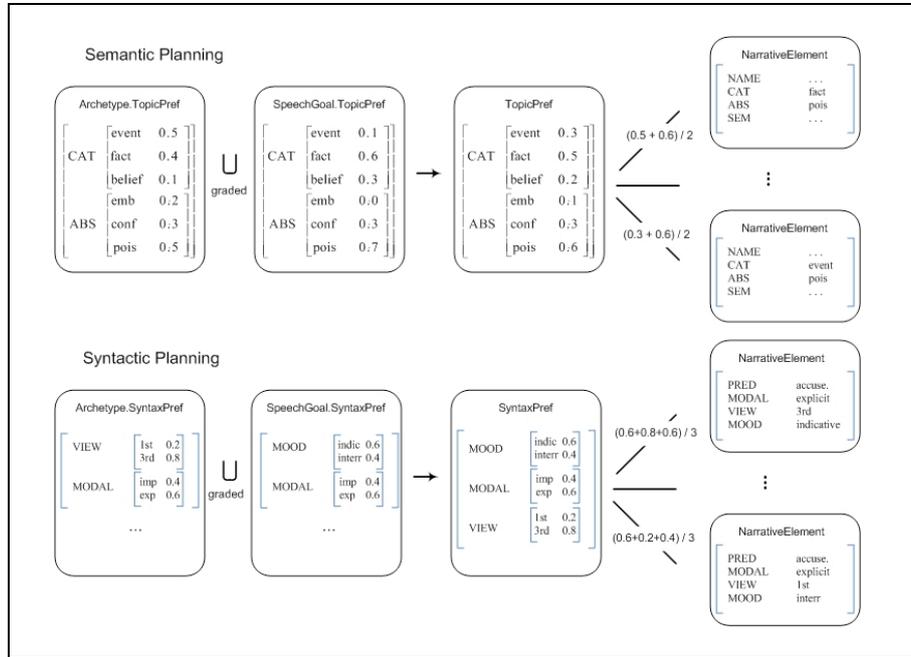
Fig. 3. Dialogue generation via probabilistic unification

Given these six sources of information, the Character Dialogue Generator accomplishes its objective by decomposing the generation task into two separate search processes, semantic planning and syntactic planning (Figure 3). *Semantic planning* is the process of searching for a topic (a narrative element) within the Narrative History that best satisfies the constraints imposed by the given preferences both of a character archetype and of a communicative goal. Similarly, *syntactic planning* is a search for a syntactic template to realize the selected topic that best satisfies the preferences of both the character archetype and of the communicative goal. Formulated as a search problem with constraints, the task of dialogue generation now confronts two challenges: (1) to effectively combine constraints from different sources, and (2) to robustly handle potential conflicts in the preferences between different sources.

To address these challenges, the character dialogue generator employs a probabilistic unification grammar framework. While probabilistic unification grammars have been applied to natural language processing tasks such as parsing [21], [22], [23] and speech recognition [24], the approach proposed here is among the first attempt to apply the framework to the task of dialogue generation. Probabilistic unification grammars offer three advantages over alternative approaches: they provide (1) a unified formalism to combine two different search tasks, (2) an effective method for merging different sources of information into a single structure, and (3) a principled mechanism for robustly handling potential conflicts in values between

different sources of information.

In probabilistic unification grammars, each terminal and non-terminal symbol is represented as a feature structure. A feature structure is a set of feature-value pairs in which values are either atomic symbols or nested feature structures (Figure 4). We adopt the technique of *graded unification* [21] that supports the unification of two features with different values by adjusting probabilities associated with each value. In this approach, each atomic-valued feature is associated with a probability and the probabilities of atomic values associated with a feature are normalized to sum to 1.0. When two atomic features are unified, the unification operator collects the disjoint set of all values present in the two arguments in the result and assigns a new probability to each value by computing the average of the probabilities associated with the value in the two arguments. When a value is present only in one argument, it is assumed that the probability of the value in the other argument is zero. It is through this technique that the Character Dialogue Generator exhibits robustness with respect to potential conflicts in the values from different information sources.

```
(a)                          (b)                                      (c)
⎡NAME:   elaborate_topic⎤    ⎡NAME:       businessman        ⎤        ⎡NAME:  ford_date_audrey⎤
⎢SPEAK:  Ford           ⎥    ⎢                  ⎡event:  0.5⎤⎥        ⎢CAT:   event           ⎥
⎣LISTEN: Quentin        ⎦    ⎢           CAT:  ⎢fact:   0.4⎥⎥        ⎢ABS:   emb_seq         ⎥
                             ⎢                  ⎣belief: 0.1⎦⎥        ⎢       ⎡PRED:   date  ⎤⎥
                             ⎢TOPICPREF:        ⎡emb_seq:  0.2⎤⎥       ⎢SEM:   ⎢TENSE:  past  ⎥⎥
                             ⎢           ABS:  ⎢conf_cat: 0.3⎥⎥       ⎢       ⎢ARG1:   Ford  ⎥⎥
                             ⎣                  ⎣pois_ev:  0.5⎦⎦       ⎣       ⎣ARG2:   Audrey⎦⎦
```
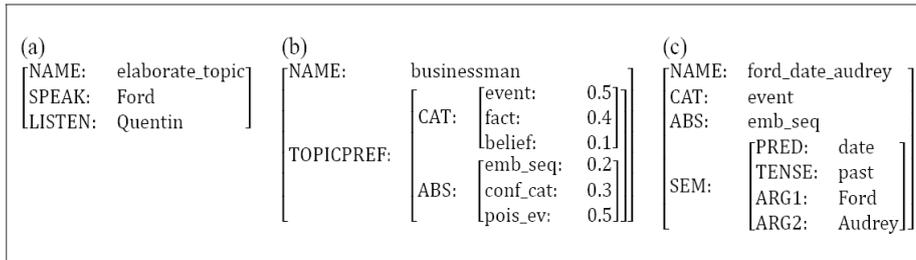
Fig. 4. Feature structure for (a) communicative goal, (b) character archetype, and (c) narrative element. In these examples, the probabilities of values equaling 1 are omitted for clarity.

Consider an illustrative example where the topic preferences of a character archetype (*Archetype.TopicPref*) and a communicative goal (*SpeechGoal.TopicPref*), are unified. Suppose that the feature *CAT* (topic category such as event, fact, and belief), common to both information sources, associates probabilities for *event* and *belief* of 0.6 and 0.4, respectively, in *Archetype.TopicPref*. The corresponding feature in *SpeechGoal.TopicPref* has only one value, *event,* with its associated probability 1.0.[1] The *graded unification* technique unifies these features into a new set with updated *event* and *belief* values. The probability of *event* is adjusted to 0.8, which is

---

[1] The current implementation of the Character Dialogue Generator uses information sources with manually authored probabilities. While the current approach demonstrates the feasibility of the probabilistic unification formalism for runtime performance, manually authoring probabilities would be prohibitively expensive. A promising direction for future work is learning probability distribution functions from observations of players' interactions with interactive narrative generators.

the average of its associated probability in each argument (0.6 in *Archetype.TopicPref* and 1.0 in *SpeechGoal.TopicPref*). The value *belief* is not present in *SpeechGoal.TopicPref*, so its associated probability in *SpeechGoal.TopicPref* is assumed to be 0.0 and its probability in the resulting set is adjusted to 0.2 (the average of 0.4 and 0.0).

Suppose that the Character Dialogue Generator is creating a dialogue for a character whose archetype is *Abuser* and a new communicative goal, *ElaborateTopic*, is posted. During semantic planning, the Dialogue Generator first computes the overall topic preference for the given dialogue turn by unifying the two features, *CAT* and *ABS* (topic abstraction such as Revenge-Catalyst, Embarrassment-Sequence), of the topic preference of the *Abuser* archetype and the topic preference of the communicative goal *ElaborateTopic* using the graded unification technique. The resulting topic preference contains new probabilities for each of the *CAT* and *ABS* features.

The combined topic preference is then used to search for a narrative element in the Narrative History by unifying the two features, *CAT* and *ABS*, of the combined topic preference with the corresponding *CAT* and *ABS* features of each narrative element in the Narrative History. Because more than one narrative element could be unified with the topic preference, the matching score of a narrative element is computed as the average probability of each feature in the topic preference that was unified with the features of the narrative element. The narrative element that yields the highest matching score is selected as the topic for the given dialogue turn.

Syntactic planning is performed in a similar manner. First, syntax preferences are computed by unifying the three features: *MOOD* (mood), *MODAL* (modality), and *VIEW* (viewpoint) of the syntax preference of the *Abuser* archetype and the syntax preference of the *ElaborateTopic* communicative goal. The combined syntax preference is matched with each syntactic template by unifying the three features. Among the templates matched with the combined syntax preference, the highest scoring one is selected as the template to realize the topic selected during semantic planning. Finally, the resulting dialogue utterance is generated by filling the slots within the selected template with the corresponding sub-features of the *SEM* (semantics) feature in the selected narrative element which represents the semantics of the corresponding narrative element.

5   An Illustrative Scenario

To illustrate the behavior of the archetype-driven approach to character dialogue generation, consider a scene in CRYSTAL ISLAND, where one of the characters, Ford, is accusing another character, Quentin, of poisoning other team members. The effects of archetype on a dialogue can be seen by holding constant the communicative goal and the narrative setting and then varying the character archetype. Table 2 shows sample dialogues created by the Character Dialogue Generator with two different settings of archetypes for Ford, the *Abuser* and the *Traitor*.

The generated dialogue shows how a character's archetype drives the generation decisions in selecting topics and syntactic templates. In the dialogue generated for the

ElaborateTopic communicative goal, the *Abuser* version of Ford selects his belief that Quentin contaminated the research team's food with salmonella to elaborate upon the current topic (accusing Quentin of poisoning) whereas the *Traitor* version selects an event in which he observed Quentin entering the pathogen room. This reflects the difference between the two archetypes: the *Abuser* makes passionate, unfounded accusations while the *Traitor* offers well-informed, controlled evidence to support his accusation. The results also show the difference in syntactic tendencies between the two archetypes. An *Abuser* character tends to be fiery and aggressive whereas the *Traitor* acts coolly and rationally. The *Abuser* selects templates with explicit modality for of all his dialogue turns, while the *Traitor* selects the explicit modality only twice (*ReassertTopic*, *ElaborateTopic*) out of four total turns.

Table 2. Generated dialogue

Speaker (Communicative Goal)	Character Archetype	
	Abuser	*Traitor*
Ford (*AssertTopic*)	We all know you are doing this, Quentin.	It is evident who's doing this, Quentin.
Quentin	What are you talking about, Ford?	
Ford (*ReassertTopic*)	Don't deny you poisoned Bryce and Teresa.	Did you think nobody could figure out you poisoned Bryce and Teresa?
Quentin	You are accusing me of making Bryce and Teresa sick?	
Ford (*ElaborateTopic*)	You put salmonella in Bryce's and Teresa's foods.	I saw you entering the pathogen room the other day.
Player	I don't believe you, Ford.	
Ford (*Rebut*)	Quentin poisoned Bryce, Alex.	Think about it rationally, Alex.

6 Conclusions and Future Work

Interactive narratives can benefit significantly from the implementation of robust character dialogue generators. Character dialogue is one of the primary means through which characters are revealed to an audience, but hand-authoring dialogues can become prohibitive as the space of possible dialogue interactions grows large. To address this problem, character dialogue generators for interactive narratives must meet several requirements. They must generate dialogues that are appropriate for characters' traits, such as personalities, motivations, and preferences; they must consider narrative context and history as they formulate dialogue; and they must be

able to robustly handle the large number of possible character-character and character-player interactions that may result in dialogue.

The model of dialogue generation introduced here uses probabilistic unification grammars to robustly create natural language character dialogue for interactive narrative. It is based on a computational representation of character archetypes that encodes story-independent character specifications. Character archetypes can be instantiated to compose characters whose sets of preferences over elements from the narrative history are utilized to generate character-appropriate dialogue that is situated within the current narrative context. The generated dialogues use preference information encoded within character archetype representations and yield character-specific variations in the dialogue.

The work to date has two principle limitations. First, it has focused on a few selected narrative contexts, an initial but not comprehensive archetype representation, and a small set of dialogues. Second, follow-up work needs to be undertaken to empirically investigate the fluidity and the "narrative appropriateness" of generated dialogues.

Several directions for future work appear promising. First, given the potentially rich inter-character relations that occur in compelling narratives, the range of information encoded by archetypes can be extended to represent character relationships and other features. The current feature set is preliminary in nature, and was largely chosen as a mechanism in service of the dialogue generation model discussed. Second, the current work focuses on natural language generation. In the future, it will be important to introduce character-appropriate prosodic markers in conjunction with text-to-speech technologies for speech generation. Third, it will be interesting to investigate techniques for coordinating archetype-driven character dialogue and characters' behaviors, with an emphasis on gesture and gaze. Finally, a particularly promising direction for future work is conducting extensive human subject evaluations. It is expected that these will yield important findings on the effects of character-driven dialogue on players, as well as insights into the set of features used to encode character archetypes and parameterize dialogues. Furthermore, we expect that human subject evaluations will furnish corpora for learning the probability distributions for the archetype representations and the unification grammars.

Acknowledgements

References

1. Egri, L.: The Art of Dramatic Writing: Its Basis in the Creative Interpretation of Human Motives. Simon & Schuster, New York (1960)
2. Aylett, R. and Figueiredo, R., Silva, A., Dias, J., A. Paiva.: Making It Up as You Go Along - Improvising Stories for Pedagogical Purposes. In: Proc. IVA. pp. 304–315 (2006)
3. Marsella S., Johnson, W. L., LaBore, C.: Interactive Pedagogical Drama for Health Interventions. In: Proc. of AI in Education (AIED). pp. 341-348 (2003)
4. Gratch, J., Marsella, S.: A Domain-independent framework for modeling emotion. Journal of Cognitive Systems Research 5, 269–306 (2004)
5. Riedl, M. O., Stern, A.: Believable Agents and Intelligent Story Adaptation for Interactive Storytelling. In: Proc. Tech. for Interactive Digital Storytelling and Entertainment. pp. 1 - 12 (2006)
6. Si, M., Marsella, S.C., Pynadath, D.V.: Thespian: Modeling Socially Normative Behavior in a Decision-Theoretic Framework. In: Proc. IVA. pp. 369-382 (2006)
7. Bates, J.: Virtual Reality, Art, and Entertainment. PRESENCE: Teleoperators and Virtual Environments 1, 133–138 (1992)
8. Cavazza, M., Charles, F.: Dialogue Generation in Character-based Interactive Storytelling. In: Proc. AI and Interactive Digital Entertainment (AIIDE). pp. 21–26 (2005)
9. Mateas, M., Stern, A.: Structuring Content in the Façade Interactive Drama Architecture. In: Proc. AI and Interactive Digital Entertainment (AIIDE). pp. 93–98 (2005)
10. Roberts, D., Isbell, C.: Desiderata for Managers of Interactive Experiences: A Survey of Recent Advances in Drama Management. In: Proc. AAMAS Wkshp on ABSHLE. (2007)
11. Rowe, J., McQuiggan, S., Lester, J.: Narrative Presence in Intelligent Learning Environments. In: Proc. AAAI Fall Sym. on Intell. Narrative Tech. pp. 126–133 (2007)
12. Schmidt, V. L.: 45 Master Characters: Mythic Models for Creating Original Characters. Writer's Digest Books: Cincinnati, OH (2001)
13. Loyall, A. B., Bates, J.: Personality-Rich Believable Agents That Use Language. In: Proc. Autonomous Agents, First International Conference. pp. 106-113(1997)
14. Rizzo, P., Veloso, M., Miceli, M., Cesta, A.: Goal-based personalities and social behaviors in believable agents. Applied Artificial Intelligence 13, 239-271 (1999)
15. Mosher, B., Magerko, B.: Personality Templates and Social Hierarchies Using Stereotypes. In: Proc. Tech. for Interactive Digital Storytelling and Entertainment. pp. 207–218 (2006)
16. Piwek, P., Hernault, H., Prendinger, H., Ishizuka, M.: T2D: Generating Dialogues Between Virtual Agents Automatically from Text. In: Proc. IVA. pp. 161–174 (2007)
17. Walker, M., Cahn, J., Whittaker, S. Improvising Linguistic Style: Social and Affective Bases for Agent Personality. In: Proc. Auton. Agents, 1st Int. Conf. pp. 96-105 (1997)
18. Isard, A., Brockmann, C., Oberlander, J. Individuality and Alignment in Generated Dialogues. In: Proc. Natural Lang Generation, 4th Int. Conference. pp. 25 - 32 (2006)
19. Mairesse, F., Walker, M. PERSONAGE: Personality Generation for Dialogue. In: Proc. 45th Annual Meeting of the Assoc. for Comp. Ling. (ACL). pp. 496-503 (2007)
20. Mott, B., Lester, J.: U-Director: A Decision-Theoretic Narrative Planning Architecture for Storytelling Environments. In: Proc. AAMAS. pp. 977–984 (2006)
21. Kim, A.: Graded Unification: A Framework for Interactive Processing. In: Proc. 32nd Annual Meeting of the Association for Computational Linguistics. pp. 313–315 (1994)
22. Abney, S.: Stochastic Attribute-Value Grammars. Comp. Ling. 23, 597–618 (1997)
23. Johnson, M.: Learning and parsing stochastic unification-based grammars. In: Proc. Learning Theory and Kernel Machines (COLT / Kernel). pp. 671-683 (2003)
24. Hemphil, C., Picone, J. Speech Recognition in a Unification Grammar Framework. In: Proc. ICASSP. pp. 723–726 (1989)