

Generating Different Story Tellings from Semantic Representations of Narrative

Elena Rishes¹, Stephanie M. Lukin¹, David K. Elson², and Marilyn A. Walker¹

¹ Natural Language and Dialogue Systems Lab
University of California Santa Cruz, Santa Cruz

{erishes,slukin,maw}@soe.ucsc.edu

² Columbia University, New York City
delson@cs.columbia.edu

Abstract. In order to tell stories in different voices for different audiences, interactive story systems require: (1) a semantic representation of story structure, and (2) the ability to automatically generate story and dialogue from this semantic representation using some form of Natural Language Generation (NLG). However, there has been limited research on methods for linking story structures to narrative descriptions of scenes and story events. In this paper we present an automatic method for converting from Scheherazade’s STORY INTENTION GRAPH, a semantic representation, to the input required by the PERSONAGE NLG engine. Using 36 Aesop Fables distributed in DramaBank, a collection of story encodings, we train translation rules on one story and then test these rules by generating text for the remaining 35. The results are measured in terms of the string similarity metrics Levenshtein Distance and BLEU score. The results show that we can generate the 35 stories with correct content: the test set stories on average are close to the output of the Scheherazade realizer, which was customized to this semantic representation. We provide some examples of story variations generated by PERSONAGE. In future work, we will experiment with measuring the quality of the same stories generated in different voices, and with techniques for making storytelling interactive.

Keywords: Semantic Narrative Representation, Natural Language Generation, StoryVariation.

1 Introduction

Sharing our experiences by storytelling is a fundamental and prevalent aspect of human social behavior. A critical aspect of storytelling “in the wild” is that it is socially interactive and situation dependent. Storytellers dynamically adjust their narratives to the context and their audience, telling and retelling the same story in many different ways depending on who the listener is. For example, storytellers tell richer stories to highly interactive and responsive addressees [1], and stories told by young adults “play” to the audience, repeatedly telling a story until a satisfactory peer response is found [2].

Original Story	Scheherazade	Personage
A hungry Fox saw some fine bunches of Grapes hanging from a vine that was trained along a high trellis, and did his best to reach them by jumping as high as he could into the air. But it was all in vain, for they were just out of reach: so he gave up trying, and walked away with an air of dignity and unconcern, remarking, "I thought those Grapes were ripe, but I see now they are quite sour."	Once, a group of grapes was hanging on a vine and the vine was hanging on a high trellis. A hungry fox saw the group of grapes. The fox jumped in order to obtain the group of grapes. The fox didn't obtain the group of grapes because he wasn't able to reach the group of grapes. The fox walked away from the group of grapes with dignity and with unconcern. The fox said that he had earlier thought that the group of grapes was ripe and said that he now saw it as being sour.	The group of grapes hung on the vine. The vine hung on the trellis. The fox saw the group of grapes. The fox jumped in order for the fox to obtain the group of grapes. The fox did not obtain the group of grapes because the fox was not able to reach the group of grapes. The fox walked away from the group of grapes with dignity and unconcern. The fox said the fox earlier thought the group of grapes was ripe. The fox said the fox now saw the group of grapes was sour.

Fig. 1. "The Fox and the Grapes" Aesop fable with Generated Versions

To have this human capability of telling stories in different voices for different audiences, interactive story systems require: (1) a semantic representation of story structure, and (2) the ability to automatically generate story and dialogue from this semantic representation using some form of Natural Language Generation (NLG). However to date, much of the research on interactive stories has focused on providing authoring tools based either on simple story trees or on underlying plan-based representations. In most cases these representations bottom out in hand-crafted descriptive text or hand-crafted dialogue, rather than connecting to an NLG engine.

Prior research on NLG for story generation has primarily focused on using planning mechanisms in order to automatically generate story event structure, with limited work on the problems involved with automatically mapping the semantic representations of a story and its event and dialogue structure to the syntactic structures that allow the story to be told in natural language [3–5]. Recent research focuses on generating story dialogue on a turn by turn basis and scaling up text planners to produce larger text prose [6–9], but has not addressed the problem of bridging between the semantic representation of story structure and the NLG engine [5, 7, 6, 10, 11]. An example of this work is the STORYBOOK system [3] which explicitly focused on the ability to generate many versions of a single story, much in a spirit of our own work. The STORYBOOK system could generate multiple versions of the "Little Red Riding Hood" fairy tale, showing both syntactic and lexical variation. However, STORYBOOK is based on highly handcrafted mappings from plans to the FUF-SURGE realizer [12] and is thus applicable only to the "Little Red Riding Hood" domain.

To our knowledge, the only work that begins to address the semantic-to-syntactic mapping within the storytelling domain is Elson's Scheherazade story annotation tool [13]. Scheherazade allows naïve users to annotate stories with a rich symbolic representation called a STORY INTENTION GRAPH. This representation is robust and linguistically grounded which makes it a good candidate for a content representation in an NLG pipeline.

In this paper, we present a working model of reproducing different tellings of a story from its symbolic representation. In Sec. 2 we explain how we build on the STORY INTENTION GRAPH representation provided by Scheherazade [14] and the previous work on NLG for interactive stories based on extensions to the PERSONAGE NLG engine [15, 9]. Sec. 3 presents an automatic method for converting from Scheherazade’s STORY INTENTION GRAPH output to the input required by PERSONAGE. Using the corpus of semantic SIG representations for 36 Aesop Fables that are distributed in DramaBank [14], we train translation rules on one story and then test these rules by generating 35 other stories in the collection. Figs. 1 and 8 show two Aesop Fables, with both Scheherazade and PERSONAGE generated versions. “The Fox and the Grapes” was the development story for our work, while “The Lion and the Boar” was part of the test set. Fig. 9 gives a feel for the retellings that PERSONAGE is capable to produce once it is coupled with Scheherazade’s story representation. Sec. 4 demonstrates our evaluation of the 35 generated stories using the measures of Levenshtein Distance and BLEU score. Sec. 5 summarizes and discusses our future work, where we aim to experiment with models for narrator’s and characters’ voices, measures of retellings’ quality, and with techniques for making story telling interactive.

2 Background

Our work is based on a simple observation: a novel capability for interactive stories can be developed by bridging two off-the-shelf linguistic tools, Scheherazade and PERSONAGE. We integrated these tools in a standard NLG pipeline:

- Content planner that introduces characters and events
- Sentence planner that creates linguistic representations for those events
- Surface realizer that produces text string out of linguistic structures

Scheherazade produces the output we would normally get from a content planner and PERSONAGE plays the role of the sentence planner and surface realizer. We developed an algorithm that creates the semantic linguistic representation from a conceptual narrative structure provided by Scheherazade and generates from it using PERSONAGE. Our algorithm acts as an intermediary producing a semantic-into-syntactic mapping. The immediate goal for this study was to regenerate directly from stories annotated with Scheherazade. Our long term goal is to build on the created infrastructure and turn the intermediary into a conventional sentence planner, capable of one-to-many semantic-to-syntactic mappings, i.e. retelling a story in different voices.

We believe that the combination of Scheherazade and PERSONAGE is a perfect fit for our long-term goal due to the following aspects of the two systems. First, Scheherazade representations are already lexically anchored into WordNet and VerbNet ontologies which allows for lexical variation. Second, PERSONAGE provides 67 parameters that make it already capable of generating many pragmatic and stylistic variations of a single utterance. Below we discuss the functionality of the two systems in more detail.

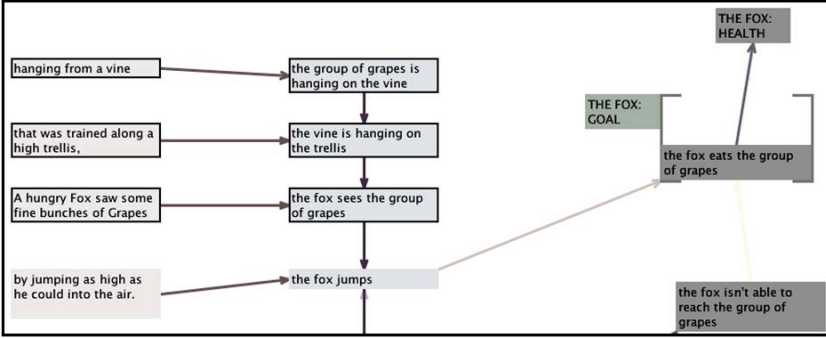


Fig. 2. Part of the STORY INTENTION GRAPH (SIG) for “The Fox and the Grapes”

Scheherazade. Scheherazade is an annotation tool that facilitates the creation of a rich symbolic representation for narrative texts, using a schemata known as the STORY INTENTION GRAPH or SIG [13]. An example SIG for “The Fox and the Grapes” development story, reproduced from Elson [14], is shown in Fig. 2. The annotation process involves sequentially labeling the original story sentences according to the SIG formalism using Scheherazade’s GUI. The annotators instantiate characters and objects, assign actions and properties to them, and provide their interpretation of *why* characters are motivated to take the actions they do. Scheherazade’s GUI features a built-in generation module in the what-you-see-is-what-you-mean paradigm (WYSIWYM) to help naïve users produce correct annotations by letting them check the natural language realization of their encoding as they annotate [16]). Scheherazade does have a built in surface generation engine, but it is inflexible with synonyms and syntax, thus why we are interested in utilizing PERSONAGE. As a baseline comparison for our method, we use Scheherazade’s generation engine to evaluate the correctness of PERSONAGE outputs.

One of the strengths of Scheherazade is that it allows users to annotate a story along several dimensions, starting with the surface form of the story (see first column in Fig. 2) and then proceeding to deeper representations. The first dimension (see second column in Fig. 2) is called the “timeline layer”, in which the story facts are encoded as predicate-argument structures (propositions) and temporally ordered on a timeline.

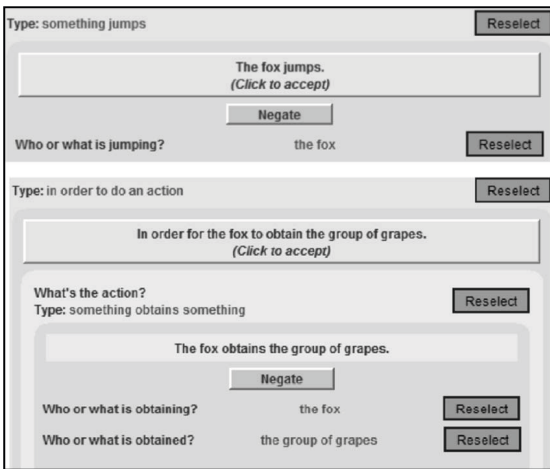


Fig. 3. GUI view of propositional modeling

The timeline layer consists of a network of propositional structures, where nodes correspond to lexical items that are linked by thematic relations. Scheherazade adapts information about predicate-argument structures from the VerbNet lexical database [17] and uses WordNet [18] as its noun and adjectives taxonomy. The arcs of the story graph are labeled with discourse relations. Fig. 3 shows a GUI screenshot of assigning propositional structure to the sentence *The fox jumped in order to obtain the group of grapes*. This sentence is encoded as two nested propositions `jump(fox)` and `obtain(fox, group of grapes)`. Both actions (`jump` and `obtain`) contain references to the story characters and objects (`fox` and `grapes`) that fill in slots corresponding to semantic roles.

The second dimension (see third column in Fig. 2) is called the “interpretative layer”; this layer goes beyond summarizing the actions and events that occur, but attempts to capture story meaning derived from agent-specific plans, goals, attempts, outcomes and affectual impacts. To date, we only utilize the event timeline layer of the SIG encoding.

Model	Parameter	Description	Example	
Shy voice	SOFTENER	Insert syntactic elements (<i>sort of, kind of, somewhat, quite, around, hungry</i>) to mitigate the strength of a proposition	<i>'It seems to me that he was rather, I think that, it seems that, it seems to me that)</i>	
	HEDGES			
	STUTTERING			<i>'The vine hung on the trellis'</i>
	FILLED PAUSES			Insert syntactic elements expressing hesitancy (<i>I mean, err, mmhm, like, you know</i>)
Laid-back voice	EMPHASIZER	Insert syntactic elements (<i>really, basically, actually</i>) to strengthen a proposition	<i>'The fox failed to get the group of grapes, alright?'</i>	
	HEDGES			
	EXCLAMATION	Insert an exclamation mark	<i>'The group of grapes hung on the vine!'</i>	
	EXPLETIVES	Insert a swear word	<i>'The fox was damn hungry'</i>	

Fig. 4. Examples of pragmatic marker insertion parameters from PERSONAGE

Personage. PERSONAGE is an NLG engine that has the ability to generate a single utterance in many different voices. Models of narrative style are currently based on the Big Five personality traits [15], or are learned from film scripts [9]. Each type of model (personality trait or film) specifies a set of language cues, one of 67 different parameters, whose value varies with the personality or style to be conveyed. Fig. 4 shows a subset of parameters, which were used in stylistic models to produce “The Fox and the Grapes” in different voices (see Fig. 9). Previous work [15] has shown that humans perceive the personality stylistic models in the way that PERSONAGE intended.

After selecting a stylistic model for each utterance, PERSONAGE uses the off-the-shelf surface realizer REALPRO [19]. PERSONAGE outputs a sentence-plan tree whose internal representations are deep syntactic structures (DSyntS) that RealPro expects as input. DSyntS provides a flexible dependency tree representation of an utterance which can be altered by the PERSONAGE parameter settings. The nodes of the DSyntS syntactic trees are labeled with lexemes and the arcs of the tree are labeled with syntactic relations. The DSyntS formalism

```

<dsyntnode class="verb" lexeme="jump" mood="ind" tense="past">
  <dsyntnode article="def" class="common_noun" lexeme="fox" number="sg" rel="I"/>
  <dsyntnode class="preposition" lexeme="in_order_to" rel="ATTR">
    <dsyntnode class="verb" lexeme="obtain" mood="inf" rel="II">
      <dsyntnode article="def" class="common_noun" lexeme="group" number="sg" rel="III">
        <dsyntnode class="common_noun" lexeme="grape" number="pl" rel="III"/>
      </dsyntnode>
    </dsyntnode>
  </dsyntnode>

```

Fig. 5. DSyntS for a sentence *The fox jumped in order to obtain the group of grapes*

distinguishes between arguments and modifiers and between different types of arguments (subject, direct and indirect object etc). Lexicalized nodes also contain a range of grammatical features used in generation. RealPro handles morphology, agreement and function words to produce an output string. Fig. 5 shows an example DSyntS structure for one of the sentences from our development story “The Fox and the Grapes”. Feature values in bold need to be obtained from the internal Scheherazade representation.

3 Method

Our method applies a model of syntax to a Scheherazade representation of a story (a SIG encoding), in order to produce a retelling of the story in a different voice. A prerequisite for producing stylistic variations of a story is an ability to generate a “correct” retelling of the story. The focus of this study is to verify that the essence of a story is not distorted as we move from one formal representation of the story (SIG) to another (DSyntS). We use Scheherazade’s built-in generator, which was customized to the SIG schemata, as a baseline to evaluate our results. Our data comes from DramaBank, a collection of Scheherazade annotated texts ranging from Aesop fables to contemporary nonfiction [14]. Aesop fables from DramaBank serve as our dataset: **one** fable (seen in Fig. 1) is used in development, and then our method is tested by automatically transforming the 35 other fables to the PERSONAGE representation. Figs. 1 and 8 show two Aesop fables, with both Scheherazade and PERSONAGE generated versions. “The Fox and the Grapes” was the development story for our work, while “The Lion and the Boar” was part of the test set. Fig. 6 shows an overview of the method consisting of the following steps:

1. Use Scheherazade’s built-in generation engine to produce text from the SIG encoding of the fable
2. Manually construct DSyntS corresponding to the text generated in step 1 (follow the right branch in Fig. 6)
3. Derive semantic representation of a fable from the SIG encoding using Scheherazade API (follow the left branch of Fig. 6)
4. Informed by our understanding of the two formalisms, develop transformation rules to build DSyntS from semantic representation

5. Apply rules to the semantic representation derived in step 3
6. Feed automatically produced DSyntS to PERSONAGE using a neutral NLG model and compare the output with that of step 1. The metrics we used for string comparison are discussed in Sec. 4.

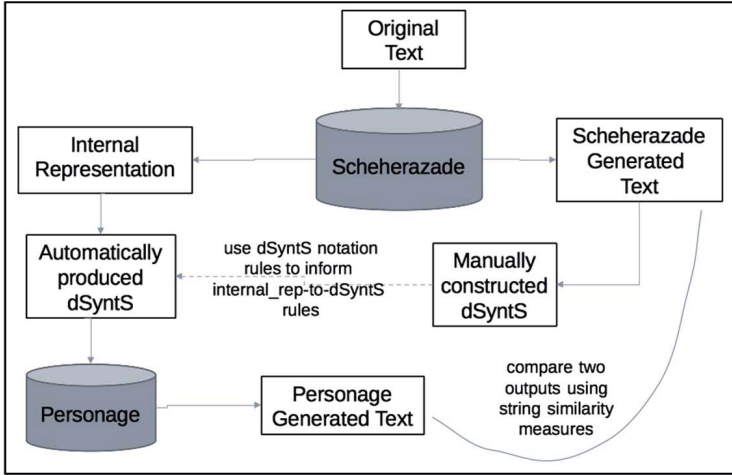


Fig. 6. Our Method

The primary technical challenge was developing a general mechanism for converting SIG semantic encodings into the DSyntS representation used by PERSONAGE’s generation dictionary. This involved enforcing syntactic tree structure over the chains of propositions. The challenge was partly due to the variety of ways that VerbNet and WordNet allow nuances of meaning to be expressed. For example, a sentence *the crow was sitting on the branch of the tree* has two alternative encodings depending on what is important about crow’s initial disposition: the crow can *be sitting* as an adjectival modifier, or can be *sitting* as a progressive action. There are also gaps in Scheherazade’s coverage of abstract concepts, which can lead to workarounds on the part of the annotators that are not easily undone by surface realization (an example is *whether later first drank* in Fig. 8, where adjectival modifiers are used to express *which of [the characters] should drink first* from the original text).

The transformation of Scheherazade’s semantic representation into syntactic dependency structure is a multi-stage process, illustrated in Fig. 7. First, a syntactic tree is constructed from the propositional event structure. Element A in Fig. 7 contains a sentence from the original “The Fox and the Grapes” fable. We use the Scheherazade API to process the fable text together with its SIG encoding and extract actions associated with each timespan of the timeline layer. Element B in Fig. 7 shows a schematic representation of the propositional structures. Each action instantiates a separate tree construction procedure. For each action, we create a verb instance (see highlighted nodes of element D in Fig. 7). We use information about the predicate-argument frame that the action

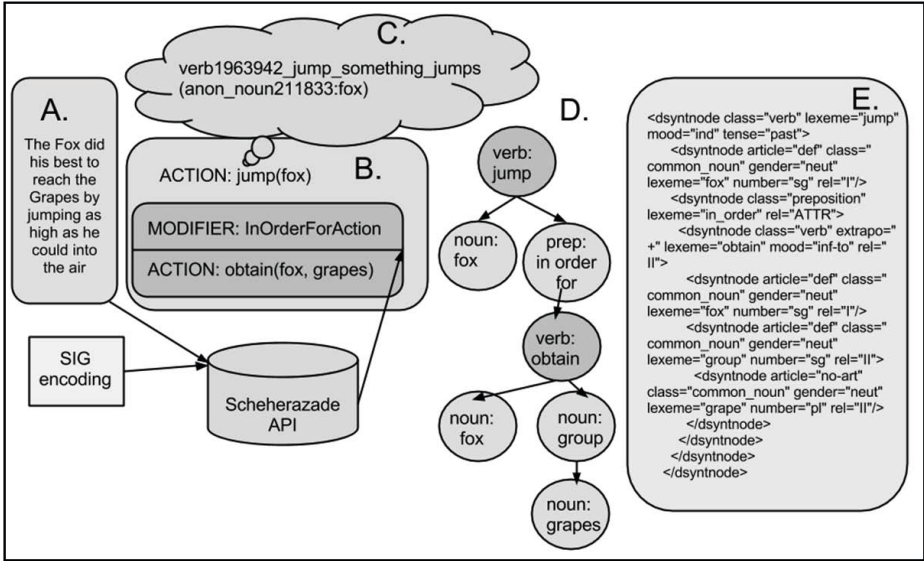


Fig. 7. Step by step transformation from SIG to DSyntS

invokes (see element C in Fig. 7) to map frame constituents into respective lexico-syntactic classes, for example, characters and objects are mapped into nouns, properties into adjectives and so on. The lexico-syntactic class aggregates all of the information that is necessary for generation of a lexico-syntactic unit in DSyntS (element E in Fig. 7). We define 6 classes corresponding to main parts of speech: noun, verb, adverb, adjective, functional word. Each class has a list of properties such as morphology or relation type that are required by the DSyntS notation for a correct rendering of a category. For example, all classes include a method that parses frame type in the SIG to derive the base lexeme. The methods to derive grammatical features are class-specific. Each lexico-syntactic unit refers to the elements that it governs syntactically thus forming a hierarchical structure. A separate method collects the frame adjuncts as they have a different internal representation in the SIG.

At the second stage, the algorithm traverses the syntactic tree in-order and creates an XML node for each lexico-syntactic unit. Class properties are then written to disk, and the resulting file (see element E in Fig. 7) is processed by the surface realizer to generate text. Fig. 8 shows the “Lion and the Boar” fable from our test set, with its generated versions.

Although the emphasis of this work is on generating a “correct” retelling of a story, this infrastructure allows us to alter the story stylistically. Fig. 9 illustrates how we can now piggy-back on the transformations that PERSONAGE can make to produce different retellings of the same story. The FORMAL voice triggered no stylistic parameters of PERSONAGE. “The Fox and the Grapes” was generated directly from DSyntS by the surface realizer. Fig. 4 provides examples of how different parameters played out for the SHY and LAID-BACK voices.

Original Story	Scheherazade	Personage
On a summer day, when the great heat induced a general thirst, a Lion and a Boar came at the same moment to a small well to drink. They fiercely disputed which of them should drink first, and were soon engaged in the agonies of a mortal combat. On their stopping on a sudden to take breath for the fiercer renewal of the strife, they saw some Vultures waiting in the distance to feast on the one which should fall first. They at once made up their quarrel, saying: "It is better for us to make friends, than to become the food of Crows or Vultures, as will certainly happen if we are disabled."	Once, the air was hot. A boar decided to drink from a spring, and a lion decided to drink from the spring. The boar quarrelled about whether later first drank from the spring, and the lion quarrelled about whether later first drank from the spring. The boar began to attack the lion, and the lion began to attack the boar. The boar stopped attacking the lion, and the lion stopped attacking the boar. The boar above saw a group of vultures being seated on some rock, and the lion above saw the group of vultures being seated on the rock. The group of vultures began to plan – if the boar were to later die, and the lion were to later die – for the group of vultures to eat. The boar sobered, and the lion sobered. The boar said to the lion that – if the lion were to not kill the boar – the group of vultures would not eat the boar, and the lion said to the boar that – if the boar were to not kill the lion – the group of vultures would not eat the lion. The boar didn't kill the lion, and the lion didn't kill the boar.	The air was hot. The lion decided the lion drank from the spring. The boar decided the boar drank from the spring. The boar quarreled. The lion quarreled. The lion attacked the boar. The boar attacked the lion. The boar above saw the group of vultures was seated on the rock. The lion above saw the group of vultures was seated on the rock. The group of vultures planned the group of vultures ate. The boar sobered. The lion sobered. The boar said the group of vultures did not eat the boar to the lion. The lion said the group of vultures did not eat the lion to the boar. The boar did not kill the lion. The lion did not kill the boar.

Fig. 8. "The Lion and the Boar" Fable from the test set, with Generated Versions

Personage FORMAL	Personage SHY	Personage LAID-BACK
The group of grapes hung on the vine. The vine hung on the trellis. The fox saw the group of grapes. The fox jumped in order to obtain the group of grapes. The fox didn't obtain the group of grapes because he couldn't reach the group of grapes.	Well, the group of grapes hung on the vine. The vine hung on the tr-tr-trellis. The fox saw the group of grapes. It seemed that he was hungry. Err... the fox jumped in order to obtain the group of grapes. The fox didn't collect the group of grapes.	Ok, the group of grapes rested on the vine. The vine hung on the trellis, didn't it? The fox saw the group of grapes. He damn was hungry. The fox jumped in order to obtain the group of grapes, okay? Oh the fox failed to get the group of grapes, didn't obtain it, because he couldn't reach the group of grapes, you see?

Fig. 9. Retellings of "The Fox and the Grapes"

The corresponding renderings of the story demonstrate lexical (*grapes hung/grapes rested*), syntactic (*didn't get/failed to get*) and stylistic variation.

4 Results

To evaluate the performance of our translation rules we compare the output generated by PERSONAGE to that of Scheherazade's built-in realizer, using two metrics: BLEU score [20] and Levenshtein distance. *BLEU* is an established standard for evaluating the quality of machine translation and summarization systems. The score between 0 and 1 measures the closeness of two documents by comparing n-grams, taking word order into account. *Levenshtein distance* is the

minimum edit distance between two strings. The objective is to minimize the total cost of character deletion, insertion, replacement that it takes to transform one string into another. In our case, we treat each word as a unit and measure word deletion, insertion, and replacement. We used word stemming as a preprocessing step to reduce the effect of individual word form variations. The results are shown in Table 1. **Scheherazade-Personage** compares the output of the PERSONAGE generator produced through our automatic translation rules to that of the Scheherazade generation. Because the rules were created on the development set to match the Scheherazade story, we expect these results to provide a topline for comparison to the test set, shown in the bottom table of Table 1.

Table 1. Mean and Standard Deviation for Levenshtein Distance (Lower is Better), and BLEU (Higher is Better) on both the DEVELOPMENT and TEST sets

DEVELOPMENT	Levenshtein	BLEU
Scheherazade-Personage	31	.59
Fable-Scheherazade	80	.04
Fable-Personage	84	.03

TEST	Levenshtein Mean (STD)	BLEU Mean (STD)
Scheherazade-Personage	72 (40)	.32 (.11)
Fable-Scheherazade	116 (41)	.06 (.02)
Fable-Personage	108 (31)	.03 (.02)

Although our rules were informed by looking at the Scheherazade generation, Table 1 also includes measures of the distance between the original fable and both Scheherazade and PERSONAGE. **Fable-Scheherazade** and **Fable-Personage** compare the original fable to the Scheherazade and to the PERSONAGE generation respectively. Note that these results should not be compared to **Scheherazade-Personage** since both of these are outputs of an NLG engine. The two-tailed Student’s t-test was used to compare the two realizers’ mean distance values to the original fables and determine statistical significance. The difference in Levenshtein distance between **Fable-Scheherazade** and **Fable-Personage** is not statistically significant ($p = 0.08$) on both development and test sets. This indicates that our rules generate a story which is similar to what Scheherazade generates in terms of closeness to the original. However, Scheherazade shows a higher BLEU score with the original fables ($p < 0.001$). We believe that this is due to the fact that our translation rules assume a simple generation. The rules did not attempt to express tense, aspect or stative/non-stative complications, all of which contribute to a lower overlap of n-grams. The assignment of tense and aspect was an area of particular focus for Scheherazade’s built-in realizer [21].

5 Conclusions and Future Work

In this paper we show that: (1) Scheherazade’s SIG annotation schemata provides a rich and robust story representation that can be linked to other generation engines; (2) we can integrate Scheherazade and PERSONAGE (two off-the-shelf tools) for representing and producing narrative, thus bridging the gap between content and sentence planning in the NLG pipeline; and (3) we have created an infrastructure which puts us in a position to reproduce a story in different voices and styles. In this paper, we presented quantitative results using Levenshtein distance and BLEU score. However, since our long term goal is to generate different retellings, these metrics will be inappropriate. Here we were primarily concerned with the correctness of the content of the generators; in future work we will need to develop new metrics or new ways of measuring the quality of story retellings. In particular we plan to compare human subjective judgements of story quality across different retellings.

There are also limitations of this work. First, the PERSONAGE realizer needs to extend its generation dictionary in order to deal with irregular forms. The current system generated incorrect forms such as *opened*, and *wifes*. Also we were not able to make tense distinctions in our generated version, and generated everything in past simple tense. In addition, we noted problems with the generation of distinct articles when needed such as *a* vs. *the*. There are a special set of surface realization rules in Scheherazade that are currently missing from PERSONAGE that adds cue phrases such as *that* and *once*. We aim to address these problems in future work.

It should be mentioned that despite being domain independent, our method relies on manual story annotations to provide content for the generation engine. DramaBank is the result of a collection experiment using trained annotators; as they became familiar with the tool, the time that the annotators took to encode each fable (80 to 175 words) as a SIG encoding dropped from several hours to 30-45 minutes on average [14]. The notion of achieving the same semantic encoding using automatic methods is still aspirational. While the SIG model overlaps with several lines of work in automatic semantic parsing, as it has aspects of annotating attribution and private states [22], annotating time [23] and annotating verb frames and semantic roles [24], there is not yet a semantic parser that can combine these aspects into an integrated encoding, and developing one falls outside of scope of this work.

In future work, we also aim to do much more detailed studies on the process of generating the same story in different voices, using the apparatus we present here. Examples of stylistic story variations presented in this paper come from modifications of narrator’s voice. In future work, we plan to apply stylistic models to story characters. An extension to Scheherazade to distinguish direct and indirect speech in the SIG will allow give characters expressive, personality driven voices. Once a story has been modeled symbolically, we can realize it in multiple ways, either by different realizers or by the same realizer in different modes.

References

1. Thorne, A.: The press of personality: A study of conversations between introverts and extraverts. *Journal of Personality and Social Psychology* 53, 718 (1987)
2. Thorne, A., McLean, K.C.: Telling traumatic events in adolescence: A study of master narrative positioning. *Connecting Culture and Memory: The Development of an Autobiographical Self*, 169–185 (2003)
3. Callaway, C., Lester, J.: Narrative prose generation. *Artificial Intelligence* 139, 213–252 (2002)
4. Turner, S.: The creative process: A computer model of storytelling and creativity. Lawrence Erlbaum (1994)
5. Riedl, M., Young, R.M.: An intent-driven planner for multi-agent story generation. In: Proc. of the 3rd International Conference on Autonomous Agents and Multi Agent Systems (2004)
6. Cavazza, M., Charles, F.: Dialogue generation in character-based interactive storytelling. In: AAAI First Annual Artificial Intelligence and Interactive Digital Entertainment Conference, Marina del Rey, California, USA (2005)
7. Rowe, J.P., Ha, E.Y., Lester, J.C.: Archetype-Driven Character Dialogue Generation for Interactive Narrative. In: Prendinger, H., Lester, J.C., Ishizuka, M. (eds.) IVA 2008. LNCS (LNAI), vol. 5208, pp. 45–58. Springer, Heidelberg (2008)
8. Lin, G.I., Walker, M.A.: All the world's a stage: Learning character models from film. In: Proceedings of the Seventh AI and Interactive Digital Entertainment Conference, AIIDE 2011. AAAI (2011)
9. Walker, M.A., Grant, R., Sawyer, J., Lin, G.I., Wardrip-Fruin, N., Buell, M.: Perceived or not perceived: Film character models for expressive NLG. In: Si, M., Thue, D., André, E., Lester, J., Tanenbaum, T.J., Zammitto, V. (eds.) ICIDS 2011. LNCS, vol. 7069, pp. 109–121. Springer, Heidelberg (2011)
10. Montfort, N.: Generating narrative variation in interactive fiction. PhD thesis, University of Pennsylvania (2007)
11. McIntyre, N., Lapata, M.: Learning to tell tales: A data-driven approach to story generation. In: Proceedings of the Joint Conference of the 47th Annual Meeting of the ACL and the 4th International Joint Conference on Natural Language Processing of the AFNLP, Singapore, pp. 217–225 (2009)
12. Elhadad, M.: Using Argumentation to Control Lexical Choice: a Functional Unification Implementation. PhD thesis, Columbia University (1992)
13. Elson, D.K., McKeown, K.: A tool for deep semantic encoding of narrative texts. In: Proceedings of the ACL-IJCNLP 2009 Software Demonstrations, pp. 9–12. Association for Computational Linguistics (2009)
14. Elson, D.K.: Detecting story analogies from annotations of time, action and agency. In: Proceedings of the LREC 2012 Workshop on Computational Models of Narrative, Istanbul, Turkey (2012)
15. Mairesse, F., Walker, M.A.: Controlling user perceptions of linguistic style: Trainable generation of personality traits. *Computational Linguistics* (2011)
16. Bouayad-Agha, N., Scott, D., Power, R.: Integrating content and style in documents: a case study of patient information leaflets. *Information Design Journal* 9, 161–176 (2000)
17. Kipper, K., Korhonen, A., Ryant, N., Palmer, M.: Extensive classifications of english verbs. In: Proceedings of the Third International Conference on Language Resources and Evaluation (2006)
18. Fellbaum, C.: *WordNet: An Electronic Lexical Database*. MIT Press (1998)

19. Lavoie, B., Rambow, O.: A fast and portable realizer for text generation systems. In: Proceedings of the Third Conference on Applied Natural Language Processing, ANLP 1997, pp. 265–268 (1997)
20. Papineni, K., Roukos, S., Ward, T., Zhu, W.J.: Bleu: a method for automatic evaluation of machine translation. In: Proceedings of the 40th Annual Meeting on Association for Computational Linguistics, ACL 2002, pp. 311–318. Association for Computational Linguistics, Stroudsburg (2002)
21. Elson, D.K., McKeown, K.: Tense and aspect assignment in narrative discourse. In: Proceedings of the Sixth International Conference on Natural Language Generation, INLG 2010. Citeseer (2010)
22. Wiebe, J.M.: Identifying subjective characters in narrative. In: Proceedings of the 13th Conference on Computational Linguistics, COLING 1990, vol. 2, pp. 401–406. Association for Computational Linguistics, Stroudsburg (1990)
23. Pustejovsky, J., Castao, J., Ingria, R., Saur, R., Gaizauskas, R., Setzer, A., Katz, G.: Timeml: Robust specification of event and temporal expressions in text. In: Fifth International Workshop on Computational Semantics, IWCS-5 (2003)
24. Palmer, M., Gildea, D., Kingsbury, P.: The proposition bank: An annotated corpus of semantic roles. *Comput. Linguist.* 31, 71–106 (2005)