Learning about Good Healthcare Practices through Interactive Storytelling with Virtual Peers

Raisa Gennel Lee, Janine Bianca Regala, Angeline Tan, Janine Mica Tan and Ethel Ong[1]
Center for Language Technologies - AdRIC, De La Salle University
[1]ethel.ong@dlsu.edu.ph

ABSTRACT
Storytelling can be used to raise awareness on healthcare practices among young children. Combined with interactive environments, children can be given a virtual simulated world to learn about the effects of practicing good healthcare habits with the guidance of virtual peers. In this paper, we present Sarah, a virtual peer that shares interactive healthcare stories centering on events that may lead to symptoms and recovery from common childhood ailments. Armed with a collection of domain knowledge about healthcare facts, the virtual peer utilizes story planning and dialogue generation strategies to facilitate a text-based storytelling session with children age 7 to 10 years old. Children participate in the storytelling by using free-form input text to respond to questions posed by Sarah. Results from end-user validation showed that children prefer to interact with virtual peers that exude a more friendly persona, and that can generate relevant responses aligned to their story text.

Keywords
Virtual peer, interactive storytelling, healthcare stories, story planning and generation, domain knowledge base

1. Introduction
Healthcare organizations have been proactive in making healthcare information more accessible to the general public by publishing substantial amount of health-related educational materials online and in print. Despite this, healthcare providers are experiencing an influx of patient visits seemingly caused by their lack of understanding of the information provided. The issue lies not on the inadequacy of the amount of available healthcare information, but rather on the ability of the patients to access and process the given information [1].

Stories can be used to provide context to patient care by breaking down difficult concepts into simplified comprehensible facts [2]. Osborne [3] noted that telling a good story "can frame important messages in ways that make them memorable for the listeners". African-American women, for example, learned about breast cancer through stories that increased their abilities to classify misconceptions and implications, and acquire knowledge on possible approaches and treatments [4]. This purposeful storytelling works by bringing across factual information embedded within real-life situations to provide a meaningful and relevant yet human-centric perspective on a particular topic of interest, such as observing good healthcare practices.

In children, storytelling enhances critical thinking skills, vocabulary and language patterns. It has been used as an effective teaching tool in language and literacy development [5] to allow children to gain a deeper understanding of what is being taught [6]. Storytelling also provides an avenue for experiential learning [1], through cases, role playing and reflective writing.

Computing technologies such as Patient Voices [7] have enabled a form of digital storytelling where patients can share their personal stories to others using mixed media form – video, audio, images and text [8], to advocate change in healthcare. Interactive storytelling environments also employed virtual peers to collaborate with the users as they navigate the virtual story world. Strong user participation is encouraged, and this affects the progression of the story, allowing users to practice decision-making and see the effects of their choices in simulated environments. Such systems have been designed as tools to help preschoolers develop their language literacy skills [9]; help children with autism learn about proper social behavior [10] and language [8]; and simulate marriage counselling [11], business case studies [12] and job interviews [13].

Virtual peers are autonomous agents that when utilized properly, can have the potential to be healthy partners in literacy learning [14]. They can engage in conversations with their human users by building rapport and maintaining the latter's attention, both of which are crucial in children's acquisition of healthcare knowledge. Research suggests that the establishment of social relationship plays a major role to a child's learning and readiness or willingness to engage in communication and collaborative relationships [9].

While previous research in interactive storytelling encourage user participation by giving choices at key points in the story plot, these choices are predefined and placed a constraint on the types of input that users can provide. The main contribution of our study is in exploring the use of free-form text to enable users to input suggestions for virtual peers to perform in the virtual story world. This dialogue exchange gives users the perception that they are active participants in the story generation process, and that their inputs are taken into consideration by the story planner to affect how the events in the story world will unfold.

The rest of the paper is organized as follows. Section 2 defines the roles that our virtual peers, Sarah and Liam, play in the interactive storytelling environment. Section 3 presents the design of the storytelling knowledge, particularly the domain knowledge that gives Sarah information about healthcare concepts, and the story world model to track the dynamic attributes of story elements. Interactive storytelling employs natural language processing techniques to process user inputs and to generate story text. These are described in Section 4, specifically the planning algorithm that Sarah uses in order to generate healthcare stories. Virtual peers also employ a set of dialogue strategies to initiate, encourage and sustain user participation during the interaction. In Section 5, the generation of text-based dialogues to engage children in conversations about story events and effect of character actions to the story world is described. We present the results of our evaluation of the child-peer interaction and the quality of the healthcare stories based on content and linguistic features in Section 6. We conclude our paper with a discussion of our findings and recommendations for future work.

2. Virtual Peer in Healthcare Stories

Virtual peers are artificial agents designed to provide continuous support and to encourage learning by participating in a shared task through collaboration, competition and demonstration [15]. Sarah is a virtual peer that uses interactive healthcare stories to impart information about common childhood ailments to children. To facilitate effective human-computer interaction, Sarah has been designed to serve three roles based on the tasks it is expected to perform, namely, companion, facilitator and teacher. As a *companion*, Sarah is the playmate of the main story character Liam, and stays with him and the user throughout the story. As a *facilitator*, Sarah guides the conversation between Liam and the user, to ensure that they both stay in line with the story flow towards the achievement of the theme. As a *teacher*, Sarah takes the responsibility of explaining the possible outcome of Liam's and the user's choices and actions by stating details about the ailment as the story unfolds – the causes, symptoms, treatment and prevention.

A healthcare story delivers, in narrative form, a discussion of a specific illness portrayed within a real-life or imagined scenario. This can serve as a tool in teaching the facts and misconceptions on proper healthcare, through portrayal of events by story characters that readers can develop affinity to. Healthcare stories are also utilized in focus group discussions and therapy sessions to help patients and family members realize the implication of the disease. Similar to stories, a healthcare story contains the elements of theme, characters and plot. The *theme* is the central idea of a story that the author wants to relay to its readers. The themes of our interactive stories revolve around the top 10 ailments that commonly affect Filipino children, which include the common cough and cold, fever, acute gastroenteritis, bronchial asthma, chicken pox, dengue, measles, pneumonia, urinary tract infection and hypersensitivity.

Characters portray key roles in a story. The major story characters are those that undergo developmental experiences and whom the readers may attach to. For children, memorable characters are those described to be in their same age range or are only slightly older. In this research, the main character is a virtual peer named Liam, who is designed to have the same age as the user so that the child may be able to visualize himself/herself in Liam's place. Sarah, on the other hand, is designed to be slightly older than Liam to establish trust and credibility as a facilitator and teacher but can still be perceived by the user as approachable. In the interactive storytelling environment, Liam would be introduced as a friendly boy who wants to play with the user. As the story progresses, the user makes decisions and suggestions, in free-form text, on actions that will be executed by Liam. Every decision made would affect the health status and sickness status of Liam. Depending on the sequence of actions, Liam may become afflicted with one of the 10 ailments, and either get well or still be sick from this when the story concludes.

The *plot* establishes a unified and coherent sequence of events that takes place in the story. Our central story plot focuses on finding ways by which the user, through his/her decisions, can help Liam recover from his ailment. For children who are 7 to 10 years old, a simple plot with three distinct parts is followed, namely *start*, *middle* and *end*, as illustrated in Figure 1. The healthcare stories commence with the introduction of the virtual peers, Sarah and Liam, along with the setting where Liam is playing in a specific location. The middle is the part where Liam's ailment is revealed; the theme, which represents one of the 10 ailments, is selected by the child prior to the start of the storytelling. The story contains text that gives explanation about the ailment, how it was acquired, and ways for treatment.

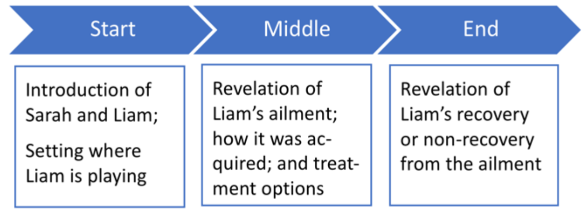

Figure 1. Interactive Story Plot

Finally, the end comes with the revelation of whether Liam is cured of his ailment or not, depending on the user's suggestions of healthcare practices that Liam should perform. Should the user's suggestions cover correct actions, the story concludes with Liam being healthy and Sarah lauds the user in coaching Liam to do what is necessary. On the other hand, had the user suggested Liam to continue playing despite his condition, the story concludes with Liam still being sick, and Sarah reminding the user that this happened because proper practice was not followed. Though happy endings are the norm, in this research, the priority is for the user to understand the implications of his/her decisions towards the recovery or non-recovery of Liam.

3. Modelling Healthcare Stories

The design of the interactive stories is inspired by the Ed-W model [16] on educational digital stories. This model outlines a number of steps to help the users in grasping the lesson or topic of discussion, beginning with the introduction of a problematic situation that uses the virtual peer's knowledge about the topic. The problem continues to worsen due to the protagonist's inappropriate actions that are influenced by the lack of knowledge or misconception of the user about the topic, leading to an escalation of the challenge faced by the main character. Finally, the protagonist is forced to reflect on his experiences and to seek for information and help to explain the problem, leading to his awareness of the misconceptions and the appropriate interventions he must take to correct the problem. Figure 2 shows the model adapted for healthcare stories.

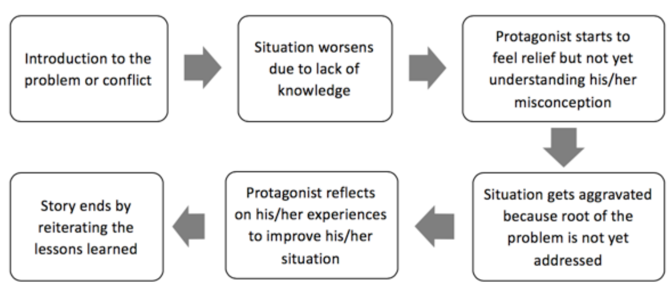

Figure 2. Model for Healthcare Stories

Healthcare stories covering the 10 common ailments of Filipino children were designed in this manner to illustrate the problems and consequences of inappropriate actions, and to emphasize the importance of proper practices in maintaining good health among children. Sarah delivers the necessary story line to impart these messages. For the virtual peer to perform all the tasks associated with its roles, it requires a set of storytelling knowledge needed to generate the stories and to adapt them as dialogue responses to the user's free-form text inputs. A knowledge base is built to store all the healthcare concepts as domain knowledge, story patterns, dialogue templates, and virtual story world state.

3.1 Domain Knowledge

Domain knowledge contains facts and events that relate concepts about the different ailments. These facts and events are stored as binary assertions following ConceptNet's [17] *semantic-relation(concept1, concept2)* format. The assertions are categorized into four, namely *setting*, *virtual peer*, *disease* and *general knowledge*.

The *setting* category contains assertions that can be used to describe the story setting, including time, location such as *isA(park, location)*; activities that can take place; and objects that exist in the story world, such as *atLocation(play-mud, park)* and *usedFor(mud, play-mud)*.

The *virtual peer* category includes assertions to describe the physical state of Liam using the semantic relation *hasProperty*, e.g., *hasProperty(hands, dirty)* and *hasProperty(body, tired)*. The actions or activities that Liam can perform in the story world are described using the semantic relation *can*, e.g., *can(hasProperty(hands, dirty), wash-hands)* and *can(hasProperty(hands, dirty), eat-food)*. The recipient of an action is modelled using *receivesAction*, as in *receivesAction(hands, wash)*, while *isEffectOf* defines the causal chain of events or the postcondition of executing an action, e.g., *isEffectOf(wash-hands, hasProperty(hands, clean))*.

The *disease* category models health-related information associated with the different ailments. Some examples of semantic relations in this category include *canLeadTo(play-in-the-rain, fever)* that states the possible consequences or symptom that may arise from performing a particular action or activity; the type of illness that can be derived from a symptom that Liam is experiencing, such as *isSymptomOf(fever, pneumonia)*; the possible causes of a particular disease, e.g., *isCausedBy(pneumonia, bacteria)*; the different treatments for a disease, e.g., *isTreatedWith (pneumonia, take-medicine)* and *isTreatedWith(pneumonia, rest)*; preventive practices that can help avoid a specific ailment, e.g., *isPreventedBy(stomach ache, avoid-drinking-raw-milk)*; the specific body parts that are harmed by a particular disease, such as *affects(pneumonia, lungs)*; and descriptions of body parts that are appropriate for the target audience, e.g., *definedAs(lungs, pair of packet-like organs inside your chest)*.

The *general knowledge* category defines the different relationships of concepts stored in the domain knowledge base. For example, the assertions *oppositeOf(tired, rested)* and *oppositeOf(clean, dirty)* are used to describe opposite concepts; while the assertions *sourceOf(potassium, banana)* and *sourceOf (vitaminC, orange)* are used to exemplify food sources that provide certain nutrients.

3.2 Virtual Story World

Character actions, occurrences of events in the story and user decisions affect the states of characters and objects in the story world. These states are considered by the virtual peer during story planning to determine the next event that may unfold. A story world model, as illustrated in Figure 3, is used to track these states. The model consists of three components – the story theme, the character representation, and the story events.

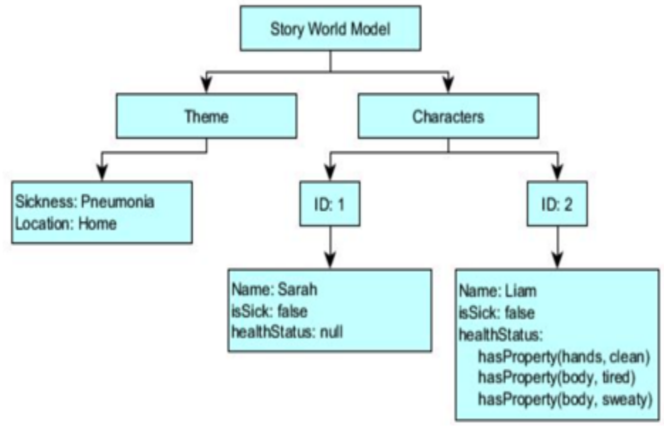

Figure 3. Story World Model

During story planning, the theme guides the selection of the sequence of events that can occur in the story world. The theme is selected based on the actions that unfold in the beginning of the story. The story begins by asking the user to specify a location where the story will take place using the interface depicted in Figure 4. The first event, chosen randomly by extracting from the domain knowledge base a list of candidate actions that can take place in the specified location, is then played. The next events that will unfold are dependent on the user's inputs or decisions with regards to the actions that Liam should do next given his current physical state. For each action that exhibits incorrect healthcare practice, its corresponding symptom will be stored in the story world model. A specific ailment is then determined from the collected symptoms, and are used to fill-in the details of the dialogues between Sarah, Liam and the user.

Figure 4. Location Selection Page

The character frame represents attributes of Sarah and Liam, specifically their name, health status and sickness status. The *name* attribute is initialized at the start of the story and remains constant. The attributes *health status* and *sickness status*, represented as assertions using the semantic relation *hasProperty*, are dynamically updated based on the associated postconditions of the actions performed by the virtual peers as directed by the story planner and in consideration of the suggestions from the user. Postconditions are adapted from the work of Ang and Ong [18] to describe the effect of the execution of an event on the state of the story world. In automated story generation, such conditions are used by the story planner to select candidate events based on the intended outcome of the story. The next section expounds on the use of postconditions when identifying the possible actions for Liam to perform.

4. Story Planning

Following the approach of Façade [11], the plot structure is divided into three (3) episodes – *introduction*, *doActivity* and *discussOutcome* -- corresponding to the beginning, middle and ending parts of an interactive healthcare story. This is illustrated in Figure 5.

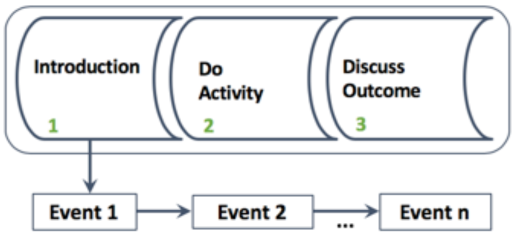

Figure 5. Episodes in a Plot Structure

An episode manager coordinates the execution of the episodes in order to ensure that the resulting story conforms to the prescribed story plot while allowing the user's contributions to affect certain parts of the story flow. An episode is deemed completed once Sarah has generated a response to the user's input. The end of one episode also triggers the start of execution of the next episode, until all episodes comprising the story plot have been completed and the story ends.

Episodes are comprised of a group of logically related events. As depicted in Figure 6, an event can be an action to be performed by Liam, or a dialogue turn used by the virtual peers to deliver a message. An *action* event is defined with a set of preconditions, postconditions, and the object needed to carry out the action. Preconditions are also adapted from the work of Ang and Ong [18] to specify the constraints that must be satisfied before the event can take place in the story world.

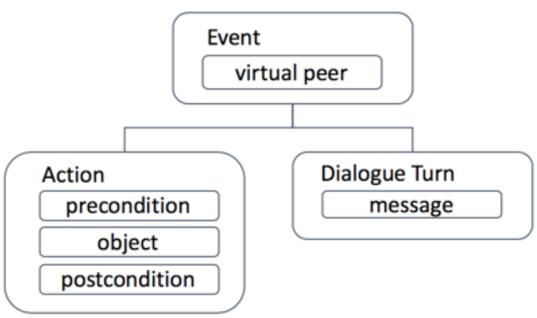

Figure 6. Types of Events

4.1 Selecting Candidate Actions

During story planning, candidate actions to be performed by Liam are selected based on the satisfaction of their preconditions or constraints for execution. Consider the event *play* depicted in Figure 7. Its preconditions for execution require the current location to be the park (*atLocation(park)*) and Liam's body is rested (*hasProperty(body, rested)*).

After the execution of this event, the postconditions specify the effect of the event to the story world – Liam becomes tired because of playing (*hasProperty(body, tired)*), his hands are dirty (*hasProperty(hands, dirty)*) and his back is sweaty (*hasProperty(back, sweaty)*). The planner may source additional knowledge from the domain knowledge to provide details to the story, such as the use of an *outdoor toy* during a *play* event.

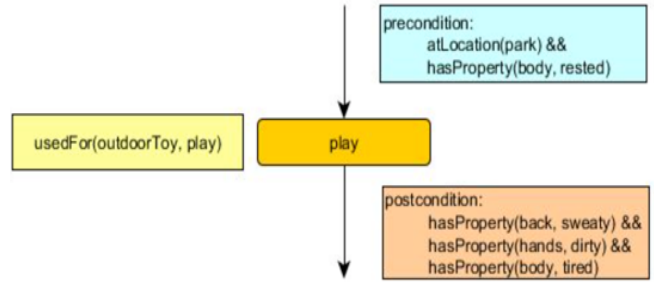

Figure 7. Episode Model for the Event "Play"

Should Liam decide to eat next, the same processes of checking the preconditions and updating the story world based on the postconditions are performed. Continuing from the event *play* which caused Liam's hands to become dirty, the planner looks for an alternative event to reverse the negative attribute, e.g., *wash*. Faced with a choice to either wash his hands first or immediately eat, Liam asks the user for the latter's advice. If the user decides that Liam should wash his hands first, Liam will end up with *clean hands*, and a *full stomach* (the outcome of *eat*). On the other hand, if the user decides to forego washing of hands, the *eat* episode would still be executed, but this time, Liam would show symptoms that correspond to the negative consequence of his (the user's recommended) action.

Depending on how these events are unraveled, the episode manager can dynamically insert additional events should the user's input fail to arrive at a target story world state. Consider the episode chain shown in Figure 8. In an ideal situation, each of the *doAction* comprising an event associated with an episode will be executed in sequence. However, because the user's input may have caused a negative story world state, such as *Liam's hands are dirty*, the next action in the chain cannot be performed. Thus, an action called *inviteReverseAction* has to be inserted to trigger a search for an action that would reverse the negative property of *Liam's hands*, for example, *wash*(*hands*). Inserted actions may cause additional actions to be dynamically created, e.g., *doReverseAction*, which leads to an update to the story world to one that is ideal for the next action in the sequence to be performed.

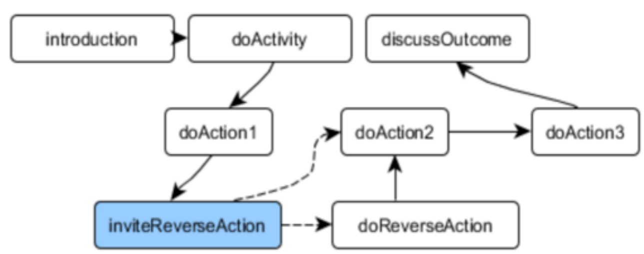

Figure 8. Episode Chain with Action Sequence

Figure 9 shows the branching story graph that illustrates the possible chain of episodes, depicted as nodes, from the sample case. Narrative episodes are indicated using green-colored nodes while decision points to be presented as questions to the user are represented as white-colored nodes. Properties represent the dynamically changing values of attributes in the story world model. Branching story graphs, a term referred to in Riedl [19], are pre-scripted at design time and are typically used in interactive storytelling systems to allow users to perform decision-making at key points in the story. From the graph, the story planner maps out the next episode by taking into account the contents of the story world model and the user's decision. It is important to note, however, that the user's control on the flow of the story is limited by the number of edges in a particular path in the graph.

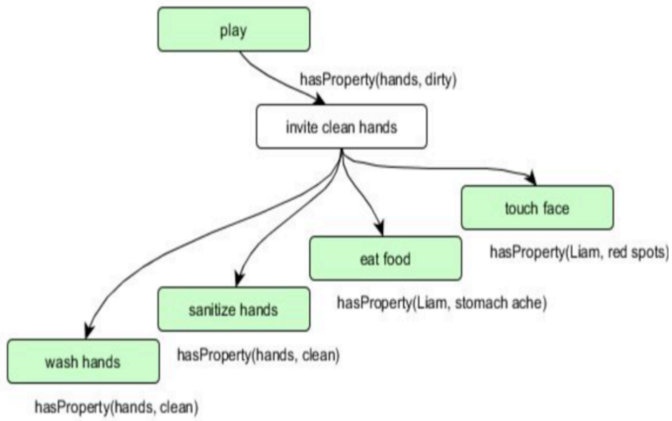

Figure 9. Branching Episode Graph

All the action events that Liam has performed are tracked to prevent the generation of redundant story text. Sarah also references these actions at the latter part of the storytelling session to generate the content for the *discussOutcome* episode. The planner generates one of two possible action sequences, as depicted in Figure 10, to explain how Liam's correct or incorrect health practices led to his ailment and recovery. In the *discussSickness* action sequence, Sarah reiterates the series of events that led to Liam's sickness and eventual recovery (or non-recovery), and summarizes the symptoms, the treatment and good healthcare practices for prevention.

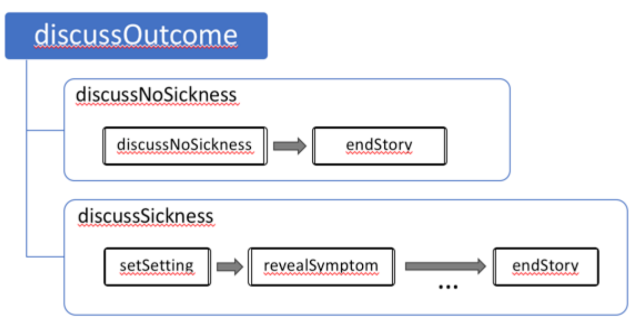

Figure 10. Alternative Action Sequence for *discussOutcome*

4.2 Generating Story Text

Interactive stories encourage user participation by performing decisions during key points in the story. This may affect the subsequent story flow, and the story text to be generated. Sarah classifies user inputs into (3) categories: positive, negative and neutral. The virtual peer's question topic along with the category of the user's input are then used to associate a discourse act to the input. This is illustrated in the partial mapping in Figure 11.

Based on the identified discourse act, the outcome of the story is determined. As each event is unraveled, Sarah generates the story text by going through each story episode, each story event and each story dialogue, and sources additional knowledge from the domain knowledge and the story world model.

The domain knowledge provides details to the story, such as the use of an *outdoor toy* during a *play* event, or to elaborate health-related information, such as defining *lungs* as a *pair of packet-like organs inside your chest*. On the other hand, the story world model is used to keep the story context in place, such as *playing* an *outdoor toy* is performed in the *park* or explaining how Liam may have gotten *pneumonia* by *playing* without *resting*.

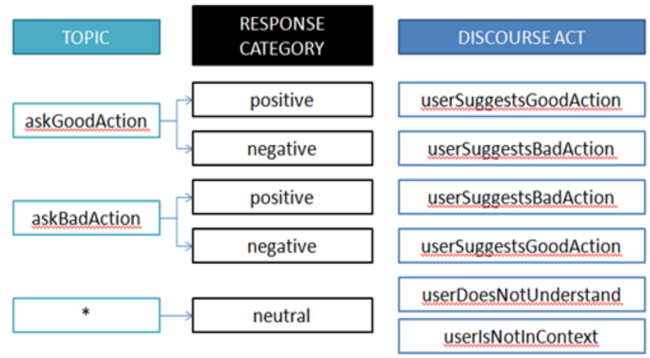

Figure 11. Partial Mapping of Topic-Pattern to Discourse Acts

The planner also supports variations in the generated story with the use of predefined templates and by combining sources from both the domain knowledge base and the story world model. For example, indicating a specific *outdoor toy* as an object during a *play* event can generate any of the story text in Listing 1.

Listing 1. Sample variants of generated story text

Variant #1	Playing with sand in the park
Variant #2	Playing with the ball in the park
Variant #3	Playing in the park

To accomplish this, the planner looks at all the current state of entities in the story world model to match them against event preconditions and aggregating the results from the candidate templates. The planner then tries to match the results again until it can form a detailed story event based on the different sentence patterns defined for that event.

5. Dialogue Generation

A dialogue is a written passage of conversation between two or more characters directed towards a particular subject. It can be used as a mechanism to keep the story moving forward, especially if a new information comes to light [20]. Dialogues also add the illusion of reality to stories and help bring characters to life. They can enhance characterization by revealing things that cannot be understood otherwise. Furthermore, the characters' interactions through dialogues serve as an avenue for establishing and maintaining character-to-character relationships.

To establish user-peer relationships, informational dialogues can be utilized to maintain the user's attention while the virtual peer delivers all necessary health information [21]. But to do so, the informational dialogue should be augmented with social and relational dialogue to manifest a caring and empathic virtual peer. This includes calling the user by name, engaging in social chat or small talks at the beginning of each interaction, providing appropriate feedback at every opportunity, and referencing information that has been discussed from past interactions to give a sense of continuity.

5.1 Types of Dialogues

Sarah uses text-based dialogues as a mechanism to keep users informed of changes in the story, such as the occurrence of an event, or the change in state of a story element (another character, object or location). The use of dialogues to describe what is happening in the story provides an alternative approach to using long descriptive narratives that tend to be very dull when it is not accompanied by any action or character interaction [3]. Dialogues may also help Sarah appear less intimidating, and bring forth a form of characterization to enhance the user's affinity to the peer.

Sarah supports text-based persuasion and inquiry dialogues. Adapting Walton and Macagno's [22] definition, persuasive dialogues aim to resolve conflicting opinions and to reach a compromise given differing interests, while inquiry dialogues aim to jointly prove a proposition given a subject. Persuasive dialogues are used when Sarah communicates with Liam to encourage him to perform certain actions, while inquiry dialogues are used when Sarah asks the user for suggestions on the actions that Liam should take. Examples of these are shown in Listing 2.

Listing 2. Examples of Persuasion and Inquiry Dialogues

Sarah	Liam, your body became dehydrated from dancing music. I think you should drink water so that your body will be hydrated. (persuasion)
Liam	But I don't want to drink some water. My body is just a bit dehydrated so it's okay. I can still plant some seeds. (inquiry)
Sarah	Liam should drink some water, right Hannah? (Hannah is the user)

5.2 Processing User Inputs

Users have limited interactions with the virtual peers. They cannot interrupt the peers' dialogue anytime, and they can only respond when Sarah or Liam asks a question directed towards them. Users are then expected to respond in free-form text using the space provided at the interface (see Figure 12).

Figure 12. Space for User's Free-form Text Input

Sarah breaks down the free-form text input representing a user's contribution to the story by performing normalization and applying regular expression patterns to detect keywords using the AI Markup Language (AIML). Normalization transforms the text input to the same case, removes non-alphanumeric characters and expands contractions. Through regular expression patterns defined in AIML, keywords are identified and extracted from the user's input and mapped as user decisions. Words like *clean*, *drink* and *eat* make up the positive set among several other keywords, while words like *ignore* and *run* imply a disregard for Sarah's suggested proper healthcare practice.

Sarah is also able to categorize indirect responses such as "*I do not think so"* through a set of negation keywords such as *not* and *never*. Moreover, topics are defined in the AIML patterns to keep recognized keywords aligned with the story context. This eliminates confusion when processing generic responses like "*I agree"* that may correspond to different user decision outcomes such as "*Do you think Liam should eat sandwich without washing his hands first?"* (a negative user decision) and "*Do you think that Liam should first wash his hands before he eats his sandwich?"* (a positive user decision).

The discourse act design presented in Figure 11 enables Sarah to appropriately categorize user inputs. For example, when Sarah asks "*How do you feel?*" and the user responds with "*I'm feeling great,*" the input is mapped to the discourse act *userFeelsGood* which is considered a positive response. On the other hand, if Sarah asks *"Do you think Liam should play?"* and the user answers with *"great,"* the response would not be mapped to a specific discourse act because this specific question-response combination is not defined in its AIML engine.

5.3 Response Generation

Using the keywords, AIML classifies the user input as either positive or negative, depending on the action or activity the user has suggested Liam to take. The positive or negative classification dictates the content of the dialogue that Sarah formulates in response to the user's input. To generate Sarah's dialogue turn, messages representing patterns of sentence tags are used. For instance, in Figure 13, the message pattern:

```
expressWorry + suggestRest + stateSickness
```

aggregates three sentence tags. Each tag in turn has multiple sentence structures to allow Sarah to convey the same message but in varying sentence forms. Consider the *suggestRest* tag; it has three possible sentence structures, e.g., "*You should rest.*", "*I think you need to sleep for a while.*", and "*It would be best for you to rest now rather than play.*" The sample dialogue turn that is constructed contains a message that allows Sarah to express concern over Liam's negative health status (portraying a caring persona), to provide a recommended action, and to briefly describe the ailment (portraying the teacher role).

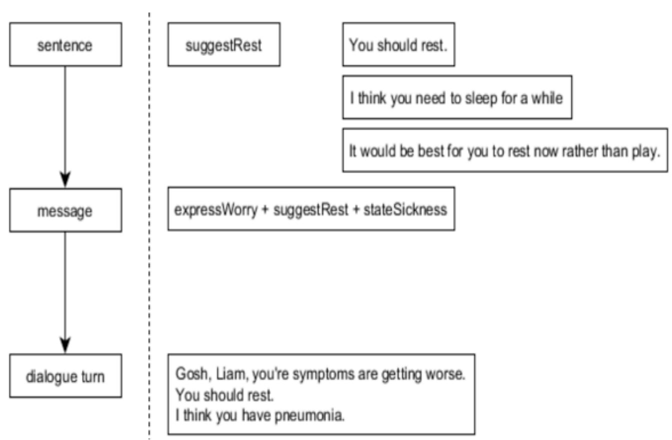

Figure 13. Generating Dialogue Turn

6. Test Results

Three types of validation were conducted to evaluate the content and linguistic features of the generated healthcare stories, the messages of the virtual peers, and the facilities for supporting the interaction between the virtual peers and the user. Participants include linguists and healthcare experts, and children from 7 to 10 years old.

6.1 Expert Evaluation

Linguists and healthcare experts reviewed the generated stories to determine their correctness and appropriateness for the target audience. The linguists found inconsistencies in spelling, errors in grammar, faulty sentence structures, inappropriate tone, and inappropriate content. These feedback were then used to adjust the words and sentence structures used by the virtual peers when conversing and sharing stories with their users. Table 1 provides examples of linguistic errors and how they were corrected to improve the use of language by the virtual peers.

Table 1. Revisions to the use of language of the peers

Text with language errors	Revised text
Sarah: "Since the kids washed their hands, they can now eat sandwich."	Incorrect verb tense. Replace with: "Since the kids *had* washed their hands, …" or "Since the kids *have already* washed their hands, …"
Liam: "It is so delicious to eat sandwich."	"This sandwich is so delicious."
Liam: "But for now, what can I do to stop measles from affecting my nose?"	"What can I do to get better from measles?"
Liam: "See, Anna thinks the same way as me!"	"See, Anna thinks the same way as I do!"
Liam: "Really now? Fine. I'm resting."	"But I don't want to rest."

It was also suggested that story text such as those in Listing 3 should be avoided. For line #1, only symptoms which can be physically noticed or felt by children should be mentioned. Two issues are raised for the dialogue shown in line #2. While the nose may be where the virus came in, it is not the only body part that is affected by measles. Furthermore, if the body part is common (e.g., nose), then there is no need to describe it as this is taught to children as early as their preschool years. Sarah does not distinguish between body parts that children may or may not know, thus, it follows its typical approach which is to generate story text to describe where the body part is.

Listing 3. Inappropriate story content

Line #1	"Liam, I noticed that your white spots inside the mouth are getting worse. You should rest. I think you have measles."
Line #2	"The cause of your measles is virus infection. You can get sick if they get inside your nose. Nose is in the middle of your face."

Healthcare experts, on the other hand, provided recommendations on the removal and replacement of complex concepts and inappropriate assertions on healthcare practices found in the domain knowledge base. Table 2 illustrates these concepts.

6.2 End-User Evaluation

30 English-literate grade school children from private schools with a basic knowledge on how to use desktop computers participated in evaluating Sarah. The validation process commenced with the parents or guardians being given an overview of the software, the validation procedure and the research ethics protocol in order to solicit their informed consent on the participation of their children in the study and the recording of the conversation log in text files. They were also the ones who described the software to the children and asked if the latter are interested in trying out Sarah.

Table 2. Revisions to the domain knowledge base

Original KB content	Recommendations from Child Healthcare Experts
(For cough and colds) Use cough drops / throat sprays / cold medicines	Never advise children to take medicine without doctor's prescription. Instead, use "*See a doctor*" or "*Drink the medicine prescribed by the doctor*"
(For chicken pox) Take painkillers for fever	"Wear loose clothing" or "Prevent itching and scratching"
(For urinary tract infection) Drink antibiotics	"Consult a doctor", "Drink plenty of water", or "Wipe from front to end after peeing"
(For stomachache) Avoid drinking raw milk / eating raw meat or shellfish	"Eat plain biscuits", "Avoid drinking milk", "Do not eat salty food", "Avoid spicy food", or "Avoid oily food"

Then, each child who expressed interest is given an orientation on how to interact with Sarah and the role of Liam. Once the child is ready, the storytelling session begins. Each session lasted from 30 – 45 minutes, depending on the reading and typing speeds of the child. A post-debriefing interview was then conducted to solicit the child's quantitative and qualitative assessment of his/her interaction with the virtual peers. A scale of 1 to 5, with 1 corresponding to *very bad*, and 5 corresponding *to very good*, is used for the quantitative assessment. Table 3 summarizes the evaluation from two iterations of testing.

Table 3. Evaluation scores for peer interaction

Evaluation Criteria	Round #1	Round #2	Average
Sarah was friendly.	3.64	4.38	4.01
Liam was friendly.	4.57	4.63	4.60
Sarah allowed me to join in the storytelling.	4.43	4.38	4.41
Liam performed my suggested actions.	4.50	4.44	4.47
I liked being asked for inputs.	4.43	4.56	4.50
Average	**4.31**	**4.48**	**4.40**

Interaction refers to user's perception of the extent of his/her participation in the storytelling session and engagement during the conversation with the virtual peer. From the results in Table 3, the children found Liam to be more friendly than Sarah, with Liam receiving an average rating of 4.60 compared to Sarah's 4.01. This may be attributed to Liam's tendency to ask children to play, whereas Sarah always poses health-related questions that children have to answer. These results support the earlier findings reported by Tartaro and Cassell in [23] where children prefer virtual peers who act as their friends rather than an entity who acts as an expert by merely stating the facts. Furthermore, in the study of Chan and Ong [24], they found that the age of the participants affect their perception of the role of the peer. Younger participants perceive peers who "asked so many questions" to portray the role of a teacher rather than a friend.

In terms of participation in the storytelling session, children gave Sarah an average score of 4.41 for allowing them opportunities to contribute to the story. Liam, on the other hand, received an average score of 4.47 for performing the actions suggested by the user. Children liked being asked for inputs and the freedom to enter free-form text because it gives them the chance to simulate a more natural, less-controlled conversation with Sarah. This is opposed to designing the virtual peer to generate a list of candidate actions that limit the choices of the user which are illustrated in the works described in [10] and in the Pizzeria Story [12]. Story text describing Liam's execution of every recommended action given by the user, and visual cues depicting the corresponding outcome, as shown in Figure 14, also contributed to the average evaluation score of 4.50. The visual cues to show symptoms experienced by Liam and the possible ailments gave children the feeling that their inputs are important for the story to progress.

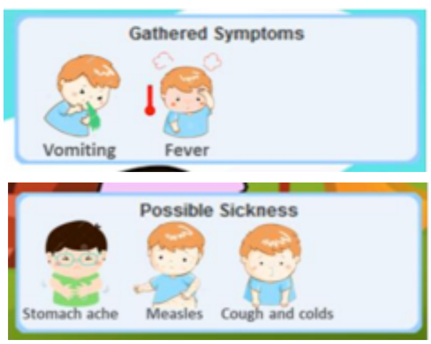

Figure 14. Visual Cues depicting Symptoms and Possible Ailments of Liam

Children also evaluated the interactive healthcare stories shared by Sarah in terms of content and linguistic features. The results are summarized in Table 4.

During the first iteration, children encountered difficulty in some of the words used by Sarah, such as *"relieving"* and *"prevention"*. Children also tend to get confused when they were asked compound questions that are of differing polarity, e.g. *"Liam should rest. Don't you think so?"* This finding correlates to the suggestion from the linguist on the use of simpler, non-compound (positive-negative or negative-positive) questions, and age-appropriate choice of words.

Despite some of these difficulties, the children found the healthcare stories to be interesting, giving a high average score of 4.70. The ability to enter free-form text input may have contributed to this, as the children need to focus on the story that Sarah is sharing in order to be able to give relevant suggestions when prompted.

Table 4. Evaluation scores for the healthcare stories

Evaluation Criteria	Round #1	Round #2	Average
I understood the words used by Sarah.	2.93	4.25	3.59
I understood Sarah's questions.	4.43	4.44	4.43
The story is interesting to me.	4.71	4.69	4.70
I learned what causes a particular sickness.	4.21	4.44	4.33
I learned how to prevent a particular sickness.	3.00	4.13	3.57
I learned the treatment to a particular sickness.	4.36	4.25	4.31
Average	**3.94**	**4.36**	**4.15**

It is worth noting that the modifications made to the lexical choices and sentence structures of Sarah's generated responses led to a significant increase in the average rating of the criterion "*I understood the words used by Sarah*", from 2.93 in the first iteration to 4.25 in the second iteration. However, while children perceived that they learned the causes and treatment for the ailments, they cannot account for learning how to prevent these sickness. This is because the planner did not generate sufficient story text to emphasize prevention practices, whereas in comparison, Sarah prompts children about causes and treatments, which facilitated retention of details.

6.3 Evaluation of Sarah's Responses

During the first iteration of testing, Sarah was programmed to provide a generic response of "*I have no answer for that*" when children's inputs are out-of-topic; that is, responses that do not answer the questions of Sarah. To improve Sarah's image with the children, user inputs were analyzed and categorized. Corresponding AIML templates were added to allow Sarah to generate more appropriate and definite responses to out-of-topic inputs. A spelling checker was also integrated to catch and correct misspellings in the user inputs. Sarah's responses to user inputs were then compared across the two rounds of user testing. The results are shown in Table 5.

Table 5. Comparison of Sarah's responses to user inputs

Category of User Input	Sarah's Response			
	Round #1		Round #2	
	Specific	Generic	Specific	Generic
Yes / No	100%	-	100%	-
In-topic	88%	12%	100%	-
In-topic but with misspellings	-	100%	72%	28%
Out-of-topic	-	100%	91%	9%
Out-of-topic and with misspellings	-	100%	64%	36%
Average	**38%**	**62%**	**85%**	**15%**

Five (5) categories of users inputs were identified based on the types of questions posed by Sarah. The *Yes/No* category requires simple *Yes/No* answers; as these are leading questions, Sarah's response is 100% accurate in both iterations. In-topic responses are those that are contextually-related to Sarah's questions. To respond appropriately, Sarah looks for specific healthcare keywords in the user's input. In the first iteration, no relevant keywords were identified in 12% of the inputs. The addition of more relevant healthcare concepts to the KB increased Sarah's accuracy to 100% in the second iteration.

User inputs frequently contain spelling errors. The use of a spelling checker caught misspellings that were previously undetected, such as "*okayyy*" and "*I'm exited*", thus allowing Sarah to provide specific responses to 72% of user inputs. The flexibility afforded to children also yielded out-of-topic input text, to which Sarah mostly replied with her generic statement. To address this, more AIML patterns were added to map out-of-topic inputs to relevant responses, enabling Sarah to provide 91% specific responses in the second iteration.

Overall, the inclusion of a spell checker and the increase in the number of AIML patterns combined to contribute to a marked decrease of 47% in the number of Sarah's generic responses, from 62% in the first iteration to 15% in the second iteration.

7. Discussion

Sarah is a virtual peer designed to present healthcare facts and good practices to children who are 7 to 10 years old with the use of interactive stories. Children play an active role by engaging in free-form conversations with Sarah regarding the actions that another virtual peer, Liam, should perform in the virtual story world. The actions are meant to address health-related concerns in order to attain and maintain a positive physical state for Liam. Positive, in this case, pertains to freedom from any symptoms that may lead to ailments.

The effect that this setup had on the children's engagement in the interactive storytelling is reflected in their evaluation of Sarah and Liam. Liam's playful persona made him appear friendly because of his tendency to ask children to play. Sarah's more reserved persona, on the other hand, is found to be attuned to her role as a teacher who imparts valuable information about the symptoms, causes, treatment and prevention of common ailments. One child described such interaction with Sarah as "*so educational*" because of the information they gleaned from the story.

The use of free-form text input and more AIML templates to classify user intent also helped in simulating a more natural and less rigid conversation between the virtual peers and the children. However, while the opportunities for inputs from the users are present, these are currently minimal and too short to make any impact on the overall story flow. Story writers identified a number of interaction points that should be considered to further enhance peer-child collaboration. These are indicated in Listing 3.

Results also show that children pose queries for more information about a particular ailment. Such queries can help them gain a better understanding of the unfolding story events. Furthermore, children's affinity with the story characters, particularly Liam, can be seen as an opportunity to encourage them to share their thoughts and feelings related to a story character or situation, similar to how people in real life do through everyday conversations.

In the aspect of story generation, Sarah needs a mechanism to show that an ailment does not happen immediately but may manifest itself through symptoms over the course of several days. The depiction of the transitions in time, e.g., from playing to showing some symptoms of ailments to eventual recovery, should be made more explicit in the narrative expressed in Sarah's dialogue turns. This can provide a more accurate simulation of events as they occur in the virtual story world.

Listing 3. Opportunities for User Input in the Story

Sarah:	It's a sunny Friday morning! Oh! Hello Hannah! I'm Sarah, and there's Liam the one in blue. Liam is excited to play with you!
Liam:	Great! Hi, Hannah! (<u>Suggested interaction point</u>: Hannah should be given a chance to respond.)
Sarah:	Liam, is there anything you want to do?
Liam:	Hannah, let's dance to the music. (<u>Suggested interaction point</u>: *Hannah, would you like to dance to the music?)*
Narrator:	Liam and Hannah are having fun dancing to the music. :
Liam:	Awesome! Now we can plant some seeds. (<u>Suggested interaction point</u>: Hannah*, what seeds should we plant?)*
Narrator:	The kids look like they're having fun planting some seeds.

8. Conclusion

In this paper, we described the use of interactive stories to depict scenarios on the occurrences and treatment of common childhood ailments. The system utilized natural language processing techniques to process children's free-form input text that dictate the actions they want characters to perform in the story world. Template-based story generation was employed in the production of events that flow from problem introduction to resolution. Dialogue turns were designed to align with the context of the story, through the use of discourse acts and AIML patterns. A domain-based KB was also utilized to provide Sarah with the necessary concepts and assertions about healthcare that it can use in generating stories and dialogues. The presence of Liam who shares the same interests as the children contributed to the latter's engagement and enjoyment of the stories. This, combined with giving children the control of what happens in the story, specifically to Liam's physical state, increased the children's involvement in the storytelling process.

Drawing from these findings, future studies can investigate the correlation between the extent of involvement in the storytelling process and the acquisition of important healthcare concepts and practices that children can apply in their daily lives. Increasing the number of interaction points between the children and the virtual peers should also be explored. Aside from responding to questions posed by Sarah, future research should work on enhancing the peer's abilities to perform its role as a storytelling partner by encouraging other forms of discourse with the user, to increase opportunities for learning through the child-agent interaction.

9. REFERENCES

[1] Day, V. 2009. Promoting health literacy through storytelling. *The Online Journal of Issues in Nursing*, 14, 3, Sep 2009.

[2] Benner, P., Sutphen, M., Leonard, V. and Day, L. 2009. *Educating Nurses: A Call for Radical Transformation*. San Francisco, CA: Jossey-Bass.

[3] Osborne, H. 2008. Tools of change: Telling and listening to stories. *On Call Magazine*, October 2008. Available at: https://healthliteracy.com/2008/10/16/telling-stories/

[4] Williams-Brown, S., Baldwin, D.M. and Bakos, A. 2001. Storytelling as a method to teach African-American women breast health information. *Journal of Cancer Education*, 17, 4, 227-230.

[5] Koki, S. 1998. Storytelling: The heart and soul of education. *Pacific Resources for Education and Learning* (PREL) Briefing Paper.

[6] Zabel, M.K. 1991. Storytelling, myths, and folk tales: Strategies for multicultural inclusion. *Journal of Preventing School Failure: Alternative Education for Children and Youth*, 36, 32-34. Tailor & Francis Online.

[7] Lambert, J. 2013. *Digital storytelling: Capturing lives, creating community*. Routledge: New York.

[8] Xu, Y., Park, H.S. and Baek, Y.K. 2011. A new approach toward digital storytelling: An activity focused on writing self-efficacy in a virtual learning environment. *Educational Technology & Society*, 14, 4, 181-191, 2011.

[9] Cassell, J., Ryokai, K. and Vaucelle, C. 2002. Literacy learning by storytelling with a virtual peer. In *Proceedings of Conference on Computer Support for Collaborative Learning: Foundations for a CSCL Community*, 352–360.

[10] Ong, E., Consignado, D.G., Ong, S.J. and Soriano, Z. 2016. Building a semantic ontology for virtual peers in narrative-based environments. In *Trends in Artificial Intelligence - PRICAI 2016 Workshops*, Lecture Notes in Computer Science, 10004, 65-76. Springer.

[11] Mateas, M. and Stern, A. 2003. Façade: An experiment in building a fully-realized interactive drama. In *Proceedings of the Game Developers Conference*.

[12] Yu Galan, S., Ramos, M.J., Dy, A., Kim, Y. and Ong, E. 2017. Towards a narrative-based game environment for simulating business decisions. In N. Nunes, I. Oakley, V. Nisi (Eds,), *Interactive Storytelling*, ICIDS 2017, Lecture Notes in Computer Science, 10690, 290-294. Springer, Cham

[13] Baur, T., Damian, I., Gebhard, P., Porayska-Pomsta, K. and Andre, E. 2014. A job interview simulation: Social cue-based interaction with a virtual character. *IEEE International Conference on Social Computing* (Jan 2014).

[14] Cassell, J., Tartaro, A., Rankin, Y., Oza, V. and Tse, C. 2005. Virtual peers for literacy learning. *Educational Technology*, Special Issue on Pedagogical Agents XLVII, 39–43.

[15] Barker, T. 2003. Collaborative learning with affective artificial study companions in a virtual learning environment. *PhD dissertation*, University of Leeds.

[16] Kordaki, M. 2014. On the design of educational digital stories: The Ed-W Model. *Procedia - Social and Behavioral Sciences*, 116, 1631-1635.

[17] Liu, H. and Singh, P. 2004. ConceptNet - A practical commonsense reasoning tool-kit. *BT Technology Journal*, 22, 4, 211–226.

[18] Ang, K. and Ong, E. 2012. Planning children's stories using agent models. In D. Richards and B.H. Kang (Eds.), Konwledge Management and Acquisition for Intelligent Systems, PKAW 2012, Lecture Notes in Computer Science, 7457, 195-208. Springer Berlin Heidelberg.

[19] Riedl, M. 2004. Narrative generation: Balancing plot and character. *PhD dissertation*, North Carolina State University.

[20] Ramet, A. 2011. *Creative writing*. Hachette UK.

[21] Garrity, T. F. 1981. Medical compliance and the clinician-patient relationship: A review. *Social Science and Medicine, Part E: Medical Psychology*, 15, 3, 215–222. Science Direct, Elsevier Ltd.

[22] Walton, D. and Macagno, F. 2007. Types of dialogue, dialectical relevance and textual congruity. *Anthropology and Philosophy*, 8, 1/2, 101–120.

[23] Tartaro, A. and Cassell, J. 2008. Playing with virtual peers: Bootstrapping contingent discourse in children with autism. In *Proceedings of the 8th International Conference for the Learning Sciences*, 2, 382–389.

[24] Chan, L. and Ong, E. 2018. Engaging Children in Conversations during Story Reading. In *Proceedings of the 26th ICCE Workshop on Innovative Technologies for Enhancing Interactions and Learning Motivation*, November 26-30, 2018. Manila, Philippines.