Commonsense Knowledge: Papers from the AAAI Fall Symposium (FS-10-02)

A Commonsense Knowledge Base for Generating Children's Stories

Ethel ChuaJoy Ong

Center for Language Technologies
College of Computer Studies
De La Salle University, Manila, 1004 Philippines
ethel.ong@delasalle.ph

Abstract

This paper presents our work in developing a commonsense knowledge source based on semantic concepts about objects, activities and their relationships in a child's daily life. This commonsense ontology is then used by our automatic story generator to output children's stories of the fable form from a given input picture. The generated story is a narration of the events of a basic plot that flows from negative to positive (rule violation to value acquisition), using themes that are familiar to children. The paper ends with descriptions of further investigations that are underway to extend the system, including using a formal upper ontology to represent storytelling knowledge, and the generation of stories from a given set of sequential scenes.

Introduction

People use storytelling as a means of sharing information about the events we experience in our daily lives. During this interchange, we are able to understand each other because we share a large body of commonsense knowledge about things, concepts and their interrelationships. In order to develop computer systems that can interact naturally with their human users, a similar collection of knowledge must be made available to them.

Researchers in the field of knowledge engineering have conducted various projects to collect and represent the knowledge in a form that can be utilized further by computer systems. ConceptNet (Liu and Singh 2004a) is one such representation of lexical and commonsense knowledge that follows a frame-based entity-relation model. It consists of three types of nodes – noun phrases to represent objects such as "*doll*" and "*lamp*", attributes or modifiers such as "*pretty*" and "*fragile*", and activity or verb phrases such as "*play*" and "*break*", with binary semantic relations connecting a pair of these nodes. The relations were extracted by the Open Mind Commonsense (OMCS) project (Liu and Singh 2004b) from commonsense concepts.

Copyright © 2010, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

Several applications have harnessed the collected information. MakeBelieve (Liu and Singh 2002) employed an interactive agent to perform logical reasoning on the commonsense knowledge extracted by the OMCS project to generate short fictions from an initial seed story. Picture Books (Hong et al 2008) is another automatic story generator that has adopted the design of ConceptNet to model concepts relevant for its set of themes reflecting daily activities of young children, such as *going to school* and *playing in the park*. Stories promoting good moral values, like *learning to share* and *being obedient*, are then generated based on an input picture containing a set of picture elements (background, characters, and objects) that have been selected by the young readers.

TPEG (Hong and Ong 2009), a pun generator, used ConceptNet to identify semantic relations that exist between words in an input pun and stores these knowledge in templates for subsequent use by the generator module to automatically produce computer puns with syntactic and semantic structures closely resembling the source human pun. FiLearn (Bayani, Palma and San Jose 2009), a learning environment on Filipino, patterned the design of its knowledge source to ConceptNet to model semantic relations between Filipino words in order to generate multiple choice questions with appropriate distractors from question templates. The questions form part of the exercises taken by primary-level children who are learning the Filipino language.

Our discussion in this paper focuses on the commonsense knowledge and the methodology involved in its development in order to support the requirements of Picture Books' story planner in its story generation task. Results from the manual evaluation of human judges on the generated stories are also presented. The paper ends with a discussion of ongoing efforts to extend Picture Books to generate stories from multiple input pictures, as well as the use of the Suggested Upper Merged Ontology or SUMO (Cua et al. 2010) to represent storytelling knowledge and the story plan.

Storytelling Knowledge

Picture Books generates fable-form stories with familiar animals as the main characters, and the story plot flows from negative to positive where the child violates a rule, experiences the consequence, and eventually learns the moral lesson. Below is an example story generated by the system for 4-year-olds, entitled "*Denise the dog learns to take a bath*".

> The day was sunny. Denise the dog was in the bedroom. She played with ball. Mommy Debbie told Denise to take bath. Denise did not want to take bath. She continued to play. Denise did not take the bath. She became dirty. Denise felt itchy. She felt hurt. Denise cried. Mommy Debbie saw that Denise was crying. She told Denise to take bath. Denise wanted to take bath. She took the bath with a soap. Denise soothed itchy skin. Taking bath results to smelling nice. It was fun. Denise was happy. From that day onwards, she always took the bath.

Notice in the story that various elements are needed to narrate a coherent story. These include those found in the picture that was provided by the child through the *Picture Editor* facility, namely the background (*bedroom*), the character (*Denise and her mommy*), and the objects (*ball*, *soap*). From these explicitly specified elements, the system randomly selects a setting (*sunny day*) and assigns a theme. Utilizing a semantic ontology, the other elements of the story are then determined, namely the activities (*playing*, *taking a bath*, *crying*), the emotions (*hurt*, *happy*), the sequence of events, and the character interactions.

According to Theune and her colleagues (2006), there are five levels of a story that must be represented in the semantic network, which are the story world knowledge, character representations, a causal and temporal network to represent plot structures, a representational model of narratological concepts, and the representation of the story's potential effects on the user. The story world, according to Swartjes (2006), is where the story takes place, and includes characters that interact in this world as well as the associated objects.

Of the five levels, only the first four are considered in the current implementation of Picture Books. Moreover, the four levels are differentiated into two types of knowledge, called the operational knowledge and the domain knowledge. The operational knowledge includes the themes, the characters, and the planning operators, while the domain knowledge represent the commonsense knowledge about the story world and the causal network of actions and events that can take place in the world.

Operational Knowledge

The operational knowledge contains theme-based planning operators comprised of author goals, character goals, and character representations. Author goals (Solis et al 2009 and Cua et al 2010), based on Minstrel (Turner 1992), are high-level tasks focusing on the narrative structure of the story to be generated. Since Picture Books follows the plot structure proposed by Machado (2003) consisting of problem, rising action, solution and climax, each of these subplots is represented by two author goals covering the major scenes of a theme. Each author goal has a main goal and a corresponding consequence. An author goal is in turn comprised of two or more character goals to represent the actions a character does to depict the main goal of the scene and the consequence of doing the goal of the scene.

A character goal represents an action that is performed by a character or an event that has occurred in the story. It has five attributes based on the action operators of Uijlings (2006) – the *action* to be performed, the *agens* representing the doer of the action, the *patiens* representing the *character* that receives the action, the *target* which is the object of the verb, and the *instrument* that is used to perform the action. This design allows a character goal to be easily translated to a declarative sentence in active voice in the resulting story using the surface realizer, simpleNLG (Venour and Reiter 2008).

Aside from the predefined character goals, the operational knowledge also includes a set of semantic rules to map dynamic character goals to action operators. Dynamic character goals are created by the planner depending on the semantic relations retrieved from the commonsense knowledge base and are used to provide more details to produce stories of varying lengths depending on the ages of the readers. This is further discussed later in the section on "Using the Storytelling Knowledge in Planning".

The characters comprise the last component of the operational knowledge. There are currently ten child characters, and for every child character, there are two corresponding adult (parent) characters. Every character has an assigned gender, a type (adult or child), and a name. The first child character placed in the picture is the main character; succeeding child characters play tertiary roles and may be included in the resulting story depending on the theme. For example, in the story below entitled "*Porky the pig learns to share*", a second child character, *Robbie the rabbit*, is needed by the theme.

> The morning was sunny. Porky the pig was at the playground. He played. Mommy Patricia told Porky to share. He did not want to share. Porky continued to play. Robbie the rabbit insisted Porky to share. Porky fought with Robbie. Robbie felt hurt. He cried. Robbie told Mommy Patricia that Porky did not want to share. She told Porky to share. Not sharing is bad. Porky felt guilty. He apologized to Robbie. Porky shared with Robbie. Sharing is fun. After that day, Porky always shared.

The adult character plays a secondary role and is always the one imparting the rule to the main character at the start of the story. The adult character is also the one providing support as the main character undergoes transformation from rule violation to experiencing the consequence and learning of the lesson.

Domain Knowledge

The domain knowledge is composed of concepts that are connected by semantic relations. The semantic relations are patterned after those defined in ConceptNet, but the actual contents have been derived in order to provide concepts that are not only familiar to children but are also relevant to the system's themes promoting moral values like sharing, honesty, and obedience. This knowledge source is utilized by the story planner as a semantic network of constraints to generate the various elements of the stories.

Story world knowledge represents concepts that involve objects and their attributes, and the various locations describing the setting where the story can take place. Causal structures represent the sequence of actions that characters can perform, events that can occur as a result of performing some actions, and emotions that characters can feel as a reaction to some events.

Objects. Objects represent things that may exist in the story world and which the character may use to pursue his/her goal. The *isA* relation is used to define classes of objects, for example, *isA(ball, toy)*. Objects are described through their components, *partOf(wheel, truck)*; properties, *property(lamp, fragile)*; composition, *madeOf(bottle, plastic)*; or purpose, *usedFor(toy, play)*.

Locations. Objects are associated to the location where they can be found and their co-located objects, for example, *locationOf(book, book store)*, *locationOf(swing, park)*, and *oftenNear(swing, slide)*.

Concepts. Abstract concepts, such as punishment, problem, and discomfort, are also modeled, for example, *isA(grounded, punishment)*, *isA(itchy, discomfort)*, and *isA(fight, problem)*. Complementary concepts are modeled using the *negate* relation, e.g., *negate(sleep early, sleep late)*, and *negate(eat healthy food, eat junk food)*.

Events. Events occur in the story world as a result of some character actions (explicit events) or as a naturally occurring phenomenon (implicit events). Picture Books currently covers only explicit events, which are represented as a series of cause and effect chain, and the storytelling task involves organizing these events representing a child's daily activities to share with others.

Predecessor and successor events are modeled using the relations *firstSubeventOf* and *lastSubeventOf*, e.g., *firstSubeventOf(itch, scratch)*, *lastSubeventOf(hurt, cry)*.

An event may lead to a target *goal event*, e.g., *eventForGoalEvent(clean up, be neat)*, or to a target *goal state*, e.g., *eventForGoalState(listen, understand)*.

Actions. Actions are activities that a character performs in order to reach a desired goal or state, or as a response to the occurrence of some events. The *capableOf* relation is used to model actions that can be performed by a character on an object, e.g., *capableOf(toy, play)*.

The *effectOf* and *effectOfIsState* relations are used to model the consequence of performing some actions, with the former requiring a verb phrase and the latter requiring a noun phrase for the second concept, e.g., *effectOf(become dirty, feel itchy)*, *effectOf(break object, be scared)*, and *effectOfIsState(irritated skin, discomfort)*.

Actions may require the use of some objects, e.g., *eventRequiresObject(eat, spoon)*.

Character Emotions. Characters may feel emotions in response to certain events. The *isA* relation is again used to represent various emotions, e.g., *isA(happy, emotion)* and *isA(scared, emotion)*. These are then associated to actions and events whose performance or occurrence will trigger the emotion, such as *eventForGoalState(play games, happy), effectOf(break object, be scared), effectOf(search, feel worried)*.

Using the Storytelling Knowledge in Planning

A character goal is the smallest unit in Picture Books that represents a specific action to be performed or an actual event that has taken place. It is represented as a five-attribute-value pairs of the form *CGid(Action:<verb>, Agens:<character>, Patiens:<character>, Target: <object>, Instrument:<object>)*. For example, the character goal *CG001(Action:"tell", Agens:%adult%, Patiens:%child%, Target: %lesson%, Instrument:null)* corresponds to "*an adult character telling the child character about a lesson*", where *%lesson%* denotes the theme of the story. The *patiens*, *target*, and *instrument* attributes are optional and may be dropped if there are no corresponding values. Keywords enclosed within two percentage symbols (%) denote derivable values that are assigned when the character goal is instantiated. The derivable values are from the input picture (location, character, object) or from the assigned theme (lesson).

A character goal can contain an inner character goal in its *Target* attribute to generate sentences with clauses. For example, the character goal *CG006(Action:"inform", Agens:%child%, Patiens:%adult%, Target:CG003 (Action:"want", Agens:%child%, Target:%object%))* corresponds to "*child character informs adult character that he wants some object*".

Queries can also be assigned as values to attributes of a character goal. A query is of the form *onto<Category> (<concept1>[, <concept2>])*, where *<Category>* refers to the semantic category (defined by ConceptNet) that classifies the various relations in the ontology, such as *Things*, *Spatial*, and *Event*. These categories place a constraint on the types of relations that will be searched for by the story planner, in order to filter irrelevant concepts and paths as the ontology is being traversed. *<concept1>* and *<concept2>* are specific words, derivable values, or another query.

A query requires searching the ontology for matching relations. For example, in the character goal *CG048(Action:"tell", Agens:%adult%, Patiens:%child%, Target:CG016(Action:"eat", Agens:%child%, Patiens: %ontoThings("food")%, Instrument:%ontoThings ("eat")%))*, two queries are present – *%ontoThings ("food")%* which searches the ontology for conceptual relations in the *Things* category that are related to *food*, and *%ontoThings("eat")%* which searches for concepts related to *eat*. The resulting values, *vegetables* and *spoon*, respectively, will then lead to the generation of the sentence "*adult character told the main character that he should eat his vegetables using a spoon*".

Search for Paths of Relations

A query in a character goal may necessitate two types of search operations – search for paths of relations and search for related concepts. The search for paths of relations is used when a relationship between two given concepts (source and destination) can be described through a series of connected relations. As the search operation proceeds, an ontology tree is constructed to store the relations that are already visited; and it continues until the destination concept is located and all concepts in the same tree level with the destination concept are in the ontology tree. This approach enables an ontology tree to contain more than one applicable path.

The search for paths of relations is often used by the story planner to relate two subplots, i.e., from initial setting to problem. The source concept can either be the opposite of the lesson that has been imparted by the adult to the child or an activity that the child is currently performing, and the destination concept is a problem or a discomfort.

Consider the query *%ontoAction(not take bath, discomfort)%* which denotes a search for a series of semantic relationships connecting *not take bath* (which is contrary to the adult's command) to *discomfort*. The resulting ontology tree contains the following path:

> not take bath =capableOf=> become dirty =effectOf=> feel itchy =lastSubeventOf=> scratch =effectOfIsState=> irritated skin =effectOfIsState=> discomfort

A set of semantic relation rules is then used to map the retrieved relations to dynamic character goals. These rules specify how the attributes of the character goal will be filled with the appropriate concepts from the resulting relations. Table 1 lists some of the semantic rules for the relations in the resulting ontology tree.

Relation	Action	Agens	Target
capableOf	Verb of concept2	Previous Agens	Complement of concept2
effectOf	Verb of concept2	Previous Agens	Complement of concept2
effectOfIsState	"cause"	Concept1	Concept2
lastSubeventOf	C2	Previous Agens	

Table 1. Sample semantic rules

Dynamic character goals are created at runtime depending on the number of relations retrieved from the knowledge source as well as the age of the user. These character goals increase the length of and also provide variances in the events that may occur in the story.

From the ontology tree derived in the sample above, the dynamic character goals will have the following forms when transformed to surface text.

> *%child% became dirty. %child% felt itchy. %child% scratched. Scratching caused irritated skin. Irritated skin caused discomfort.*

Currently, one relation maps to one character goal (or one sentence in the output story), but ongoing work is investigating the application of the Rhetorical Structure Theory of Mann and Thompson (1988) to generate a complex sentence from two or more character goals with the use of explicit discourse markers.

A pronoun generator in the latter stages of realization handles identifying the pronoun to be used at the appropriate places in the sentences before sending these to the surface realizer, simpleNLG.

Search for Related Concepts

The search for related concepts is used to locate another concept that has a direct relationship with the given concept. It is further subdivided into two, *searchConceptViaCategory* searches for a related concept given the word and the semantic category. For example, the query *%ontoThings("food")%* searches for all concepts having a semantic relation with *food* under the *Things* category, and may return any of the following results – *isA(apple, food)*, *isA(candy, food)*, *property(apple, delicious)*, and *property(candy, sweet)*.

searchConceptViaRelationship, on the other hand, searches for another concept that has a specified semantic relation with the given concept. For example, the query *%propertyOf(day)%* searches for all *propertyOf* relations with *day* as one of its concepts, and may return any of the following results – *propertyOf(day, sunny)*, *propertyOf (day, warm)*, and *propertyOf(day, windy)*. Randomly choosing any of the resulting relations can lead to different settings for the introductory part of the story, e.g., "*The day is sunny.*", "*The day is warm.*", or "*The day is windy.*"

Another usage of *searchConceptViaRelationship* is in generating object descriptions. During the initial implementation of Picture Books, objects were not described in the story, e.g., "*Roy the rabbit played near a lamp.*", resulting in low evaluation scores given by the linguists. To address this issue, *property* relations for objects were added to the ontology and are used to generate sentences of the forms, "*He played near a fragile lamp.*", as well as "*Apple is delicious.*".

Evaluation of the Generated Story

Consultations with child educators led to the identification of 9 backgrounds, 11 themes and 37 objects. Target stories depicting the themes were handcrafted, and semantic relations present in these stories were manually extracted and stored into the ontology, which currently contains 240 concepts and 369 semantic relations.

15 stories (5 for each age group) were generated for evaluation. No automatic comparison between the generated stories was performed. Instead, human judges manually read the stories and rated each from 1 to 4 based on the degree to which the individual story satisfied the property under consideration. There were two major categories – linguistic quality and coherency. Linguistic quality specifies how well the story text is written and focuses on correct syntax and pronoun usage, while coherency specifies how well the sequence of events are chained together, if the actions of the characters are believable, and the presence of object descriptions. Table 2 shows the evaluation result for coherency, which is largely dictated by the available knowledge in the ontology.

Criteria	Average
Sentences are coherent	2.67
Objects in the story were described	2.80
The story has transition	3.47
Believable character actions/responses	3.93
General Average	**3.22**

Table 2. Evaluation of the Semantic Ontology

The coherency of the sentences received a low average score of 2.67 because some stories have missing actions. In the excerpt below, the action of the child (underlined words) depicting what she did to show bravery was missing. This simply requires populating the ontology with additional world knowledge, such as the set of possible actions a character can do to achieve a desired goal, e.g., *eventForGoalState(introduce, friendship).*

She wanted to be brave. Ellen was brave. She wanted to play with others. <u>She bravely introduced herself.</u> Ellen made friends.

The presence of object descriptions also initially received a low average score of 2.80 due to two reasons – missing objects and missing properties of objects. In the example below, the object that *Porky* is playing with and being asked to share with *Robbie* was not stated explicitly.

The morning was sunny. Porky the pig was at the playground. He played. Mommy Patricia told Porky to share. He did not want to share. Porky continued to play. Robbie the rabbit insisted Porky to share.

And when objects do exist, no description is provided in the text, such as in the sentences "*Rizzy the rabbit played near a <u>lamp</u>.*", and "*Geena the giraffe saw a <u>doll</u>.*" This concern was easily remedied by adding relations denoting properties of objects into the ontology, e.g., *propertyOf(lamp, fragile)* and *propertyOf(doll, pretty).*

Although a trivial solution, this approach gave rise to another concern. Object descriptions can serve as a plot device to direct the story flow, e.g., a *fragile* lamp can lead to themes on honesty or being careful (when the lamp breaks through play), whereas an *expensive* lamp can lead to themes on the value of money or being thrifty. Representing this knowledge for use by the story planner is an open problem that we plan to address in the future.

The availability of operational knowledge in the form of predefined author goals and character goals serve to constrain the stories being generated. The ontology is used only to provide information needed to fill in the attributes in the character goals in a theme-driven story plot template. Thus, the standard system output falls into a progression of story events from rule violation to lesson acquisition, with the available resource providing varying instantiations of the same basic story plot. This approach led to an average score of 3.47 for the transition of story events with the presence of *action* and *event* relations in the semantic ontology representing the cause-effect chain of character actions and reactions. This coincides with the findings of Peinado and Gervas (2006) that "ontology-based story obtain good results on coherence because the ontology forces explicit links between events".

Although the system received a high average rating in the believable character actions criterion, the setting (time of day and weather conditions) do not have any effect on how the story progresses, resulting in the generation of stories with activities that take place at inappropriate setting, such as in the excerpt below:

The evening was warm. Ellen the elephant was at the school. She went with Mommy Edna to the school.

Related Ongoing Works

In this section, we present brief descriptions of related works that are being undertaken to extend Picture Books.

Picture Books 2

Picture Books 2 intends to generate stories for older kids (6-8 years old). Its story themes revolve around the child exploring the world and life's lessons on his own, outside the comforts of his home. The list of suported backgrounds includes the *camp*, *grocery*, *classroom*, and *street*. Adult (parent) characters are no longer required in every story and the *Picture Editor* facility now provides support for the child to define at least three pictures (scenes). This requires representing knowledge in the semantic ontology about concepts on character and object existence (appearance and disappearance) and movement across two adjacent scenes.

Another factor being investigated is character believability, which according to Riedl and Young (2004), is an essential property of narratives because "the events that occur in the story are motivated by the beliefs, desires and the goals of the characters". Children easily connect to characters in the story when they believe they have something in common with the characters. Although the current system made attempts to model character actions and emotions that are as realistic as possible, the character traits are part of the operational knowledge and are dictated by the theme, which is assigned based on the background and objects (of the input picture).

Child educators believe that embodying the story characters with traits (e.g., a dog is loyal, a fox is cunning, an elephant is studious, and a cat is mischievous) would help children to relate to the story better. Picture Books 2 is thus working on enhancing the character representations by providing individual traits to the characters and considers these as another factor in determining the theme.

SUMO Stories

We are also investigating the use of SUMO to represent the storytelling knowledge of Picture Books and the resulting fabula, and takes advantage of Sigma to infer facts and events about the story being created (Cua et al 2010). We hope to find out if using a more formal upper ontology based on first-order logic can produce better and more variety of stories, as well as transfer the search control from the story planner to the inference engine.

Conclusion

We have presented our story generator, Picture Books, its semantic ontology based on commonsense knowledge, as well as how this knowledge source is used by the story planner to generate children's stories. Evaluation results from human judges show that using an ontology produce coherent stories, which concur with the findings of Peinado and Gervas (2006).

Picture Books' ontology is extensible and allows the designer to add new concepts to increase the variation of possible stories that can be generated. As long as the new concepts use the standard semantic relations defined in ConceptNet and are connected with the existing concepts in the ontology, the coherence of the resulting story is guaranteed. Note further that the semantic relations rules allow new relation types to be used, as long as the mapping of the relations to dynamic character goals are also defined. More experiments are needed to validate these claims.

Finally, it would be significant to the future directions of this research if more experiments will be conducted to measure the *interest* property of the generated stories with the actual users, the children.

References

Bayani, G., Palma, D., and San Jose, R. 2009. *FiLearn: Filipino Language Learning Software for Kids*. Undergraduate Thesis, College of Computer Studies, De La Salle University, Manila, Philippines.

Cua, J., Manurung, R., Ong, E., and Pease, A. 2010. Representing Story Plans in SUMO. In *Proc. NAACL-HLT 2010 Workshop on Computational Approaches to Linguistic Creativity*, 40-48, Los Angeles, California: Association for Computational Linguistics.

Gervas, P. 2009. Computational Approaches to Storytelling and Creativity. *AI Magazine*, 30(3): 63-70, AAAI.

Hong, A.J., Siy, J.T., Solis, C.J., and Tabirao, E. 2008. *Picture Books: An Automated Story Generator*. Undergraduate Thesis, College of Computer Studies, De La Salle University, Manila, Philippines.

Hong, B.A., and Ong, E. 2009. Automatically Extracting Word Relationships as Templates for Pun Generation. In *Proc. NAACL-HLT 2009 Workshop on Computational Approaches to Linguistic Creativity*, 24-31, Boulder, Colorado: Association for Computational Linguistics.

Liu, H., and Singh, P. 2002. MakeBelieve: Using Commonsense Knowledge to Generate Stories. In *Proc. 18th National Conference on AI*, 957-958, Edmonton, Alberta: National Conference on Artificial Intelligence.

Liu, H., and Singh, P. 2004a. Commonsense Reasoning in and over Natural Language. In *Proc. 8th International Conference on Knowledge-Based Intelligent Information and Engineering Systems*, 293-306, Wellington, New Zealand: Springer Berlin.

Liu, H., and Singh, P. 2004b. ConceptNet — A Practical Commonsense Reasoning Tool-Kit. *BT Technology Journal*, 22(4): 211-226, Springer Netherlands.

Machado, J. 2003. Storytelling. *Early Childhood Experiences in Language Arts: Emerging Literacy*, 304-319. Clifton Park, New York: Thomson/Delmar Learning.

Mann, W. and Thompson, S. 1988. Rhetorical Structure Theory: Towards a Functional Theory of Text Organization. *TEXT*, 8(3): 243-281.

Peinado, F. and Gervás, P. 2006. Evaluation of Automatic Generation of Basic Stories. *New Generation Computing,* 24(3): 289-302, Ohmsha, Ltd. and Springer.

Riedl, M. O. and Young, R. M. 2004. An Intent-Driven Planner for Multi-Agent Story Generation. In *Proc. 3rd International Joint Conference on Autonomous Agents and Multi-Agent Systems*, 186-193. Washington, DC, USA: IEEE Computer Society.

Solis, C., Siy, J.T., Tabirao, E., and Ong, E. 2009. *Planning Author and Character Goals for Story Generation*. In *Proc. NAACL HLT 2009 Workshop on Computational Approaches to Linguistic Creativity*, 63-70, Boulder, Colorado: Association for Computational Linguistics.

Swartjes, I. 2006. *The Plot Thickens: Bringing Structure and Meaning into Automated Story Generation*. Master's Thesis, University of Twente, The Netherlands.

Theune, M., Nijholt, A., Oinonen, K., and Uijlings J. 2006. Designing a Story Database for Use in Automatic Story Generation. In *Proc. 5th International Conference Entertainment Computing*, Cambridge, UK. In *Lecture Notes in Computer Science*, 4161: 298-301, Heidelberg: Springer Berlin.

Turner, S.R. 1992. *Minstrel: A Computer Model of Creativity and Storytelling*. University of California Technical Report CSD920057.

Uijlings, J.R.R. 2006. *Designing a Virtual Environment for Story Generation*. Master's Thesis, University of Amsterdam, The Netherlands.

Venour, C., and Reiter, E. 2008. A Tutorial for Simplenlg. http://www.csd.abdn.ac.uk/~ereiter/simplenlg.