

Towards Interest and Engagement: A Framework For Adaptive Storytelling

Michael Garber-Barron
Cognitive Science Department
Rensselaer Polytechnic Institute
barrom2@rpi.edu

Mei Si
Cognitive Science Department
Rensselaer Polytechnic Institute
sim@rpi.edu

Abstract

A storyteller builds a narrative that captivates the audience, immersing them in the story. Storytelling is an interactive process. Though the listeners cannot affect what happens in the story, a good narrator observes the audience's responses and adjusts his/her storytelling accordingly. We present an automated storytelling agent that is aimed at achieving the same effect. While presenting a story, the user is given chances to give comments or ask questions. The agent estimates the user's preferences towards various topics from these responses and weighs the factors of novelty, current interest, and consistency for generating the next part of the narration. We describe the components of the agent, and an example of applying it for narrating a Chinese fantasy story.

Introduction

Storytelling is an important aspect of the human experience. As a child, we spend a lot of time listening to stories. Soon, we start to tell stories ourselves.

When storytelling happens face to face, it is a dynamic process. The listener interrupts, asks questions, and gives comments. They also express their experience through facial expressions and body language. The storyteller picks up this feedback and adjusts his/her storytelling accordingly.

When a story is told through media, however, it is often delivered as a canned piece - a book, a movie, or a recording of voices, and the interactive component is missing. In this work, we aim at creating an AI agent which can tell stories interactively, not only by answering the listeners' questions, but also by inferring the listeners' interests and tuning its storytelling towards them.

Of course, the inferred interests cannot be the only factor to consider. Examining novel and unfamiliar content can also increase the listener's feelings of immersion and engagement (Wade 1992). New content can attract a person, leading them to explore different material which can develop alternate interests (Loewenstein 1994). Coupled alongside these notions is the need to balance between novel and familiar topics as reexamining familiar and related material can also develop and deepen a person's preexisting interests (Arnone et al. 2011). Based on these findings, our storytelling agent is designed to engage the user by balancing

Copyright © 2012, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

novel content with a deeper exploration of the inferred preferences and interests of the user.

Related Work

Several AI systems have been proposed for storytelling that take into account the user's experience. The Papous system (Silva, Raimundo, and Paiva 2003) proposes to adapt storytelling based on the listener's emotional experiences. The Virtual Storyteller (Theune et al. 2003) uses autonomous agents to interact with each other and present a story to the audience. The interaction is steered by the positive or negative feedback from the audience members. Similar to Papous, the storybook system (SSAU) (Tanenbaum and Tomizu 2008) explored adapting storytelling to the user's mood. In IRST (Nakasone, Prendinger, and Ishizuka 2009), the story is developed around the notion of interest. A fixed story is presented to the user, but optional information can be presented to add details about characters or events.

Like IRST, we also try to adapt a story around the interests and preferences of the user. In this work, we go further by trying to balance between providing the user more material along the lines of where he/she has already shown interest and presenting materials that can potentially trigger new interest from the user. We also developed algorithms for automatically inferring the user's interest from his/her questions and comments.

Example Domain

We use as our example domain the Chinese fantasy tale - "The Painted Skin" (Sung-Ling 1916) to demonstrate the application of our agent.

The story follows the meeting of a man - Wang - and a demon disguised as a girl who eventually murders Wang.

This storytelling scenario occurs in a virtual Chinese tea house created using the Unity game engine (Figure 1) and is intended as part of a larger project on using a virtual environment to teach Mandarin Chinese to college students. The narrator only speaks Chinese. Translations are provided as needed.



Figure 1: Virtual Environment

Agent Architecture

We want to allow the user to explore content they already find appealing as well as to discover new interests. To create this effect, the agent performs in an iterative fashion. As the story proceeds the user is given opportunities from time to time to ask questions or comment on the transpiring events. The user does so by selecting from a list of statements displayed on the screen. This feedback updates the agent’s model of the user’s interests through the profiling module. The event selection module determines the content to present to the user by taking into account his/her interests and the progression of the story.

Before presenting how these modules work, we will describe how the story is represented in our design.

Story Representation

The story is composed of a series of connected events. An event describes the characters’ background, beliefs, and motivations. For example, an event can be an explanation for Wang’s wife’s insistence that the girl leave their home or Wang’s emotional response to realizing the girl is a demon. Each event corresponds to a segment of text that is spoken by the narrator.

To ensure that the story can only progress in a logical and coherent manner, events have preconditions. Preconditions are defined as a combination of events that can and cannot have already happened in the story.

For enabling the storytelling agent to *understand* the story and infer the user’s interests, the events described by the narrator and the statements the user can say to the narrator are tagged with labels designed by the author. The labels allows us to provide a comparable semantic representation of the story. For example, the sentence “Wang’s wife’s response was out of fear that the girl would take her place” is labeled with *Wang, Fear, Wife, Love*. This is represented as a vector where if a label is associated with the sentence, the corresponding position of the label is marked with the value of one.

Interest Profiles

The user’s interests are represented as a weighted linear combination of multiple user profiles. Each *profile* is defined as a vector representing the user’s possible interest towards each of the labels. Table 1 depicts two profiles we used for modeling “The Painted Skin” story. The values associated

with each label of a profile are randomly assigned a value at initialization. We used twenty-four labels in total for modeling this story. Only five labels are listed in the table because of space constraints.

The reason we want to use multiple profiles is to allow for each profile to represent a different prediction of what materials (labels) are most interesting to the user and only get updated along this direction. The preference values are used to weight the multiple profiles and in doing so represent their relative importance to the listeners.

		Profile 1	Profile 2
	Preference Value	25%	75%
Label Values	Home	0.24	0.07
	Gain	0.78	0.00
	Fear	0.14	0.08
	Demon	0.16	0.1
	Harm	0.75	1.00

Table 1: Example Profiles

More specifically, when the user selects a statement, its corresponding label vector is compared to each profile, and used to update the most similar one – the profile with the smallest Euclidean distance from the label vector. Including the statement’s label vector, the number of times each label has appeared in the story is recorded. The profile’s label values are then updated by recalculating the average occurrence of each label. This is done by dividing their occurrence in the profile by the number of times the profile has been chosen by the agent. Furthermore, this value is then weighted by a novelty bias so that less common labels are given relatively greater increases in value. This is done by multiplying the label value by $1 + (1 - \text{label value})$. At the same time, the preference value is increased. Figure 2 provides an example of this process.

Event Selection

The narrator agent selects the next event to present by evaluating the future trajectories following each of the current candidate events. The candidate events are those events whose preconditions are satisfied and therefore can immediately follow in the story. Similarly, their future trajectories are formed by selecting events that can follow the current event.

Each trajectory is assigned a score based on how likely it will engage/interest the user by comparing the events in it with each of the profiles and obtaining an overall score by summing the values generated in each comparison. Algorithm 1 provides the pseudo code for this process. The process for evaluating a trajectory takes into account each profile along with two factors: novelty and interest consistency.

Novelty: When calculating the score of an event, a novelty metric is applied by increasing the influence of less emphasized labels in the scoring process. As future events along a trajectory are compared to the profiles, those events that contain labels that have occurred less frequently in the profiles

Algorithm 1 EventEvaluation(EventTrajectory, Profiles)

```
1: PreviousProfile = null
2: for each Event in EventTrajectory do
3:   Profile = GetClosest(Profiles, Event)
4:   Distance = GetDistance(Profile, Event)
5:
6:   #EventLocInStory: The event's location in the story.
7:   #AvgStoryLength: The average length the story can take.
8:   # Events farther along in the story with less common labels
   are given a greater impact.
9:   ProfileScore[Profile] +=
   | 1 -  $\frac{\text{EventLocInStory} - \text{Distance}}{\text{AvgStoryLength}}$  |
10:  if PreviousProfile == Profile then
11:    Consistency[Profile] +=  $\frac{1.0}{\text{EventTrajectory.Length}}$ 
12:    PreviousProfile = Profile
13:  for each Profile in Profiles do
14:    Score +=
    Consistency[Profile]*ProfileScore[Profile]*
    Profile.PreferenceValue
15:  Return Score
16:
17: #GetDistance(Profile, Event): Returns the distance between
   a profile and an event.
18: #GetClosest(Profiles, Event): Returns the profile with
   the smallest distance to an event.
```

are weighted more positively in the scoring process (Algorithm 1, Line 9). This causes the agent to incorporate events with more uncommon and novel labels.

Consistency of Interests: Rather than oscillate between topics, the agent prefers trajectories that fully explore one topic before moving to the next. This is realized by giving higher scores to events in succession that are similar to the same associated profile (Algorithm 1, Line 11).

The final value of the whole trajectory is calculated as the weighted sum of the user's preference towards each interest profile with the values generated for the events relative to that profile (Algorithm 1, Line 14). Once a value for each candidate event is generated, the agent chooses the event with the highest calculated value to present.

Example Process

In this section we provide an example detailing the event selection and profile update process. Figure 2 presents a successful example of the agent incorporating user feedback to determine between whether to discuss the demon's intention to murder Wang or to elaborate on the wife's actions.

Discussion and Future Work

In this work we have proposed a computational approach for adaptive storytelling. Our AI agent is able to adapt the narrative of the story by inferring the interest's of the user, and also accounting for aspects of novelty, and the consistency of the user's interests during storytelling. The agent's design has been implemented and has been used to tell an example story. Formal evaluation has been planned to evaluate the effectiveness of our agent towards engaging the audiences

Current Profiles:

Profile 1: Home 0.24 Fear 0.14 Harm 0.75 Wife 0.00 Lust 0.50
Profile 2: Home 0.50 Fear 0.79 Harm 0.77 Wife 0.22 Lust 0.31
Profile 1 Preference Value: 0, Profile 2 Preference Value: 0

User selects a statement with the labels: *Wife, Harm*.

Profile 1 distance: 1.25, Profile 2 distance: 1.01

Profile 2 is chosen, and its Preference Value is incremented.

Profile 1 Preference Value: 0, Profile 2 Preference Value: **1.0**

The values of the corresponding labels in Profile 2 were increased to suggest their consistency with the user's preferences. We apply a novelty bias which is in favor of content that is less seen in the story. The details of the two algorithms cannot be presented here due to space limitations.

Profile 1: Home 0.24 Fear 0.14 Harm 0.75 Wife 0.0 Lust 0.50

Profile 2: Home 0.50 Fear 0.79 **Harm 1.0 Wife 1.0** Lust 0.31

The agent then scores the possible events to present, assuming there are 2 candidate events with the following labels:

Event 1 Labels: Fear, Harm, Wife, Girl

Event 2 Labels: Wang, Fear, Wife, Love

Calculating the scores using Algorithm 1. The value i refers to the profile ID.

Score =

$$\sum_{i=1} \text{Consistency}_i * \text{Profile_Score}_i * \text{Preference_Value}_i$$

$$\text{Event 1 Score} = (.17 * .25 * 0) + (.31 * .41 * 1) = 0.13$$

$$\text{Event 2 Score} = (.26 * .34 * 0) + (.36 * .30 * 1) = 0.11$$

The agent selects Event 1.

Figure 2: Example Process

of an automated storyteller. We also plan to explore automated approaches for labeling statements in a story, and for assigning the initial weights of the labels in each profile in a meaningful way.

References

- Arnone, M.; Small, R.; Chauncey, S.; and McKenna, P. 2011. Curiosity, interest and engagement in technology-pervasive learning environments: a new research agenda. *Educational Technology Research and Development* 59:181–198.
- Loewenstein, G. 1994. The psychology of curiosity: A review and reinterpretation. *Psychological Bulletin* 116:75–98.
- Nakasone, A.; Prendinger, H.; and Ishizuka, M. 2009. Isrst: Generating interesting multimedia stories on the web. *Applied Artificial Intelligence* 23(7):633–679.
- Silva, A.; Raimundo, G.; and Paiva, A. 2003. Tell me that bit again... bringing interactivity to a virtual storyteller. In *Proceedings of the 2nd International Conference on Virtual Storytelling*, 146–154.