Generating Text Descriptions in the Alex Interactive Storytelling System using a Semantic Ontology

Rancelli Jane Roxas[1], Dolleen Lour Huang, Bruce Elmer Peralta,
Sashmir Yap, Ethel ChuaJoy Ong[2]
College of Computer Studies, De La Salle University
Manila, Philippines
rancelli.jane.roxas@gmail.com[1], ethel.ong@delasalle.ph[2]

ABSTRACT
Semantic ontologies of commonsense concepts have been used to provide the necessary knowledge needed by computer systems to perform various tasks, such as in automatic story generation, interactive scavenger hunt games, virtual museum and storytelling. This paper presents our work in the development of a semantic ontology to support the generation of narrative text in the Alex interactive storytelling system. The text generated involves simple textual description of the objects found in the virtual world, the current options that are presented to the user to move the story forward, and a simple text feedback about the choices made by the user.

Results from user evaluation showed that the semantic ontology can be used to model concepts representing the various elements of a story. These include characters, locations, objects and their attributes, as well as story events. However, the text generation process can still be improved specifically in the production of text that describes the importance of acquiring an item, keeping track of the progress of the user towards fulfilling the goal of the story, and detailing how the previous and current choices of the user affected the story flow.

Keywords
Semantic ontology, Interactive Storytelling, Text Generation

1. INTRODUCTION
Interactive storytelling involves the participation of the user in story development [10]. It is a story-based experience where the progression of story events is dependent on the combined decision making of the computer and the user. To support various branching of the storyline and the dynamic generation of story events, the computer must be given a sufficiently large repository of knowledge from which to generate stories from. But what kind of model can be used to represent this knowledge needed by the computer to support the generation of story text in an interactive storytelling environment?

Storytelling involves the use of commonsense knowledge for story development and reasoning. For example, in this one line in a story, "Yesterday, I did something wrong", the word "wrong" can trigger human storytellers to think of different causal events to continue with the story, such as "forgetting to do your homework" or "being late in class". For a computer to be able to make sense of this situation and use it to participate in collaborative storytelling with its human users, it must be given the same collection of commonsense knowledge relevant to the themes of the stories it is trying to engage the user in.

There are various mechanisms that can be utilized to create a story. In an interactive storytelling environment where user participation is of prime importance, the story text to be generated should revolve around three main areas - 1) the set of possible actions that a user may be allowed to perform in order to drive the story forward; 2) the set of objects present in the virtual world that the user may utilize as he/she moves around the story world; and 3) the consequences of his/her action to the state of the story world. In this paper, we highlight the representation of commonsense knowledge that is needed to generate descriptions of objects present in the virtual world, and the causal relations between two or more events (including events that arise as a result of a user action). This commonsense knowledge can be represented with the use of a semantic ontology.

An ontology is a form of knowledge representation that shows the relationship between concepts given a domain. Existing systems use an ontology for story generation, interactive scavenger hunt games, virtual museum, storytelling, and knowledge acquisition. Picture Books [6], MAKEBELIEVE [7], and OMAdventure [5] are some story-based applications that make use of a semantic ontology to enable their story planners to fulfill the task of generating some story text in a given domain.

The rest of this paper is organized as follows. Section 2 gives a review of related works, specifically the use of semantic ontology in story-based environments. Section 3 then presents the design of our ontology and how the interactive storytelling system, Alex, uses this to generate narrative text that describes objects and events in the virtual world. Section 4 contains a discussion of the evaluation results and findings. An analysis of the flexibility of the design to support other domains by simply extending and populating the ontology with the appropriate semantic relations and concepts is also presented in this section. The paper ends with a summary of our findings and recommendations for further work to enhance the emergent narrative and the interactive storytelling aspects of Alex.

2. RELATED WORKS
An ontology is used as a form of knowledge representation to show the relationship between concepts in a given domain [8]. Several story-based environments use a semantic ontology to model the knowledge needed by their planners to generate complete or partial story text.

Picture Books [6] is an automated story generation system that adapted the theories behind ConceptNet [8]. Its semantic ontology is populated with concepts about everyday objects and events familiar to and suitable for children age four to six years old. These concepts are also aligned with the themes of the stories that the system generates, specifically those focusing on promoting good behavior such as sleeping early and being neat and clean, as well as teaching values like bravery and honesty.

From an input picture containing elements of the story (background, characters, objects) that the child specified using the Picture Editor facility of Picture Books, the story planner creates a story by searching the semantic ontology for a path of relationships between two input concepts [9]. For example, using the two concepts: "break object" (starting concept node) and "punishment" (target concept node), a path of semantic conceptual relations retrieved from the ontology is returned containing the following binary relations:

break object: lastSubeventOf: get punished
get punished: lastSubeventOf: not allowed to play today
not allowed to play today: isA: punishment

This leads to the generation of a story text about breaking an object that can lead to receiving a corresponding punishment, specifically, being disallowed from playing, as shown in the story excerpt in Listing 1.

Listing 1. Story excerpt from Picture Books

> Title: Rizzy the Rabbit learns to be Honest
>
> Rizzy broke the lamp. She was scared.
>
> Mommy Francine saw that the lamp was broken. Rizzy told Mommy Francine that Pinky broke the lamp.
>
> Pinky got punished. Mommy Francine told Pinky that she was not allowed to play today. Pinky cried.

A succeeding implementation, Picture Books 2, also utilized a semantic ontology to allow a theme-based causal planner [1] to generate a sequence of story events depicting the cause-and-effect relationships across a chain of character actions and resulting events in the story world. For example, using the two concepts: "far" (starting concept node derived from the attribute of the story setting camp) and "scared" (target concept node derived from the non-trait of the main story character who is not brave), the following path of semantic conceptual relations is retrieved from the ontology using a chain of "EffectOf" binary relations:

tired: effectOf: far
hungry: effectOf: tired
eat: effectOf: hungry
sleepy: effectOf: eat
sleep: effectOf: sleepy
hear: effectOf: sleep
scared: effectOf: hear

An excerpt of the resulting story is shown in Listing 2. Picture Books 2 was later enhanced to use an agent-based planner [2] whereby its plot agent and a set of character agents are working together, accessing the same semantic ontology, to produce more logical stories. A semantic ontology that provides supplementary domain-specific knowledge to augment existing resources like ConceptNet, VerbNet and WordNet was also explored for Picture Books 2 in the work of Yu and Ong [12].

The planning agent in MAKEBELIEVE [7] performed logical reasoning on the causal relations found in the commonsense ontology extracted from the Open Mind Common Sense (OMCS) project [8] to generate short fictions spanning 5 to 20 lines of story text from an initial seed story. OMCS [11] is a large-scale knowledge acquisition project that turned to the public community instead of the experts to collect commonsense.

Listing 2. Story excerpt from Picture Books

> Title: Danny the Dog learns to be Brave
>
> The camp is far. Danny the dog feels tired. He wanted to eat, since he felt hungry.
>
> Danny the dog ate the marshmallow, therefore he felt sleepy. He slept in a tent.
>
> He heard a sound. Danny the dog feels scared.

The work of Chua and Ong [4] also used crowdsourcing to collect commonsense knowledge, this time from children. The knowledge is then used to provide the relevant concepts needed by story generators to produce story text that are appropriate for them. The system instantiates story templates that have been defined (by child educators, story writers or any adult) using its built-in story scripting language to generate stories with blanks that children can fill-in, as shown in Listing 3. The blanks are associated with relation extraction templates, shown in Listing 4. Once the blanks are filled up, the resulting conceptual relations are stored into an ontology. Only those relations that passed a given validation threshold are then used to further instantiate story templates to enable the system's ontology to grow.

Listing 3. Story with blanks for children to fill-up [4]

> Today I fell happy. _____ (1) _____ is something I do when I am happy.
>
> I am happy because I _____ (2) _____. Another word for happy is _____ (3) _____.

Listing 4. Relation templates for the story in Listing 2 [4]

> (1) ? | EffectOf | $emotion
> (2) ? | Causes | $emotion
> (2) ? | IsA | "action"
> (3) ? | Synonym | $emotion
>
> where $emotion = "happy" (as queried from the existing content of the ontology)

OMAdventure [5], an interactive scavenger hunt game developed at the MIT Media Lab, also uses the knowledge collected from the OMCS project to dynamically generate a virtual story world environment that users can navigate in. As the game progresses, the system queries the OMCS to determine adjacent locations as well as objects that can be found in a given location.

To be able to create the dynamic game environment, OMAdventure [5] needed to filter and extract necessary information from OMCS. Because of the information extraction from OMCS, filtering and getting rid of unwanted information became a problem to consider. OMAdventure also had limited gameplay. It gave the user the freedom to explore the world but it did not provide the ability for the user to interact with objects in the story world.

3. DESIGN OF ALEX

Since interactive storytelling involves the participation of the user, or player, in story development, the Alex interactive storytelling system incorporates collaboration, emergent stories and a virtual peer. Collaboration between the virtual peer (Alex) and the user is used to achieve the goal of the adventure-themed story, in this case, preparing for an upcoming storm, where user participation is necessary in creating different storylines. An emergent story, defined in [3] as storylines generated from the interaction between a number of characters and objects in the virtual world, can branch out to multiple story paths and is used in Alex to make collaboration possible.

Throughout the adventure, the user can interact with the different non-playing characters (NPCs), objects, inventory and items in the virtual world. User inputs can be in any of the following two forms - 1) an option that the user would choose from a candidate list of possible actions provided by Alex; or 2) the interactions that the user would make in the virtual world, such as clicking on an object or navigating to another location. Every selected option and activity performed by the player in the virtual world would affect the outcome of the story. As the adventure progresses, Alex guides and critics the player's chosen options and interactions.

For the system to incorporate interactive storytelling, a representation of the causal relationship of the events that can take place in the story world is needed. The approach employed is through a semantic ontology. Gottlieb and Juster [5] noted that commonsense knowledge can be useful for interactive applications, including those that utilize intelligent agents. The semantic ontology in Alex is used for the representation of its story structure; the attributes of inventory items (those that the user can collect) and objects; the location of characters, inventory items and objects in the virtual world; and the attributes and capabilities of items and NPCs.

3.1 Ontology Design

An ontology is used in knowledge representation to show the relationship between concepts in a given domain through a collection of assertions. Assertions are in the form of binary relations between two given concepts. For example, in the assertion "*apple isA fruit*", "*apple*" and "*fruit*" are two concepts that are related through the relation "*isA*" [8].

Two types of knowledge are represented in the ontology of Alex, namely domain knowledge and operational knowledge. This classification was adopted from the storytelling knowledge representation of Picture Books [9]. Domain knowledge is composed of assertions that can be used to generate the story content and includes concepts on characters, objects, locations and events, and their attributes. Operational knowledge, on the other hand, contains rules that define the story structure and are used to guide and constrain the story planning process.

3.1.1 Domain Knowledge

Five types of concepts comprise the domain knowledge of Alex.

Characters. Stories contain static non-player characters that can either aid in the development of the story or distract the player from accomplishing his/her goal. Pre-defined static characters can be found in different places in the story world, but they remain in one place throughout the entire story. These characters' dialogue can change depending on the current event and can contain vital information to aid the player in accomplishing a certain task.

Inventory Items. Inventory items are entities in the story world that the player can collect and store in his/her inventory repository. Because Alex's main theme is to collect needed items to prepare for an upcoming storm, the inventory items must all be collected to accomplish the goal of the story**.**

Objects. Objects are static items that populate the story world such as trees, lakes, river, and flowers. These items cannot be collected or added to a player's inventory.

Item and Object Descriptions. The semantic ontology contains assertions that describe objects and items, such as the uses of an object ("*usedFor*"), definition of an object ("*isA*"), and things that you can do with an object ("*canBe*"). Table 1 shows sample assertions for each type of relations currently supported.

Table 1. Assertions for describing Items and Objects

Concept1	Relation	Concept2
flashlight	usedFor	giving light
first-aid kit	usedFor	healing wounds
flashlight	canBe	keep
flashlight	canBe	leave behind
flower	isA	woody plant
flashlight	usedDuring	storm
first-aid kit	usedDuring	storm

Attributes. Attributes are concepts that are used to describe an object or an inventory item. For example, the object "*flower*" is related to the attribute concept "*type of plant*" through the binary relation "*isA*".

Location. The location concept points to where an item or an NPC can be found in the story world, using the relation "*locatedAt*", as shown in Table 2.

Table 2. Assertions for describing Locations

Concept1	Relation	Concept2
flashlight	locatedAt	Store
first-aid kit	locatedAt	house
Emily	locatedAt	playground
flower	locatedAt	garden

3.1.2 Operational Knowledge

The operational knowledge refers to those aspects of the ontology that are used to describe the story structure. The story structure consists of assertions made up of story events, system events and criterions. These events can be triggered when a criterion is met and/or an event is triggered by means of user interaction with an NPC or an object.

Story Events. The "*parentOf*" relation represents the causal chain of story events. For example, if the story event "*Juan Give Item*" has a relation "*parentOf*" to "*Play Emily*", then if the event "*Play Emily*" does not have any criteria or requirements to be accomplished, this event will be triggered after the "*Juan Give Item*" event.

System Events. These are used to call certain functions available in the system. For example, the event *End1* can trigger the "*DRAW_ART*" function, which draws an image on the visual interface of the story world whenever the event is called. System events can also be used to check the contents of the inventory. The event "*INVENTORY_<no_of_items>_<relation>_ <concept2>*" can be used as a criterion that specifies the number of items with a specified relation to *concept2* must be found in the inventory. For example, *INVENTORY_3_ISA_ RECYCLABLE* requires the inventory to have 3 items that are "recyclable".

Criteria. These are events that trigger another event that need to be accomplished in order to move the story forward. For example, in order for the positive ending of the story to be reached, the player must have collected all the needed inventory items.

Figure 1 shows a visual representation of the story structure. Nodes labeled N1, N2, N3, N4, N5 and N6 can be either a story event or a system event. Each event can be associated with 0 or more criterions. In the given figure, event N2 has a set of criteria that must be satisfied for the event to be triggered.

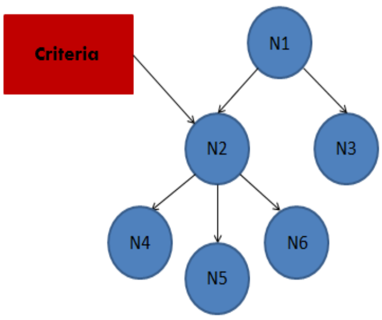

Figure 1. A Visual Representation of the Story Structure

3.1.3 Event Triggers

A story begins from the start node that points to one or more starting events. Three kinds of trigger can force the occurrence of an event in the story world: the NPC trigger, the inventory trigger, and the timer trigger.

NPC Trigger. When the player interacts with the NPC, one or more events can be triggered next. The NPC uses the relation "*consistsOf*" to indicate if a certain event needs to trigger one or more events. If there are multiple succeeding events, the player may be asked to decide which of the candidate events in the "*consistsOf*" relation to perform.

Consider the sample assertions in Table 3 where the current event "*Play Emily*" which has two candidate "*consistsOf*" relations with the events "*Continue with scavenger hunt*" and "*Play Simon Says*". Choosing the first event will move the story forward and closer to the goal, while choosing the second event will disrupt the current story flow as the player decides to play a non-relevant (in relation to the story goal) game instead.

Inventory Trigger. An interaction with an inventory item may necessitate the system to generate a text that prompts the player if he/she wants to store the item into his/her repository or to leave it behind. Depending on the type of items that must be collected, choosing to keep or to leave an item may bring the player closer to achieving the story goal, as indicated in the number and type of items that must be collected to satisfy the criterion leading to a positive ending of the story.

Table 3. Options prompted by NPC

Concept1	Relation	Concept2
Play Emily	consistsOf	Continue with the scavenger hunt
Play Emily	consistsOf	Play Simon Says

Timer Trigger. The system includes a built-in function that triggers an event when one or more criteria are met and/or the pre-requisite *<seconds>_MIN* runs out. For example, at the start of the story, we set up a concept to 2_MIN and made an assertion with the current node using the "*prerequisite*" function. This means that after 2 minutes, the next node will be triggered.

Consider the story plot in Listing 5. A visual representation of the semantic ontology for the given story plot is shown in Figure 2. The circular nodes are story events while the square nodes are criterions. "*Start*", "*Play Emily*", "*Play Simon Says*", "*Key Taken*", "*End1*" and "*End2*" are story events, while "*Emily Interaction*", "*Keep Key*" and "*INVENTORY_8_ITEMS*" are criterions. The "*Play Simon Says*" is a story event that triggers a system event as shown later in Figure 3, when the interaction between the player and the NPC is discussed.

Listing 5. Sample story plot of Alex

The player can talk to Emily (an NPC) or keep a key. If the player chooses to talk to Emily, two candidate events will be made available (based on the sample assertions in Table 1) :

i) continue with the scavenger hunt; or,

ii) play Simon Says.

Assume the player chooses to continue with the scavenger hunt. Certain criteria must be met for the story to develop. If the player collects all the inventory items needed within the prescribed time limit, the first possible ending of the story would take place. But if the player plays a game, the second ending may take place instead.

3.1.4 Virtual Peer

A virtual peer prompts and guides the participation of the user in the progress of the story. The virtual peer, Alex, serves two roles, namely, facilitator and critic. As a facilitator, Alex guides the player as he/she navigates the story world by informing him/her of his/her goal (e.g., collecting this item), and by acting as a medium to interact with NPCs in the story world. As a critic, Alex is able to influence the decision of the player by assessing and criticizing the choices made by the latter.

At every point in the story where the player makes a decision, the role of the peer is determined by calculating the values of the recent events. Each event has an associated value that is represented in the ontology using the "*hasValue*" assertion, as shown in Table 4. *Concept1* models the event concept while

Concept2 is the value. Positive-valued events can bring the player closer to the goal, while negative-valued events will move the player away from achieving the goal of the story. The role of the peer, and consequently the feedback that it provides to the player, depends on the sum of the values of the recently selected actions and events.

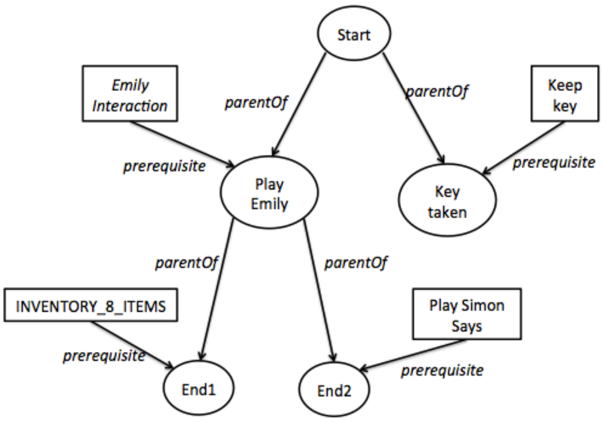

Figure 2. Visual Representation of the Semantic Ontology

Table 4. Assertions Representing the Role of the Peer

Concept1	Relation	Concept2
Play Simon Says	hasValue	-3
Key taken	hasValue	4

3.2 Using the Ontology in Generating Text

Using knowledge from the semantic ontology, Alex generates different types of text at various points in the story, including simple textual description of the objects found in the virtual world and items in the player's inventory, the current options that are presented to the player to move the story forward, and feedback from the virtual peer regarding the player's choices.

Templates are used to provide a generic sentence structure for different purposes like describing an item or an object. Concepts in the ontology are mapped to a lexicon, which is then used to generate the surface form of the story text.

3.2.1 Lexicon

A lexicon is composed of words and phrases that represent the surface form of the relations and the concepts found in the ontology. It acts as a dictionary containing words in the target language. This design of separating the model of actual words to be presented to the user from the concepts stored in the ontology allows for the same set of concepts and planning algorithm to be used to produce story text in varying languages by simply changing the contents of the lexicon.

A concept in the semantic ontology can be mapped to multiple entries in the lexicon, allowing for the generation of less redundant text. A sample set of lexicon entries for the relations found in the semantic ontology of Alex is shown in Table 5.

Listing 6 shows sample sentences that can be generated given the assertions in Table 1.

Table 5. Sample Lexicon

Relation	Synonym
usedFor	is helpful in
usedFor	can be used in
isA	is
isA	is a kind of
usedDuring	used during

Listing 6. Sample text descriptions in Alex

A flashlight is helpful in giving light.

A flashlight can be used in giving light.

A flower is a woody plant.

A flower is a kind of woody plant.

3.2.2 Templates

There are four different kinds of templates: templates used for object and item description, templates used for generating candidate actions, and templates used for generating peer feedback.

Generating Object and Item Descriptions. Since the semantic ontology contains assertions that describe objects and inventory items, the templates and lexicons can be used to generate texts. For example, from the assertions given in Table 1, "*flower-isA-woody plant*", when the flower object is selected, the template with the type "*forObj*" (shown in Table 6) will be used. The *<relation>* tag is used to identify which surface form (represented in the Synonym column in Table 5) is to be used to fill up the template. In this case, the surface form for the "*isA*" relation is the phrase "*is*", leading to the generation of the text "The flower is a woody plant".

Table 6. Sample Templates

Message	Group Reference	Type
What do we have here, it's <article> <object>	1	forItem
Did you know that <article> <object> is <usedFor>.	1	forItem
Oh look! This is <article> <object>.	2	forItem
This is <usedFor>	2	forItem
What should we do with the <object>? Should we <canBe>?	1	forItem
<article> <concept1> <relation> <article> <concept2>	1	forObj

The Type and the Group Reference columns (in Table 6) are used to concatenate generated sentences. For example, when the "*flashlight*" object is selected, the three "*forItem*" templates will be used to generate the story text shown in Listing 7. The content of the <*canBe*> tag will be discussed next.

Listing 7. Sample text generated for "flashlight"

> What do we have here, it's a flash light. Did you know that a flash light is helpful in giving light? What should we do with the flash light? Should we <canBe>?

Generating Candidate Actions. Candidate actions to move the story forward are presented to the player in the form of options. These options are triggered when the player interacts with an NPC or with an inventory item in the virtual world. These options may trigger further events for the story to develop. In this case, the ontology is used to support the branching of the story or to enable the user to participate in developing the storyline.

For example, in Figure 2, when the story reaches the node "*Play Emily*", the player is given a set of candidate actions to choose from, as shown in Figure 3. Candidate actions are story events and represented as circular nodes. The *consistsOf* relations are used to retrieve the set of candidate actions for a given node concept from the ontology. In the given example, the candidate actions are "*continue with scavenger hunt*" and "*play Simon says*".

As indicated earlier, a candidate node representing an event may be associated with a numeric value through the "*hasValue*" relation. This is indicated as a square node in Figure 3. The numeric concept indicates how the role of the peer may change when its associated event has been selected.

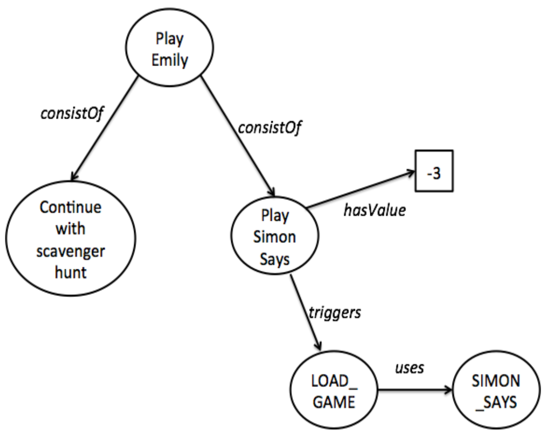

Figure 3. Visual Representation of Candidate Actions

Aside from than the NPC interaction, the options are also triggered by interacting with an inventory item. An example of this is when the user clicks an inventory item. Table 1 shows the relation "*canBe*" associating the concept "flashlight" to the concepts "keep" and "leave behind". This inventory item interaction has a template (found in Table 6) that has the tag "*canBe*". The tag <*canBe*> is used to query the ontology to derive assertions for the "*canBe*" relation. This can lead to the generation of the text shown in Listing 8.

Listing 8. Sample text generated for the tag <canBe>

> What should we do with the flash light? Should we <u>keep</u> or <u>leave behind</u>?

Generating Peer Feedback. The virtual peer, Alex, serves two roles in the system: facilitator and critic. As a facilitator, Alex guides the player as he/she navigates the story world by informing him/her of his/her goal (e.g., collecting this item), and by acting as a medium to interact with NPCs in the story world. As a critic, Alex is able to influence the decision of the player by assessing and criticizing the choices made by the latter.

The virtual peer also acts as an interface between the story world and the user. For example, Alex only appears when the user is interacting with an inventory item or an NPC.

As the story develops, Alex evaluates the story event nodes and checks for the sum of their values. Alex uses the current sum of the value in the numeric concept of the relation "*hasValue*" to evaluate which dialogue to show to the user depending on the threshold. The threshold is set using the "*thresholdGreaterThan*" and "*thresholdLessThan*" relations. A combination of tags and canned phrases are used for the feedback of the peer, as shown in Table 7.

Table 7. Templates for Peer Feedback

Dialogue	Relation	Concept2
We're doing a really good job. Remember we still have to look for the following: <ITEM_USEDDURING_STORM>	threshold GreaterThan	1
I think we should focus on collecting items needed before the storm comes.	thresholdLess Than	-11

For example, if the current threshold is -12, Alex assumes the critic role by generating a feedback that reminds the player to focus on collecting the needed items.

On the other hand, if the current threshold is 2, then Alex switches to the facilitator role and generates an encouraging dialogue using the first template in Table 7. Notice that this template has a tag <*ITEM_USEDDURING_STORM*>. This tag is used to query the ontology, in this case, to search for assertions involving the "*usedDuring*" relation and "*storm*" as the second concept. Depending on the current concepts available in the ontology, the results may include "*flashlight*" and "*first-aid kit*" (given the example assertions in Table 1).

4. TEST RESULTS AND FINDINGS

Alex's knowledge base has been manually populated with the relevant assertions and templates needed to generate one story with two alternative endings. Currently, the system has ten templates that can be used to generate the dialogue of the NPCs, the description of items and objects, and the feedback of the virtual peer. The semantic ontology contains 20 unique concepts and 19 unique relations.

There are 166 assertions available in the ontology distributed as follows: 11 assertions and 3 relation names for objects; 66 assertions and 6 relation names for inventory items; 22 assertions and 4 relation names for characters; and 67 assertions and 10 relation names for the story structure.

Two types of testing were conducted to validate the appropriateness of the design of the semantic ontology. The text generation engine was tested to determine if the correct semantic concepts can be retrieved from the ontology to produce the different types of text, e.g., item and object descriptions, and candidate actions. The second testing involves validating if the storytelling engine can generate different types of stories with the basic theme of “collecting a specified set of items”.

34 children age 6 to 8 years old from different schools (San Ildefonso Learning Center, Grace Christian College, Hope Christian High School and Angelicum College) participated in the user testing. At the start of the user testing, the participants were given an orientation and demonstration of the different features of the system. During the first few minutes of the testing, the child participants were guided in using the system. Each child was then allowed to use the system freely under close observation by the team. A post-debriefing was then conducted.

The children’s individual responses, their interaction with the system and their comments were all observed and summarized. A set of evaluation questions was developed to guide the team during the observation and debriefing sessions.

The story used in the validation process is about collecting items in preparation for an upcoming storm. If the player is able to collect all the items before the storm comes, then the story would end with Alex and his mom deciding to stay safely in their house. But if the player failed to help Alex collect all the items, then Alex and his mom would be forced to go to the evacuation center.

4.1. Validating the Use of the Ontology in Generating Text

The text generation engine was validated by means of creating different template messages with each template containing different tags. The sample test data are shown in Table 8.

Object and Item Description. The players find the generated text that describe objects and inventory items to be not redundant, as seen in the results in Table 9. 88.24% of the participants found the descriptors to be helpful because the definitions were descriptive enough for them to understand the usage of the selected item or object. On the other hand, 11.76% of the participants found the definitions to be not helpful, as the case for the text “Oh look! This is a phone. It used in calling people”. Some of them, especially the children who are 6 years old, did not fully realize the importance of acquiring the phone. Since the target age group is children aged six to eight years old, simple text may not provide enough details as to the importance of the item to the story.

Options. Table 10 shows the results of evaluating the options that were provided to the player. 79.41% of the participants find the options to be stated clearly. Furthermore, the same number found the dialogue to make sense, and is both believable and sufficient. However, 20.59% of the participants did not realize that they needed to make a choice because it was not clear who the speaker is, nor was it clear how the selection of an option would affect the story flow.

For example, an NPC named *Gabby* generated the dialogue shown in Listing 9.

Listing 9. Sample dialogue by an NPC

```
Hi Alex! Do you want to help me in sorting the trash?

    - Help in cleaning
    - Help later
```

Interaction. This criterion seeks to determine if the player is able to participate in story development through the different interactions present in the story world, such as when the player interacts with an object or with an NPC. The system should generate a description of the object in the former case, and a dialogue in the latter case.

As can be seen in Table 11, 94.12% of the participants provided an affirmative answer, while 5.88% stated that the system is lacking in interaction. The system provided the necessary descriptions and dialogue, however, it did not provide simple gameplay interactions, such as giving the player the ability to trigger the state of an inventory item.

Peer Feedback. This criterion is used to determine if the player is able to distinguish the two roles of Alex. It is also used to determine if the players find the presence of a virtual peer to be useful in facilitating them on what to do throughout the story. As shown in Table 12, 76.47% of the players can distinguish the two roles of the virtual peer. 23.53% did not notice that the peer had two roles because the generated texts were too similar to one another.

Table 8. Sample Concepts Used to Fill Up Template Tags

Type of Tag	Tags	Concepts to fill up the tags
For item	<object>	Flash light
For object	<concept1>	Flower
For object definition	<concept2>	woody plant
For relation	<relation>	Used For
For article	<article>	a / an
For options (item interaction)	<canBe>	Keep, Leave behind
For options (NPC interaction)	<can>	Continue with the scavenger hunt, Play Simon Says
For peer feedback	<ITEM_<relation> _<concept2> e.g. <ITEM_USEDDU RING_STORM>	Flash light, first aid kit

Table 9. Test Results: Object and Item Description

Criteria	Yes	No
The sentences are not redundant.	100.00%	0.00%
The definitions are helpful.	88.24%	11.76%

Table 10. Test Results: Options

Criteria	Yes	No
Options are clear	79.41%	20.59%
Dialogue makes sense	79.41%	20.59%
Dialogue is believable	79.41%	20.59%
Dialogue is sufficient	79.41%	20.59%

Table 11. Test Results: Interaction

Criteria	Yes	No
The system is interactive.	94.12%	5.88%

For example, the following generated feedback texts are too similar that users failed to distinguish the two roles portrayed by the peer.

Critic Role: *I think we should focus on collecting items needed before the storm comes.*

Facilitator Role: *We're doing a really good job. Remember we still have to look for the following: flashlight and first-aid kit.*

Only 17.65% of the participants like the presence of the peer while 82.35% of the participants did not like having the peer around because the feedback provided by the peer did not sound like a storyteller. Moreover, the children wanted to explore the story world on their own especially those who are 7 to 8 years old.

Less than half or 44.12% of the participants find it necessary to have a virtual peer for them to understand the story because the peer gave them feedback and instructions on what to do. On the other hand, 55.88% did not think the peer was necessary in understanding the story because the peer only told the story and gave them feedback as the event unfolds. After a while, some of them forgot what had happened or the past decisions that they had made. In this case, the participants did not fully realize what they did wrong or what led to the portrayal of the critic role because they lost track of what was happening in the story.

Another reason that may rationalize as to why 55.88% of the participants did not find the peer a necessity is because of its interaction model. Alex only shows up when the user interacts with an inventory item or an NPC. A different storyteller interaction model may change the result.

Table 12. Test Results: Peer Feedback

Criteria	Yes	No
The child can distinguish two roles	76.47%	23.53%
The child likes having the peer around	17.65%	82.35%
The peer is necessary to understand the story	44.12%	55.88%

Overall, the use of a semantic ontology enabled the flexible representation of the story structure and helped in text generation. Test results showed that there are still a lot of things to consider in the design of the semantic ontology such as the portrayal of the personality of the peer as a storyteller and as someone, who encourages the user to participate; the consistent personality of the NPC; the descriptiveness of the sentences generated; and the logical formulation of reasonable sentences based on the story development and that consider past decisions made by the player.

4.2. Validating the Use of the Ontology in Generating Stories

To validate the flexibility of the ontology design, two additional stories, each with three alternative endings, were created. The first story involves a race to collect the most number of recyclable and non-recyclable materials. The second story involves retrieving a blue box for the user's mother.

In order to create these kinds of story, the semantic ontology must be populated with the necessary knowledge. However, only one goal per story can be supported by the system. Further note that the ontology design only supports stories with the theme "find or collect something", with or without a specified time limit.

4.2.1. "Racing to Collect the Most Number of Materials" Story

This story involves a race to collect the most number of biodegradable and non-biodegradable materials. The current capacity of the inventory is set to 3. Two criteria were defined using the system event *INVENTORY_<no_of_items>_ <relation>_<concept2>*, one for recyclable and the other for non-recyclable items.

Four events can happen in this story: the Instructions, the Start of the story, the Collection of 3 Recyclable items, and the Collection of 3 Non-Recyclable items. To represent the sequence of the three main events, the "*parentOf*" relation is used, as shown in Figure 4.

Recyclable and non-recyclable items are defined through the "*isA*" relations in the semantic ontology, as shown in Table 13. The "*canBe*" relation is used to store items into the inventory when the player chooses the "*keep*" option.

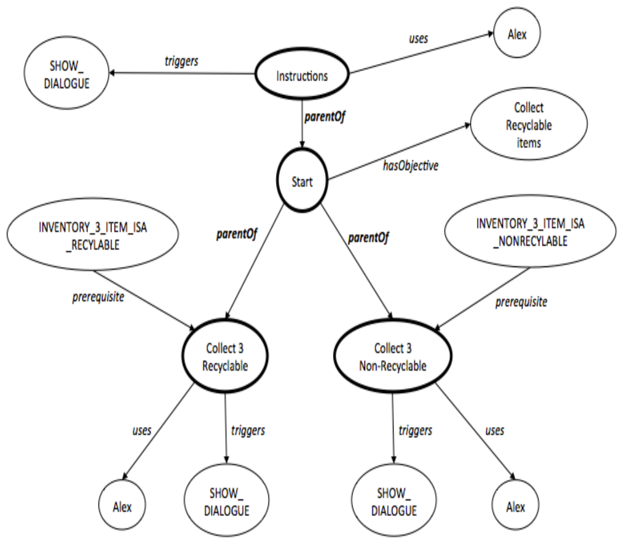

Figure 4. Visual Representation of the Story "Racing to Collect the Most Number of Materials"

Table 13. Partial Ontology for Collecting Recyclable Materials

Concept1	Relation	Concept2
apple	isA	non-recyclable
plastic bottle	isA	recyclable
plastic bottles	canBe	keep
plastic bottles	canBe	leave behind
plastic bottles	usedFor	storing water to drink

Table 14. Sample Table Representation of the Candidate Actions

Concept1	Relation	Concept2
Gabby	can	PROVIDE_CLEAN_OPTIONS
PROVIDE_CLEAN_ OPTIONS	consistsOf	Help in cleaning
PROVIDE_CLEAN_ OPTIONS	consistsOf	Help later
Help in cleaning	triggers	LOAD_GAME
LOAD_GAME	uses	GAME_3

In this story, helping an NPC clean up the environment can triger a system game. In this case, the option is represented using a semantic ontology (shown in Table 14).

4.2.2. "Retrieving a Blue Box" Story

This story is about retrieving a blue box for Alex's mother. Three main events can happen in the story: the Start of the story, the Retrieval of the Blue box, and the Retrieval of the Red box. The visual representation of the story is shown in Figure 5. The "*parentOf*" relation is used to represent the three main events.

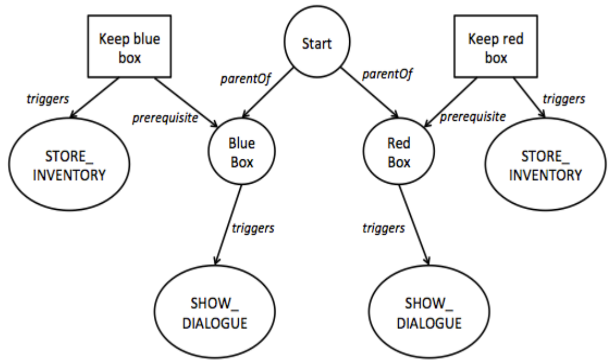

Figure 5. Visual Representation of the Story "Retrieving of a Blue Box"

The main events can happen depending on which item the user collects. The dialogues that will be presented will also depend on the item that the user has collected. The criterions, represented as square nodes in Figure 5, must be set using the "prerequisite" relation. The corresponding dialogues are defined using the "*triggers*" relation.

Table 15. Partial Ontology for Retrieving a Blue Box

Concept1	Relation	Concept2
Blue box	color	blue
Blue box	canBe	keep
Blue Box	canBe	leave behind
Red box	color	red
Red box	canBe	keep
Red box	canBe	leave behind

In order for the options to appear, the "*canBe*" relation is set to the blue box and the red box. For the inventory to store the item, the keep option must be selected. The "keep" concept will then trigger the "*STORE_INVENTORY*" system event. The partial ontology for this story is shown in Table 15.

5. CONCLUSION

In this paper, we presented the design of a semantic ontology to support the generation of text for the Alex Interactive Storytelling system. Specifically, these include the generation of descriptions for objects in the story world and items in the player's inventory, the current options or candidate actions available to the player to move the story forward, and the feedback provided by the virtual peer based on the player's actions or decisions. The ontology contains concepts and relations that describe the characters, locations, objects, items and events needed to generate text at various points in the story.

Alex is currently targeted towards children aged 6 to 8 years old. Test results showed that the participants find the definitions of objects and items to be helpful, the candidate actions are presented clearly, and the simple dialogue are sufficient for storytelling. Mixed results were received regarding the presence and usefulness of a virtual peer.

The primary purpose of using a semantic ontology in an interactive storytelling environment is to reduce redundancy in storing concepts and their relations, and to generate varying text to prevent monotony in the stories produced by the system. Test results showed that the design of the ontology allows the definition of other stories with the basic theme of "collecting items in the story world". However, the semantic ontology currently supports only one goal for each defined story. Incorporating one or more goals depending on the user's actions and decisions can be the subject of a future research. Furthermore, the knowledge representation is flexible but is limited in the art assets available in the system.

On the other hand, limitations may arise in designing a well-balanced story. The delivery of the story versus the size or the number of branches in the story should be considered. Having multiple branches may involve having a big knowledge base. Further research may involve what kind of problems will arise when long stories having multiple branches are involved.

Simple user interactions that can change the state of an object, such as pressing a button to change the state of a light bulb to on or off, should also be explored to increase interactivity. There might be a difference if the state of objects is represented as part of the semantic ontology.

Peer feedback can also be improved on using the information given by the semantic ontology. Further work on the personality that can be portrayed by the virtual peer to better emulate a human storyteller should be conducted, as well as the effect of the personality to the types of dialogue and lexical terms to be used in the text to be generated.

Future research should also consider expanding the design of the ontology to support various themes beyond the simple collection task. The use of a player's activity log can also be explored to support the generation of text describing the previous actions performed by the user in the virtual world, and how this led to the progression of the story. Context-sensitive dialogue can then be generated to make the interaction between the peer and the player sound more natural.

6. REFERENCES

[1] Ang, K., Yu, S. and Ong, E. 2011. Theme-Based Cause-Effect Planning for Multiple-Scene Story Generation. In Proceedings of the 2nd International Conference on Computational Creativity, 48-53, Mexico City, Mexico, April 27-29 2011.

[2] Ang, K. and Ong, E., 2012. Planning Children's Stories using Agent Models. In D. Richards and B.H. Kang (Eds.), 2012 International Pacific Rim Knowledge Acquisition Workshop (PKAW 2012), LNAI 7457, 195-208, September 5-6 2012, Malaysia. Springer-Verlag Berlin Heidelberg.

[3] Bosse, T., Memon, Z., and Treur, J. 2007. Emergent Storylines Based on Autonomous Characters with Mindreading Capabilities. In Proceedings of the Sixth International Conference on Intelligent Agent Technology, IAT'07. IEEE Computer Society Press, 2007, pp. 207-214.

[4] Chua, R. Jr. and Ong, E. 2012. Commonsense Knowledge Acquisition through Children's Stories. In D. Richards and B.H. Kang (Eds.), 2012 International Pacific Rim Knowledge Acquisition Workshop (PKAW 2012), LNAI 7457, 244-250, September 5-6 2012, Malaysia. Springer-Verlag Berlin Heidelberg.

[5] Gottlieb, D., and Juster, J. Generating a Dynamic Gaming Environment using OMCS. Unpublished Technical Paper, MIT Media Laboratory, Massachusetts Institute of Technology.

[6] Hong, A.J., Solis, C., Siy, J.T., Tabirao, E., and Ong, E. 2009. Planning Author and Character Goals for Story Generation. In Proceedings of the NAACL Human Language Technology 2009 Workshop on Computational Approaches to Linguistic Creativity, 63-70, Colorado, USA. Morristown, NJ, USA: Association for Computational Linguistics.

[7] Liu, H., and Singh, P. 2002. MAKEBELIEVE: Using Commonsense to Generate Stories. In Proceedings of the Eighteenth National Conference on Artificial Intelligence,957-958. AAAI Press.

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

[9] Ong, E. 2010. A Commonsense Knowledge Base for Generating Children's Stories. In 2010 AAAI Fall Symposium on Common Sense Knowledge, pp. 82-87. AAAI, CA.

[10] Porteous, J., Cavazza, M., and Charles, F.. 2010. Applying Planning to Interactive Storytelling: Narrative Control Using State Constraints. ACM Transactions on Intelligent System Technology (TIST), 1(2):1-21, November 2010. New York: ACM. http://dx.doi.org/10.1145/1869397.1869399

[11] Singh, P., Barry, B., and Liu, H. 2004. Teaching Machines about Everyday Life. BT Technology Journal, 22(4):227-240, October 2004. Kluwer Academic Publishers.

[12] Yu, S. and Ong, E., 2012. Using Common-Sense Knowledge in Generating Stories. In P. Anthony, M. Ishizuka and D. Lukose (Eds.), 12th Pacific Rim International Conference on Artificial Intelligence (PRICAI 2012), LNAI 7458, 838-843, September 5-7 2012, Malaysia. Springer-Verlag Berlin Heidelberg.