

Tell a Story About Anything

Mei Si^(✉)

Rensselaer Polytechnic Institute, Troy, NY 12180, USA

sim@rpi.edu

Abstract. With the fast development of internet technology, people can have easy access to a massive amount of information. The goal of this project is to provide a personal assistant for helping people explore large network of information by using narrative technologies. We propose an automated narration system that takes structured information and tailors the presentation to the user. It is aimed at presenting the information as an interesting and meaningful story by taking into consideration a combination of factors including topic consistency, novelty, user interests, and the user's preferences in exploration style. We present preliminary results of using this system for presenting information about the 2008 Summer Olympics Games, followed by discussion and future work.

Keywords: Data exploration · Novelty · User preferences

1 Introduction and Motivation

The development of internet technologies has enabled people to have access to a massive amount of information. For example, there are more than 4 million entries currently in the English Wikipedia. Unlike a traditional dictionary, these articles share many links among each other and form a huge network of information. Similarly, an incredible amount of information is being generated on Twitter and other social media at every moment. These posts are connected with each other by shared topics, user names, etc. A key difference between exploring such a large network of information and reading a book or an article is that there is not a clear thread for how the reading should proceed. Because the information is highly interconnected, a reader can be easily distracted by new topics. As a result, the structure in the data may become less obvious to the reader, and the reader may feel overwhelmed more easily.

To provide the user a more structured experience of information exploration while not restricting the user, we propose a narrative based agent with an integrated visualization tool as a personalized assistant. Visualization has been used widely for illustrating relationships in data. Similarly, narrative is a powerful tool for helping people understand and organize information [1, 5, 6]. In this work, an automated narrative agent is developed to take structured data and present them piece by piece to the user. It proactively constructs the narratives using information relevant to the user's interests, and also takes into consideration a combination of factors including topic consistency, novelty and the user's preferences for exploration style.

2 Example Domain and User Interface

We will demonstrate the application of our systems by introducing information about the 2008 Summer Olympic Games in Beijing, China. We represent the network of information as topics and their relationships. Each topic is treated as a node in a graph, which has a description and links to other nodes. For example, the node “Table Tennis Competition” and the node “Peking University Gymnasium” are linked by the “happened” relationship. Currently, the example domain for this work includes about 200 nodes and 1400 links among the nodes, representing 27 types of relationships. This information is kept in an XML file.




Fig. 1. Interactive visualization interface

Figure 1 shows a snapshot of the end-user’s interface. This interface is synchronized with the narrative agent. In the center of the screen, the current topic is displayed. All of the topics that are linked to the current topic are shown in a circle around it. The user can mouse over a topic to get detailed information about it. Without input from the user, the next topic is decided by the narrative agent. The user can click on any topic on the screen to make it the current topic, and the narrative agent will reconstruct its presentation accordingly. To help the user explore the knowledge base, on the right side of screen two lists of topics are suggested. One list contains the topics the narrative agent is most likely to talk about next based on the history of the interaction and the user’s preferred information exploration style. The other list contains the topics that are considered as most novel to the user based on the history of the interaction. Using a set of slider bars, the user can indicate his/her preferred information exploration style by changing the relative weights of the factors that the narrative agent uses for picking the next topics, such as topic consistency and novelty. The algorithms of how the narrative agent works are described in Sect. 3. Finally, the interface allows the user to directly ask questions either by voice or text. Currently, we only support simple queries such as “Who is Michael Phelps?” by matching the user’s input with predefined templates.

3 Storytelling by Balancing Objectives

Starting from any point in the knowledge base, the narrative agent can present the domain by introducing the topics one by one. The agent strives to go over

the topics in a meaningful order, while being attentive to the user's interests. The agent decides what to present next by balancing a number of objectives. For the domain of introducing the 2008 Olympics games, the agent has four objectives, which consist of maintaining hierarchical ordering consistency in the description, maintaining spatial ordering consistency in the description, adhering to the user's interests, and introducing more novel topics. These objectives may affect the narrative agent's behavior in different directions. For example, the system's current topic is introducing the Beijing National Aquatics Center. At the same time, there are novel topics the system can present, such as the US tennis players. Should the agent move its current topic to the tennis players or find a topic closely related to the Beijing National Aquatics Center? In our system, each objective is weighted with a relative importance. By default, we give high importance to those objectives that arrange the topics in a meaningful order. For our domain, that includes hierarchical ordering consistency and spatial ordering consistency. The user can change the weights at any time in the user interface. The agent picks its next topic by maximizing its achievement of these four objectives with their relative importance factored in. Next, we briefly describe how each of the objectives is evaluated.

First, for knowledge that is structured hierarchically, the agent prefers to talk about topics at the same hierarchical level together. For example, the 2008 Beijing Olympic Games have multiple venues, such as the Bird's Nest and the Water Cube, and within each venue there are several subtopics. For keeping hierarchical ordering consistency, the agent will not mix topics from different hierarchical levels, nor jump between subtopics that belong to different topics at a higher level.

Similarly, when describing a spatial environment, people typically follow an order, e.g. clockwise instead of randomly jumping around [4]. For keeping spatial ordering consistency, the agent ensures that the spatial relationship between the new topic and the current topic, and that between the current topic and the previous topic are the same.

For both maintaining hierarchical ordering consistency and spatial ordering consistency, a topic can receive a score of 1, 0 or -1. 1 means the ordering consistency is kept; -1 means the ordering consistency is violated; and 0 means the consistency check does not apply.

Thirdly, the agent wants to take the user's interest into consideration. Currently, we use a simple heuristics for estimating the user's interest – if the user has asked about a topic in the previous five steps, the user is interested in the topic and its closely related ones. The agent uses breadth first search for computing the distance between every topic in the knowledge base and the user's most recently queried topic. The shorter the distance is, the more coherent the topic is to the user's interest. If the user has not asked a question in the previous five steps, then 0 will be returned.

Finally, we want to keep the user engaged. In our previous work on creating digital storytellers, we found many users enjoyed the stories more when new topics were introduced from time to time, e.g. a new person or item [2]. We try

to do the same thing here by requesting the agent to pick novel topics to present to the user. This objective also serves the purpose of allowing the user to explore the knowledge base more efficiently because the agent will present a wider range of topics to the user. We evaluate the novelty of a topic by considering whether the topic has been presented before, what percentage of the topic's immediately related topics have been presented before, and the distance between the topic and the current topic.

An example is provided below. In this example, all of the four objectives have the same weight for the simplicity of the presentation. For computing novelty, the three factors we consider are weighted as -10, 1, and 1 respectively, i.e. if a topic has been presented before, it is usually not regarded as novel.

Node: 2008 Summer Olympic Games
 Hierarchical Ordering Consistency: 0
 Spatial Ordering Consistency: 0
 Adhere to User's Interest: 0.8
 Have been Presented Before: 1 (true)
 Percentage of Related Topics Presented: 0.6
 Distance from Current Topic: 0.2
 Novelty ($-1*10 - 0.6*1 + 0.2*1$): -10.4
 Total Score (Hierarchical Ordering Consistency*1 + Spatial Ordering Consistency*1 + Adhere to User's Interest*1 + Novelty*1): -9.6

4 Discussion and Future Work

Narrative and dialogue are good ways for engaging people and helping people organize and memorize information. This project is aimed at helping people explore and consume a large network of information in a narrative form. A preliminary version of the system has been implemented and is presented in this paper.

For future work, we want to improve the algorithms for evaluating the objectives. In particular, we currently calculate the distance between two topics simply using breadth first search. This distance may not represent the distance of the two concepts in the user's mind. In the future, we will consider using a semantic tool such as concept net [3] to help us get a better estimation. Secondly, we want to explore how to automatically create the knowledge base by using existing data extracting tools such as DBpedia. This will allow us to have larger and richer domains for testing our algorithms. Having more domains encoded will also help us design a more general set of objectives for the narrative agent. Finally, we want to evaluate our system and test whether it can in fact help people explore information more effectively and whether people enjoy using it.

References

1. Abbott, H.P.: The Cambridge introduction to narrative. Cambridge University Press, Cambridge (2008)
2. Garber-Barron, M., Si, M.: Adaptive storytelling through user understanding. In: Ninth Artificial Intelligence and Interactive Digital Entertainment Conference (2013)
3. Liu, H., Singh, P.: ConceptNet-a practical commonsense reasoning tool-kit. *BT Technol. J.* **22**(4), 211–226 (2004)
4. Taylor, H.A., Tversky, B.: Perspective in spatial descriptions. *J. Mem. Lang.* **35**(3), 371–391 (1996)
5. Thorndyke, P.W.: Cognitive structures in comprehension and memory of narrative discourse. *Cogn. Psychol.* **9**(1), 77–110 (1977)
6. Wilkens, T., Hughes, A., Wildemuth, B.M., Marchionini, G.: The role of narrative in understanding digital video: an exploratory analysis. In: The Annual Meeting of the American Society for Information Science, pp. 323–329 (2003)