

M2D: Monolog to Dialog Generation for Conversational Story Telling

Kevin K. Bowden, Grace I. Lin, Lena I. Reed,
Jean E. Fox Tree, and Marilyn A. Walker

Natural Language and Dialog Systems Lab
University of California, Santa Cruz
{kbowden, glin5, lireed, foxtree, mawwalker}@ucsc.edu

Abstract. Storytelling serves many different social functions, e.g. stories are used to persuade, share troubles, establish shared values, learn social behaviors, and entertain. Moreover, stories are often told conversationally through dialog, and previous work suggests that information provided dialogically is more engaging than when provided in monolog. In this paper, we present algorithms for converting a deep representation of a story into a dialogic storytelling, that can vary aspects of the telling, including the personality of the storytellers. We conduct several experiments to test whether dialogic storytellings are more engaging, and whether automatically generated variants in linguistic form that correspond to personality differences can be recognized in an extended storytelling dialog.

Keywords: ICIDS, Dialog, Natural Language Generation, Personality, Conversational Storytelling, Analysis and Evaluation of Systems

1 Introduction

Storytelling serves many different social functions, e.g. stories are used to persuade, share troubles, establish shared values, learn social behaviors, and entertain [33, 24]. Moreover, stories are often told conversationally through dialog [39, 38] where the telling of a story is shaped by the personality of both the teller and the listener. For example, extraverted friends actively engage one another in constructing the action of the story by peppering the storyteller with questions, and by asking the listener to guess what happened [39, 38]. Thus the same story can be told in many different ways, often achieving different effects [22, 40].

A system capable of telling a story and then retelling it in different settings to different audiences requires two components: 1) a deep representation of the STORY and 2) algorithms that render the story content as different DISCOURSE instantiations. A deep representation of the story’s content, often called the STORY or FABULA, must specify the events, characters, and props of the story, as well as relations among them, including reactions of characters to story events. This is accomplished through EST [31], a framework that bridges the story annotation tool SCHEHERAZADE and a natural language generator (NLG).

The DISCOURSE representation is the surface rendering of the FABULA, an instantiated expressive telling of a story as a stream of words, gestures or expressions [3, 8, 28, 29]. This paper presents M2D, a framework with algorithms that manipulate the story content to retell the story as a conversational dialog between two people. An example of the original, monologic, and dialogic telling of the Garden Story is shown in Figure 1. Note that highlighted areas indicate the same information being presented differently at different stages.

Garden Story: Original	
<p>Today when I arrived at my community garden plot, it actually looked like a garden. Not a weedy mess with maybe some stuff growing in it if you know where to look. We had hit the typical mid-summer mess of fast-growing weeds and no time to do anything about it. Plus all the rain had made a huge swamp and it was hard to get a moment to work when it wasn't actively pouring. I put in a bunch of time this past week, and it's paying off. Along with free-standing non-weed-choked plants, I have now re-planted three of the beds with salad greens, spinach, and chard. And while the viability of the seeds was questionable, I accidentally unearthed some from the bed I planted 2 days ago and they had already started to sprout! ... <i>etc.</i></p>	
Garden Story: Monolog generated by EST	Garden Story: Dialog generated by M2D
<p>The radishes charmed the butterflies. The communal garden was weedy. It rained. The communal garden was swampy. It rained. The productive gardener planted the plants. The gardener planted the chards the lettuces and the spinach. The pleased gardener did not expect for the chards the lettuces and the spinach to grow. The chards the lettuces and the spinach sprouted. The gardener mistakenly dug the chards the lettuces and the spinach. The surprised gardener saw for the chards the lettuces and the spinach to sprout. The communal garden was not weedy. The communal garden was not swampy. The communal garden was productive. The gardener was proud. The eager gardener wanted to reap the lettuces. The radishes were droopy. The gardener planned to remove the radishes. The thoughtful gardener thought the flowers charmed the butterflies.</p>	<p>S1: The radishes charmed the butterflies. Technically, the damn communal garden was not weedless. Err ... it rained. Great right, it was really swampy.</p> <p>S2: Yeah, it rained. The productive gardener planted the plants. I mean, she planted the chards the lettuces and the spinach.</p> <p>S1: She did not expect for them to grow. The chards the lettuces and the spinach sprouted, didn't they?</p> <p>S2: Unfortunately, the gardener mistakenly dug them. She saw for the chards the lettuces and the spinach to sprout.</p> <p>S1: The communal garden was not weedy. It was pretty productive and not swampy.</p> <p>S2: Mmhm ... the gardener was proud and wanted to reap the lettuces. I thought everybody knew that they were quite droopy? The radishes were droopy.</p> <p>S1: I see, well, she planned to remove them. The thoughtful gardener thought the flowers charmed the butterflies.</p>

Fig. 1. Garden Story: Original Version and Monologue/Dialog Generation. Highlighted areas indicate examples of the same information.

We build on the publicly available PersonaBank corpus ¹, which provides us with the deep story representation and a lexico-syntactic representation of its monologic retelling [14]. PersonaBank consists of a corpus of monologic personal narratives from the ICWSM Spinn3r Corpus [6] that are annotated with a deep story representation called a STORY INTENTION GRAPH [13]. After annotation, the stories are run through the EST system to generate corresponding deep linguistic structure representations. M2D then takes these representations as input and creates **dialog with different character voices**. We identify several stories by hand as good candidates for dialogic tellings because they describe events or experiences that two people could have experienced together.

Our primary hypothesis is **H1**: Dialogic tellings of stories are more engaging than monologic tellings. We also hypothesize that good dialog requires the use

¹ Available from nlds.soe.ucsc.edu/personabank

of narratological variations such as direct speech, first person, and focalization [14]. Moreover, once utterances are rendered as first-person with direct speech, then character voice becomes relevant, because it does not make sense for all the characters and the narrator to talk in the same voice. Thus our primary hypothesis **H1** entails two additional hypotheses **H2** and **H3**:

H2: Narratological variations such as direct speech, first person, and focalization will affect a readers engagement with a story.

H3: Personality-based variation is a key aspect of expressive variation in storytelling, both for narrators and story characters. Changes in narrator or character voice may affect empathy for particular characters, as well as engagement and memory for a story.

Our approach to creating different character voices is based on the **Big Five** theory of personality [1, 9]. It provides a useful level of abstraction (e.g., extraverted vs. introverted characters) that helps to generate language and to guide the integration of verbal and nonverbal behaviors [16, 21, 11].

To the best of our knowledge, our work is the first to develop and evaluate algorithms for automatically generating different dialogic tellings of a story from a deep story representation, and the first to evaluate the utility and effect of parameterizing the style of speaker voices (personality) while telling the story.

2 Background and Motivation

Stories can be told in either dialog or as a monolog, and in many natural settings storytelling is conversational [4]. Hypothesis **H1** posits that dialogic tellings of stories will be more engaging than monologic tellings. In storytelling and at least some educational settings, dialogs have cognitive advantages over monologs for learning and memory. Students learn better from a verbally interactive agent than from reading text, and they also learned better when they interacted with the agent with a personalized dialog (whether spoken or written) than a non-personalized monolog [20]. Our experiments compare different instances of the dialog, e.g. to test whether more realistic conversational exchanges affects whether people become immersed in the story and affected by it.

Previous work supports **H2**, claiming that direct, first-person speech increases stories’ drama and memorability [34, 37]. Even when a story is told as a monolog or with third person narration, dialog is an essential part of storytelling: in one study of 7 books, between 40% and 60% of the sentences were dialog [7]. In general narratives are mentally simulated by readers [35], but readers also enact a protagonist’s speech according to her speech style, reading more slowly for a slow-speaking protagonist and more quickly for a fast-speaking protagonist, both out-loud and silently [43]. However, the speech simulation only occurred for direct quotation (e.g. *She said “Yeah, it rained”*), not indirect quotation (e.g. *She said that it had rained*). Only direct quotations activate voice-related parts of the brain [43], as they create a more vivid experience, because they express enactments of previous events, whereas indirect quotations describe events [42].

Several previous studies also suggest **H3**, that personality-based variation is a key aspect of storytelling, both for narrators and story characters. Personality traits have been shown to affect how people tell stories as well as their choices of stories to tell [17]. And people also spontaneously encode trait inferences from everyday life when experiencing narratives, and they derive trait-based explanations of character’s behavior [30, 32]. Readers use these trait inferences to make predictions about story outcomes and prefer outcomes that are congruent with trait-based models [30]. The finding that the behavior of the story-teller is affected by the personality of both the teller and the listener also motivates our algorithms for monolog to dialog generation [38, 39]. Content allocation should be controlled by the personality of the storyteller (e.g. enabling extraverted agents to be more verbose than introverted agents).

Previous work on generation for fictional domains has typically combined STORY and DISCOURSE, focusing on the generation of STORY events and then using a direct text realization strategy to report those events [18]. This approach cannot support generation of different tellings of a story [23]. Previous work on generating textual dialog from monolog suggests the utility of adding extra interactive elements (dialog interaction) to storytelling and some strategies for doing so [2, 36, 25]. In addition, expository or persuasive content rendered as dialog is more persuasive and memorable [41, 26, 27]. None of this previous work attempts to generate dialogic storytelling from original monologic content.

3 M2D: Monolog-to-Dialog Generation

Figure 2 illustrates the architecture of M2D. The EST framework produces a story annotated by SCHEHERAZADE as a list of Deep Syntactic Structures (**DsyntS**). DsyntS, the input format for the surface realizer RealPro [12, 19], is a dependency-tree structure where each node contains the lexical information for the important words in a sentence. Each sentence in the story is represented as a DsyntS.

M2D converts a story (as a list of DsyntS) into different versions of a two-speaker dialog using a parameterizable framework. The input parameters control, for each speaker, the allocation of content, the usage of questions of different forms, and the usage of various pragmatic markers (Table 1). We describe the M2D parameters in more details below.

Content allocation: We allocate the content of the original story between the two speakers using a content-allocation parameter that ranges from 0 to 1. A value of .5 means that the content is equally split between 2 speakers. This is motivated by the fact that, for example, extraverted speakers typically provide more content than introverted speakers [16, 38].




Fig. 2. M2D Pipeline Architecture.

Table 1. Dialog Conversion Parameters

Parameter	Example
Aggregation	
Merge short sents	The garden was swampy, and not productive.
Split long sents	The garden was very swampy because it rained. The garden is very large, and has lots of plants.
Coreference	
Pronominalize	The gardener likes to eat apples from his orchard. They are red.
Pragmatic Markers	
Emphasizer_great	Great, the garden was swampy.
Downer_kind_of	The garden was kind of swampy.
Acknowledgment_yeah	Yeah, the garden was swampy.
Repetition	S_1 : The garden was swampy. S_2 : Yeah, the garden was swampy.
Paraphrase	S'_2 : Right, the garden was boggy.
Interactions	
Affirm Adjective	S_1 : The red apples were tasty and --- S_2 : Just delicious, really. S_1 : Yeah, and the gardener ate them.
Correct Inaccuracies	S_2 : The garden was not productive and ---. S_1 : I don't think that's quite right, actually. I think the garden was productive.
Questions	
Provoking With_Answer	I don't really remember this part, can you tell it? S_1 : How was the garden? S_2 : The garden was swampy.
Tag	The garden was swampy, wasn't it?

Character and Property Database: We use the original source material for the story to infer information about actors, items, groups, and other properties of the story, using the information specified in the DsyntS. We create actor objects for each character and track changes in the actor states as the story proceeds, as well as changes in basic properties such as their their body parts and possessions. **Aggregation and Deaggregation:** We break apart long DsyntS into smaller DsyntS, and then check where we can merge small and/or repetitious DsyntS. We believe that deaggregation will improve our dialogs overall clarity while aggregation will make our content feel more connected [15].

Content Elaboration: In natural dialog, speakers often repeat or partially paraphrase each other, repeating the same content in multiple ways. This can be a key part of entrainment. Speakers may also ask each other questions thereby setting up frames for interaction [38]. In our framework, this involves duplicating content in a single DsyntS by either 1) generating a question/answer pair from it and allocating the content across speakers,

or 2) duplicating it and then generating paraphrases or repetitions across speakers. Questions are generated by performing a series of pruning operations based on the class of the selected node and the relationship with its parent and siblings. For example, if STORE in Figure 3 is selected, we identify this node as our question. The class of a node indicates the rules our system must follow when making deletions. Since STORE is a noun we prune away all of the ATTR siblings




Fig. 3. The DsyntS tree for *The man ran to the big store.*

that modify it. By noticing that it is part of a prepositional phrase, we are able to delete STORE and use TO as our question, generating *The man ran where?*.

Content Extrapolation: We make use of the deep underlying story representation and the actor database to make inferences explicit that are not actually part of the original discourse. For example, the actor database tracks aspects of a character’s state. By using known antonyms of the adjective defining the current state, we can insert content for state changes, i.e. the alteration from *the fox is happy* to *now, the fox is sad*, where the fox is the actor and happiness is one of his states. This also allows us to introduce new dialogic interactions by having one speaker ask the other about the state of an actor, or make one speaker say something incorrect which allows the second speaker to contradict them: *The fox was hungry* followed by *No, he wasn’t hungry, he was just greedy*.

Pragmatic Markers: We can also insert pragmatic markers and tag questions as described in Table 1. Particular syntactic constraints are specified for each pragmatic marker that controls whether the marker can be inserted at all [16]. The frequency and type of insertions are controlled by values in the input parameter file. Some parameters are grouped by default into sets that allow them to be used interchangeably, such as downtoners or emphasizees. To provide us more control over the variability of generated variants, specific markers which are by default unrelated can be packaged together and share a distributed frequency limit. Due to their simplistic nature and low number of constraints, pragmatic markers prove to be a reliable source of variation in the systems output.

Lexical Choice: We can also replace a word with one of its synonyms. This can be driven simply by a desire for variability, or by lexical choice parameters such as word frequency or word length.

Morphosyntactic Postprocessing: The final postprocessing phase forms contractions and possessives and corrects known grammatical errors.

The results of the M2D processor are then given as input to RealPro [12], an off-the-shelf surface text realizer. RealPro is responsible for enforcing English grammar rules, morphology, correct punctuation, and inserting functional words in order to produce natural and grammatical utterances.

4 Evaluation Experiments

We assume H1 on the basis of previous experimental work, and test H2 and H3. Our experiments aim to: (1) establish whether and to what degree the M2D engine produces natural dialogs; (2) determine how the use of different parameters affect the user’s engagement with the story and the user’s perceptions of the naturalness of the dialog; and (3) test whether users perceive personality differences that are generated using personality models inspired by previous work. All experimental participants are pre-qualified Amazon Mechanical Turkers to guarantee that they provide detailed and thoughtful comments.

We test users’ perceptions of naturalness and engagement using two stories: the Garden story (Figure 1) and the Squirrel story (Figure 4). For each story, we generate **three** different dialogic versions with varying features:

M2D-EST renders the output from EST as a dialog by allocating the content equally to the two speakers. **No variations** of sentences are introduced.

M2D-BASIC consists of transformations required to produce a minimally natural dialog. First we apply **pronominalization** to replace nouns with their pronominal forms when telling the story in the third person. We then manipulate **sentence length** by breaking very long sentences into shorter ones, or by combining repetitious short sentences into one sentence. This is motivated by the fact that utterances in dialog tend to be less formal and use less complex syntactic structures [5]. The last transformation is **morphosyntactic postprocessing** as described in Section 3.

M2D-CHATTY adds interactive features to M2D-BASIC such as the insertion of **pragmatic markers** (acknowledgements, disfluencies, hedges) and question-answer generation (Table 1).

Squirrel Story: Monolog generated by EST
I placed the steely bowl on the deck in order for Benjamin to drink the bowl's water. The steely bowl was popular. The birds drank the bowl's water. The birds bathed themselves in the steely bowl. The birds organized themselves on the deck's railing in order for the birds to wait. The squirrels drank the bowl's water. The squirrel approached the steely bowl. The crazy squirrel was startled because the squirrel saw the squirrel's reflection. The crazy squirrel leaped because the squirrel was startled. The crazy squirrel fell over the deck's railing because the squirrel leaped because the squirrel was startled. The crazy squirrel held the deck's railing with the squirrel's paw. The squirrel's paw slipped off the deck's railing. The crazy squirrel fell.
Squirrel Story: Dialog generated by M2D
S1: I placed the steely bowl on the deck. Benjamin drinks its water. S2: The steely bowl was very popular, technically. Typical. The birds drank its water. They bathed themselves in the steely bowl! The birds organized themselves on the deck's railing in order for them to wait. S1: Yeah, the birds organized themselves on the deck's railing in order for them to wait. The squirrels drank it's water. The squirrel approached the steely bowl. S2: Technically, it was startled. I am delighted to say that the squirrel saw its reflection. It jumped, actually and was startled. Basically, the squirrel was literally startled and fell over the deck's railing. It leaped because it was startled. S1: The crazy squirrel held the deck's railing with its paw. S2: Its paw slipped off its railing pal. The particularly crazy squirrel fell mate.

Fig. 4. Squirrel Story: Monolog/Dialog Generation.

Each pairwise comparison is a Human Intelligence Task (HIT; a question that needs an answer), yielding 6 different HITs. We used 5 annotators (Turkers) per HIT to rate the levels of engagement/naturalness on a scale of 1-5, followed by detailed comments justifying their ratings.

We create several subsets of features that work well together and recursively apply random feature insertion to create many different output generations. These subsets include the types of questions that can be asked, different speaker interactions, content polarity with repetition options, pragmatic markers, and lexical choice options. A restriction is imposed on each of the subgroups, indicating the maximum number of parameters that can be enabled from the associated subgroup. This results in different styles of speaker depending on which subset of features is chosen. A speaker who has a high number of questions along with hedge pragmatic markers will seem more inquisitive, while a speaker who just repeats what the other speaker says may appear to have less credibility than the other speaker. We plan to explore particular feature groupings in future work to identify specific dialogic features that create a strong perception of personality.

4.1 M2D-EST vs. -Basic vs. -Chatty

The perceptions of engagement given different versions of the dialogic story is shown in Figure 5. A paired t-test comparing M2D-CHATTY to M2D-EST shows that increasing the number of appropriate features makes the dialog more engaging ($p = .04$, $df = 9$). However there are no statistically significant differences between M2D-BASIC and M2D-EST, or between M2D-BASIC and M2D-CHATTY. Comments by Turkers suggest that the M2D-CHATTY speakers have more personality because they use many different pragmatic markers, such as questions and other dialogically oriented features.

The perception of naturalness across the same set of dialogic stories is shown in Figure 6. It shows that M2D-BASIC was rated higher than M2D-EST, and their paired t-test shows that M2D-BASIC (inclusion of pronouns and agg- and deaggregation) has a positive impact on the naturalness of a dialog ($p = .0016$, $df = 8$). On the other hand, M2D-BASIC is preferred over M2D-CHATTY, where the use of pragmatic markers in M2D-CHATTY was often noted as unnatural.




Fig. 5. Mean scores for engagement.




Fig. 6. Mean scores for naturalness.

4.2 Personality Models

A second experiment creates a version of M2D called M2D-PERSONALITY which tests whether users perceive the personality that M2D-PERSONALITY intends to manifest. We use 4 different stories from the PersonaBank corpus [13] and create introverted and extroverted personality models, partly by drawing on features from previous work on generating personality [16].

We use a number of new dialogic features in our personality models that increase the level of interactivity and entrainment, such as asking the other speaker questions or entraining on their vocabulary by repeating things that they have said. Content allocation is also controlled by the personality of the speaker, so that extraverted agents get to tell more of the content than introverted agents.

We generate two different versions of each dialog, an extroverted and an introverted speaker (Table 2). Each dialog also has one speaker who uses a default personality model, neither strongly introverted or extraverted. This allows us to test whether the perception of the default personality model changes depending on the personality of the other speaker. We again created HITs for Mechanical Turk for each variation. The Turkers are asked to indicate which personality best describes the speaker from among extroverted, introverted, or none, and then explain their choices with detailed comments. The results are shown in Table 3, where Turkers correctly identified the personality that M2D-PERSONALITY aimed to manifest 88% of the time.

Table 3. Personality Judgments

	Extro	Intro	None
Extro	16	0	0
Intro	3	12	1

Turkers’ comments noted the differential use of pragmatic markers, content allocation, asking questions, and vocabulary and punctuation. The **extroverted** character was viewed as more dominant, engaging, excited, and confident. These traits were tied to the features used: exclamation marks, questions asked, exchanges between speakers, and pragmatic markers (e.g., *basically, actually*).

The **introverted** character was generally timid, hesitant, and keeps their thoughts to themselves. Turkers noticed that the introverted speaker was allocated less content, the tendency to repeat what has already been said, and the use of different pragmatic markers (e.g. *kind of, I guess, Mhmm, Err...*).

Table 4 shows Turker judgements for the speaker in each dialog who had a default personality model. In 53% of the trials, our participants picked a personality other than “none” for the speaker that had the default personality. Moreover, in 88% of these incorrect assignments, the personality assigned to the speaker was the opposite of the personality model assigned to the other speaker. These results imply that when multiple speakers are in a conversation, judgements of personality are **relative** to the other speaker. For example, an introvert seems more introverted in the presence of an extravert, or a default personality may seem introverted in the presence of an extravert.

5 Discussion and Future Work

We hypothesize that dialogic storytelling may produce more engagement in the listener, and that the capability to render a story as dialog will have many practical applications (e.g. with gestures [10]). We also hypothesize that expressing personality in storytelling will be useful and show how it is possible to do this in the experiments presented here. We described an initial system that can translate a monologic deep syntactic structure into many different dialogic renderings.

Table 4. Default Personality Judgments

	Extro	Intro	None
Other is Extro	1	7	8
Other is Intro	8	1	7

Table 2. Feature Frequency for Extra. vs. Intro.

Not all lexical instantiations of a feature are listed.

Parameter	Extra.	Intro.
Content allocation		
Content density	high	low
Pragmatic markers		
Adjective softeners	low	high
Exclamation	high	low
Tag Questions	high	low
Acknowledgments: <i>Yeah, oh God</i>	high	low
Acknowledgments: <i>I see, well, right</i>	low	high
Downtoners: <i>Sort of, rather, quite, pretty</i>	low	high
Downtoners: <i>Like</i>	high	low
Uncertainty: <i>I guess, I think, I suppose</i>	low	high
Filled pauses: <i>Err..., Mhmm...</i>	low	high
Emphasizers: <i>Really, basically, technically</i>	high	low
In-group Markers: <i>Buddy, pal</i>	high	low
Content elaboration		
Questions: Ask & let other spkr answer	high	low
Questions: Rhetorical, request confirmation	low	high
Paraphrase	high	low
Repetition	low	high
Interactions: Affirm adjective	high	low
Interactions: Corrections	high	low
Lexical choice		
Vocabulary size	high	low
Word length	high	low

We evaluated different versions of our M2D system. The results indicate that the perceived levels of engagement for a dialogic storytelling increase proportionally with the density of interactive features. Turkers commented that the use of pragmatic markers, proper pronominalization, questions, and other interactions between speakers added personality to the dialog, making it more engaging. In a second experiment, we directly test whether Turkers perceive that different speaker’s personalities in dialog. We compared introvert, extrovert, and a speaker with a default personality model. The results show that in 88% of cases the reader correctly identified the personality model assigned to the speaker. The results show that the content density assigned to each speaker as well as the choice of pragmatic markers are strong indicators of the personality. Pragmatic markers that most emphasize speech, or attempt to engage the other speaker are associated with extroverts, while softeners and disfluencies are associated with introverts. Other interactions such as correcting false statements and asking questions also contribute to the perception of the extroverted personality.

In addition, the perceived personality of the default personality speaker was affected by the personality of the other speaker. The default personality speaker was classified as having a personality 53% of the time. In 88% of these misclassifications, the personality assigned to the speaker was the opposite of the other speaker, suggesting that personality perception is relative in context.

While this experiment focused only on extrovert and introvert, our framework contains other Big-Five personality models that can be explored in the future. We plan to investigate: 1) the effect of varying feature density on the perception of a personality model, 2) how personality perception is relative in context, and 3) the interaction of particular types of content or dialog acts with perceptions of a storyteller’s character or personality. The pragmatic markers are seen as unnatural in some cases. We note that our system currently inserts them probabilistically but do not make intelligent decisions about using them in pragmatically appropriate situations. We plan to add this capability in the future. In addition we will explore new parameters that improves the naturalness and flow of the story.

Acknowledgments

We would like to thank Johnnie Chang and Diego Pedro for their roles as collaborators in the early inception of our system.

References

1. Gordon W. Allport and Henry S. Odbert. Trait names: a psycho-lexical study. *Psychological Monographs*, 47(1, Whole No. 211):171–220, 1936.
2. Elisabeth André, Thomas Rist, Susanne van Mulken, Martin Klesen, and Stephan Baldes. The automated design of believable dialogues for animated presentation teams. *Embodied conversational agents*, pages 220–255, 2000.
3. Mieke Bal and Eve Tavor. Notes on narrative embedding. *Poetics Today*, pages 41–59, 1981.

4. Janet B Bavelas, Linda Coates, and Trudy Johnson. Listeners as co-narrators. *Journal of personality and social psychology*, 79(6):941, 2000.
5. D. Biber. *Variation across speech and writing*. Cambridge Univ Pr, 1991.
6. Kevin Burton, Akshay Java, and Ian Soboroff. The icwsm 2009 spinn3r dataset. In *Proc. of the Annual Conf. on Weblogs and Social Media (ICWSM)*, 2009.
7. Pieter De Haan. More on the language of dialogue in fiction. *ICAME JOURNAL*, 20:23–40, 1996.
8. Gérard Genette. *Nouveau discours du récit*. Éd. du Seuil, 1983.
9. Lewis R. Goldberg. An alternative “description of personality”: The Big-Five factor structure. *Journal of Personality and Social Psychology*, 59:1216–1229, 1990.
10. Zhichao Hu, Michelle Dick, Chung-Ning Chang, Kevin Bowden, Michael Neff, Jean Fox Tree, and Marilyn Walker. A Corpus of Gesture-Annotated Dialogues for Monologue-to-Dialogue Generation from Personal Narratives. *Proc. of the Tenth International Conf. on Language Resources and Evaluation (LREC 2016)*, 2016.
11. Zhichao Hu, Marilyn Walker, Michael Neff, and Jean E. Fox Tree. Storytelling agents with personality and adaptivity. *Intelligent Virtual Agents*, page 30, 2015.
12. Benoit Lavoie and Owen Rambow. A fast and portable realizer for text generation systems. In *Proc. of the Third Conf. on Applied Natural Language Processing, ANLP97*, pages 265–268, 1997.
13. Stephanie Lukin, Kevin Bowden, Casey Barackman, and Marilyn Walker. Personabank: A corpus of personal narratives and their story intention graphs.
14. Stephanie Lukin and Marilyn Walker. Narrative variations in a virtual storyteller. *Intelligent Virtual Agents*, page 30, 2015.
15. Stephanie M Lukin, Lena I Reed, and Marilyn Walker. Generating sentence planning variations for story telling. In *16th Annual Meeting of the Special Interest Group on Discourse and Dialogue*, page 188, 2015.
16. F. Mairesse and M.A. Walker. Towards personality-based user adaptation: psychologically informed stylistic language generation. *User Modeling and User-Adapted Interaction*, pages 1–52, 2010.
17. D.P. McAdams, N.A. Anyidoho, C. Brown, Y.T. Huang, B. Kaplan, and M.A. Machado. Traits and stories: Links between dispositional and narrative features of personality. *Journal of Personality*, 72(4):761–784, 2004.
18. J Meehan. *The metanovel: writing stories by computer*, pages 197–258. 1976.
19. Igor A. Melčuk. *Dependency Syntax: Theory and Practice*. SUNY, Albany, New York, 1988.
20. Roxana Moreno, Richard Mayer, and James Lester. Life-like pedagogical agents in constructivist multimedia environments: Cognitive consequences of their interaction. In *World Conf. on Educational Media and Technology*, number 1, pages 776–781, 2000.
21. M. Neff, Y. Wang, R. Abbott, and M. Walker. Evaluating the effect of gesture and language on personality perception in conversational agents. In *Intelligent Virtual Agents*, pages 222–235. Springer, 2010.
22. Monisha Pasupathi and Timothy Hoyt. The development of narrative identity in late adolescence and emergent adulthood: The continued importance of listeners. *Developmental Psychology*, 45(2):558–574, 2009.
23. Federico Peinado and Pablo Gervás. Evaluation of automatic generation of basic stories. *New Generation Computing, Special issue: Computational Creativity*, 24(3):289–302, 2006.
24. James W Pennebaker and Janel D Seagal. Forming a story: The health benefits of narrative. *Journal of clinical psychology*, 55(10):1243–1254, 1999.

25. P. Piwek and S. Stoyanchev. Generating expository dialogue from monologue: motivation, corpus and preliminary rules. In *Human Language Technologies: The 2010 Annual Conf. of the North American Chapter of the Association for Computational Linguistics*, pages 333–336. Association for Computational Linguistics, 2010.
26. Paul Piwek, Hugo Hernault, Helmut Prendinger, and Mitsuru Ishizuka. T2d: Generating dialogues between virtual agents automatically from text. In *Intelligent Virtual Agents*, pages 161–174. Springer, 2007.
27. Paul Piwek and Svetlana Stoyanchev. Data-oriented monologue-to-dialogue generation. In *Proc. of the 49th Annual Meeting of the Association for Computational Linguistics: Human Language Technologies: short papers-Volume 2*, pages 242–247. Association for Computational Linguistics, 2011.
28. Gerald Prince. *A Grammar of Stories: An Introduction*. Number 13. Walter de Gruyter, 1973.
29. Vladimir Propp. *Morphology of the Folk Tale*. University of Texas Press, second edition, 1969. Trans. Laurence Scott. Originally published 1928.
30. D.N. Rapp, R.J. Gerrig, and D.A. Prentice. Readers’ Trait-Based Models of Characters in Narrative Comprehension. *Journal of Memory and Language*, 45(4):737–750, 2001.
31. Elena Rishes, Stephanie M Lukin, David K Elson, and Marilyn A Walker. Generating different story tellings from semantic representations of narrative. In *Interactive Storytelling*, pages 192–204. Springer International Publishing, 2013.
32. L. Ross. The intuitive psychologist and his shortcomings: Distortions in the attribution process. *Advances in experimental social psychology*, 10:173–220, 1977.
33. Kimiko Ryokai, Catherine Vaucelle, and Justine Cassell. Virtual peers as partners in storytelling and literacy learning. *Journal of Computer Assisted Learning*, 19(2):195–208, 2003.
34. Deborah Schiffrin. Tense variation in narrative. *Language*, pages 45–62, 1981.
35. Nicole K Speer, Jeremy R Reynolds, Khena M Swallow, and Jeffrey M Zacks. Reading stories activates neural representations of visual and motor experiences. *Psychological Science*, 20(8):989–999, 2009.
36. S. Stoyanchev and P. Piwek. Harvesting re-usable high-level rules for expository dialogue generation. In *Proc. of the 6th International Natural Language Generation Conf.*, pages 145–154. Association for Computational Linguistics, 2010.
37. Deborah Tannen. *Talking voices : repetition, dialogue, and imagery in conversational discourse*. CUP, 1989.
38. A. Thorne. The press of personality: A study of conversations between introverts and extraverts. *Journal of Personality and Social Psychology*, 53(4):718, 1987.
39. Avril Thorne, Neill Korobov, and Elizabeth M Morgan. Channeling identity: A study of storytelling in conversations between introverted and extraverted friends. *Journal of research in personality*, 41(5):1008–1031, 2007.
40. Avril Thorne, Kate C McLean, and Amy M Lawrence. When remembering is not enough: Reflecting on self-defining memories in late adolescence. *Journal of personality*, 72(3):513–542, 2004.
41. Kees Van Deemter, Brigitte Krenn, Paul Piwek, Martin Klesen, Marc Schröder, and Stefan Baumann. Fully generated scripted dialogue for embodied agents. *Artificial Intelligence*, 172(10):1219–1244, 2008.
42. E. Wade and H.H. Clark. Reproduction and demonstration in quotations. *Journal of Memory and Language*, 32(6):805 – 819, 1993.
43. Bo Yao and Christoph Scheepers. Contextual modulation of reading rate for direct versus indirect speech quotations. *Cognition*, 121(3):447–453, 2011.