The Victorian Laptop: Narrative Engagement through Place and Time

J. Cassell & J. Smith
MIT Media Lab
E15-315, 20 Ames Street
Cambridge, MA 02139
{justine, jsmith}@media.mit.edu

Abstract

We introduce a class of systems that motivate and aid users to think and write about their personal experiences by presenting them with similar first-person narratives, which allows them to reflect on the connection between their experience and the story presented. These systems rely upon the narrative intelligence of humans to establish the relevance connections, and so can make use of story indexing techniques which do not require the system to have full story understanding. We present the Victorian Laptop, an implementation of such a system in the domain of travelogue. This system is designed to give users a richer connection to a place that they are visiting, and deeper engagement with their own experience of that site, by presenting them with first-person travelogues, which allows them to reflect on the connection between their own travel and the travel of others who have come before.

Introduction

Narrative is the primary form through which we understand and make meaning of our experience. It is in telling the narrative of our lives (to ourselves and to others) that we construe our role in the world. Speaking or writing about personal experiences is one way to gain a deeper understanding of the world and our relation to it. Throughout time, people have written about their travel as an important way of placing themselves in the world, and understanding the story of their personal development with respect to that world. In recording a story on site which documents our impressions of places we visit and people we meet, we push ourselves to understand those events more fully. Or, if we are recording after the fact, we relive those experiences and reflect upon them and how they can be integrated with the rest of our experience and knowledge. In creating technologies to participate in the act of storytelling, therefore, we argue for systems which "stand back" and encourage a user to be creative and actively engaged with the world, rather than substituting in any way for that engagement.

Reading personal stories made by others can also encourage reflection. This is particularly the case when we experience a sense of connection or identification with the story. This can happen when the events in the story are similar to events we have experienced, when the places in the story are familiar to us, when we feel we have something in common with the author of the story. This sense of personal connection encourages active engagement with the story, reflection on its meaning and on our own experience, which brings a deeper understanding of both the story and our own experience. Thus, a system which helps create that personal connection can make the experience of reading and writing more active, richer and more engaging.

This ability to make experience into stories and to actively understand others' stories is what we call human narrative intelligence. An important component of this intelligence is the ability to create meaningful links between pieces of discourse. It is a basic human assumption that discourse is meaningful and coherent, and so when the links are not obvious from the text itself, we create meanings that will fill in the blanks and maintain coherence. (Hobbs 1979) This creation of coherence links occurs at the inter-sentence level, but also at the larger level of linking stories. For example, in traditional storytelling activities, one person may tell a story about the history of the group as a response to a story about personal experience told by a group member. It is the job of the listeners to find the (sometimes quite inferentially challenging) relationships between the two stories. And it is in the process of understanding the relationship between the two stories that a new "story" is created, that relies on the shared features, and inferences about human experience and cultural values.

The system approach described in this paper is motivated by a desire to give users the valuable and satisfying experience of reflecting on their experience through writing the story of their own lives, and reading stories which are relevant to their own experience. This approach uses indexing techniques and natural language processing to understand a user's story on some level, and then uses that understanding to select from a database a story which is similar to the user's story in a way which will promote a sense of personal connection with the selected story.

We present the Victorian Laptop, a system we have built following this approach, whose domain is travelogue. This

domain has the advantage that experience of keeping a travel journal or writing letters home, is familiar to many people. It also provides promising opportunities for a sense of personal connection, since first-person stories about travelling are rich in descriptions of place and experience. In addition, many people are stimulated by reading the writings of other travelers, past and present, since it allows them to travel in their minds to other times, and to understand in an immediate way the history of a place that they are just discovering. The interface for the system we present here is an antique lapdesk writing box, an important component of a context conducive to writing about place and thinking about time.

In this paper, we first review work in areas with related goals and techniques. We then outline the basic system approach and describe its implementation in the Victorian Laptop. Finally, we discuss our ongoing research and future directions.

Related Work

The goal of the Victorian Laptop is to engage a user in the act of travel writing, and to deepen that experience through providing similar stories to reflect upon. The challenge of "providing similar stories" stands at the intersection of two major fields: information retrieval and narrative understanding, and this section reviews related work from those fields. From the field of narrative understanding we review research in case-based reasoning and modeling human story understanding. We also discuss standard statistical information retrieval and the application of natural language processing to information retrieval. Finally, we review a story-matching system for children.

Our system has some commonalities with systems developed by Shank, Domeshek, and others, working in Case-Based Reasoning (CBR). A core observation underlying this approach is that people often learn from other people by exchanging stories which are related in some way. In order to model this process, systems have been developed in which a large number of stories is stored and indexed using a complex indexing scheme. This allows those systems to match input stories with a story which is similar in a way that is relevant to the domain. For example, Domeshek (1996) built a system called ABBY which takes in a story about a lovelorn user's relationship problem, and uses an indexing system to select a story from its database to offer as social advice.

In a similar vein, some systems have been developed which aim to model human creative understanding of stories. These approaches are interesting to us since a system which could understand both its own stories and the user's stories on some level may be able to make a good match between them. Moorman and Ram (1993) have reviewed several of these systems and proposed a set of six tasks involved in creatively understanding stories. This delineation can help us clarify which aspects of story understanding we wish to implement. Of those six tasks, the most relevant for this task are "language understanding", which includes syntactic processing, "story structure understanding", which includes tasks such as character identification and genre identification, and perhaps "episodic understanding", which identifies actions and the agents involved in them. It is important to keep in mind when deciding which of these tasks to implement that the lower the level of understanding, the broader the domain to which it can be applied with equal computational effort. The field of information retrieval also has much to offer this project. Information retrieval systems select documents from a database for a user based on a user's information need, as expressed in a query. While we seek to offer our users inspiration rather than information, the goal of finding a match between pieces of text is very relevant. The standard method of information retrieval uses statistics computed on the set of all terms (often referred to as a "bag of words") in a query and document to find a relevant match, and several components of this technique are useful for us. One of the systems we have tested is the Okapi (Robertson 1995) standard IR retrieval algorithm, as implemented in Rhodes' Remembrance Agent (1996). Other researchers in this field are investigating the application of natural language processing techniques to information retrieval. This can be done at the lexical, syntactic, semantic, or discourse level of language processing (Smeaton 1991). Comparisons have indicated that adding NLP techniques to information retrieval often results in an increase in precision (Strzalkowski and Sparck Jones 1996). Since our system only returns one story at a time, precision is more important than recall, so adding NLP techniques is a very promising direction.

SAGE (Bers & Cassell 1998) has a motivation very similar to ours. This system is intended to help children explore their identity by prompting them to tell personal stories and offering traditional inspirational stories in return as comfort and material for further reflection. SAGE also uses a simple keyword matching algorithm to select stories from its database to offer to the user. SAGE differs from our system, however, in that the interaction is framed by an embodied conversational character who engages the child in conversation and relates the selected story. In addition, the character explicitly prompts the user to find a connection between the two stories and reflect on that connection. This prompting, together with the rich context and social script created by the presence of the embodied character, helps the user make a connection between her own story and the reply story. However, the important lesson that we draw from SAGE is that the system itself need not have human-like narrative intelligence. In the domain of storytelling, we can - and should -- rely on the human's own ability to derive

narrative coherence to complement the intelligence of the system.

System Design

In this section, we describe the system approach in general, and how it is implemented in the current Victorian Laptop. There are four important components: interface, input collection, story matching, and database.

Interface

The physical interface draws the user in and sets the mood for the interaction, and so it must induce a feeling of connection with the story domain and be aesthetically pleasing to use. This applies to both the method of entering a story and the appearance of displayed response stories.

The Victorian Laptop is housed in an antique mahogany Victorian lapdesk writing box. The lapdesk serves as an introduction to the historicity of the system, and also a lighthearted allusion to the lapdesk -- laptop lineage. The lower half of the open box contains a digitizing tablet with paper laid over it. The user composes travel stories on the paper using a real pen which doubles as a mouse. The upper half of the box contains a tablet computer whose screen displays the user's cumulative writings and the output stories. This screen is made to look like an open book with aged paper, to evoke the feeling of reading an old, precious diary. The user's story is displayed on the left page in a modern font, while the response stories are displayed on the right page in an old-fashioned font. Figure 1 contains a photograph of the Victorian Laptop.

Figure 1. The Victorian Laptop physical interface.

Input

Input to the system is composed of text from the user and contextual information. The system keeps a record of all the user's input in a given session to build a discourse history. The most recent part of the story is weighted more strongly when matching, but the entire discourse history is also an important part of the input. This aids in developing cohesion across the entire interaction. Examples of contextual information which can be used are location (using GPS), orientation (digital compass), time of day or year, and information about the user (previous locations visited, and so forth).

Database

The database consists of a collection of first-person historical writings. For the moment, the documents (diaries, surveyors' notes, writings from the WPA writers' project, and so on) have been segmented into travel stories. If the database is stored in a central place and the devices are networked in the future, the corpus could also include stories written by others during their interactions with the system. While our current implementation emphasizes automatic processing of the texts to find keywords, we do not rule out the possibility of pre-processing the stories according to an indexing scheme as used in CBR systems.

Matching

Matching is triggered when the input module sends a signal to the matcher, requesting a reponse. In the current system, this signal is issued directly by the user, who taps an on-screen button to elicit a response story. However, this is not our ideal triggering method. In future, we plan to continuously monitor the input story, checking for matches, and return a matching story once the strength of the match exceeds a given threshold. We may also use a length cutoff as a trigger for matching, so that a response will arrive after a given number of sentences.

The matching module selects an appropriate story from the database for presentation to the user, based on an analysis of all the inputs. In general terms, the story selected should be similar to the user's story. For example, it may be about the same place the user has described. It may describe similar weather, time of day, or time of year. Matching on the length of the story is also important (Grice 1989). The current system uses a keyword-matching algorithm. The list of keywords includes a set of words determined to be key by statistical comparison with a large reference corpus (Guardian). In addition, the list includes all the proper nouns in the corpus, since we believe connection with a place is important in the domain of travel stories. Each keyword is also assigned an importance weight. At this early stage, we are only beginning to experiment with weighting keywords differently. We have obtained encouraging results by giving a higher weight to words associated with travel, such as "arrive" and "visit". This direction is discussed further in the next section.

In the first step of matching, keywords in the user's story are extracted and stemmed following Porter (1980). Response stories containing keywords are then found, and the story with the highest cumulative weight (normalized by length) is selected.

The current system, since it analyzes only at the word level, produces mixed results. Gough's (1990) observations about the relevance of conversational contributions explains this and points the way for improvement. He distinguishes between "inter-contextual" links, in which a story is related to the conversational domain, and "subcontextual" links, in which a story is tied only to a non-topic word of the previous story, and notes that a participant who makes too many of the latter type of link is percevied as rambling. In the following section, we discuss our current research on improving the system, as well as specific directions for improvement in matching.

Example

Below is an extract from an actual interaction with the Victorian Laptop. The user typed in the first paragraph, and then paused, which led the system to respond with output from Amelia Burnett. The user continued his story and when he again paused the system contributed a segment of a story by Eva Blackwelder. The input and output were displayed on the upper screen of the laptop as described above.

Input:

We (a co-worker and I) flew in late Thursday. A quick, aggressive cab ride from Logan Airport to the Copley Square Hotel didn't leave me with much in the way of initial impressions. Freeways and traffic, the scourges of every American city. At least the Bostonians are trying to hide theirs with the "Big Dig" -- a ten-billion-dollar project to bury all the expressways under the city.

Output:

Amelia Burnett, visiting from Fayal, Portugal, in 1846-1848 excerpts from letters written journal form to her parents.

Saturday 31st Oct., at 5 A.M. a sailor from the mast's head, called out 'Land' delightful sound, and it was amusing to watch the general movement which followed the intelligence. I sat on deck till dark, when I could perceive the distant light on Thatcher's island off Cape Ann. We went to our berths early, but at half-past 10 I was woke up, as I had requested, by the mate, to see Boston Light which had just appeared in sight. I went on deck, where Clara and Olivia joined me, and there remained until we cast anchor in the stream close to the wharf, at 1-1/2 Sunday morning.

Input:

The Copley Square hotel is small and pleasant -- nice, but not quite charming enough to be called quaint. The rooms are functional but the hotel is definitely lacking some of the amenities of its larger brethren. $160 wasn't exactly a bargain -- but then again, it wasn't too bad.

Output:
Eva Blackwelder, tour of Massachusetts, 1938-1939.

We finally went to the hotel Westminster ($3.50) in Copley Square across from the Copley Plaza Hotel, a gorgeous hotel taking up a whole square block. (rates $7-8-12) We expected a F.H.A. a rate at the Westminster.

Future Work

A primary research focus is the creation of a story-matching system. We believe that this will fall between two extremes. At one extreme, a story could be selected completely at random from the corpus. At the other extreme, the system would fully model human narrative intelligence and select a story in the way that conversing humans do. It will be helpful to explain why we do not consider either of these extremes appropriate for this task.

In the SAGE system, Bers & Cassell (1998) investigated children's ability to find a connection between their own story and the system's reply story. She found that children were equally able to make a rich connection when the stories were selected at random, and when they were selected based on her matching algorithm. Clearly, humans have narrative intelligence and are able to create meaning out of almost anything (Grice 1989). However, the present system differs from SAGE along several dimensions which make random stories less effective for us. As described above, in the SAGE system the presence of a character representing a wise storyteller provides a rich social script which prepares the user to interpret a story which is relevant, even if ambiguous at first glance. In addition, the type of stories SAGE was intended to include, such as parables and myths, are designed to lend themselves easily to many different interpretations and applications. Travelogues, on the other hand, are much more specific and so require more effort to connect coherently with other stories.

However, the fact that random stories were so effective in SAGE indicates that a full knowledge-based story understanding system may not be necessary, or even desirable in this system. The true narrative intelligence lies in the user; our job is to support it. The stories must be close enough that the user can make the cohesive link without great effort, but the very act of making that connection is an important component of the satisfying experience we wish to support. Since humans are so adept at making meaning and connections, a complex indexing system which models deep aspects of human understanding, while we do not rule it out on principle, is likely to be more artificial intelligence than we need. In

addition, it is difficult to see how such a system could be implemented in the near future. Our system would require the input story to be understood in real time. With a domain as broad and vague as travelogue, the knowledge base required to fully understand the story could be prohibitively large.

Given that neither random stories not an expert system approach is ideal for this system, what would work well? We would like to borrow aspects of various IR, NLP, and story understanding approaches and combine them into a workable system. In particular, one fruitful direction that we are pursuing derives from the indexing problem of case-based reasoning. There are certain common themes in travel stories (such as arrival, visiting friends, visiting landmarks, travel romance), and receiving a story with the same theme as yours could be very stimulating. Instead of using complex frames or other indexing structures, these sub-genres may be identifiable on the basis of surface features, such as word choice, verb tense, or distribution of parts of speech. Guesses about the theme of an input story on the basis of such features may improve story selection and the discovery of such feature sets is the subject of ongoing investigation.

In the future, our work can also be improved by the addition of NLP techniques, which can increase precision. At the syntactic level, keyword matching could be improved by an analysis which extracts the main subject of each sentence, or head-modifier pairs. At the discourse level, the story could be analyzed for style, genre, or mood. Anaphoric resolution might also improve retrieval, although there is evidence that this does not improve IR as much as expected (Smeaton 1991). At the semantic level, the stories could be coded in frames or other semantic structures, as Domeshek (1992) did in his implementation of ABBY.

Finally, we would like to incorporate more contextual information into the input. A GPS and digital compass linked with a map would tell the system what the user might be looking at, and use this to weight stories about that place more highly. Information about the time of day, time of year, or weather could also be used to maximize the similarity between the user's current context and the historical writer's context, increasing the sense of connection with that person.

Conclusion

We have presented an approach to the role of technology in storytelling that relies on the technology to situate the act of storytelling in its social, site-specific and temporal context, but relies on the user's own narrative intelligence to make story meaning. The Victorian Laptop is an example of such a "story listening" system, that situates users' travel writing in the larger historical context, and encourages them to reflect deeply about the site they are visiting, and the temporal context in which their own lives fit.

Acknowledgements

We would like to thank Petra Chong and Jason Bau for their work on implementation and research, Timothy Bickmore for feedback on this paper, and the members of the Gesture and Narrative Language Group for their collaboration and help.

References

Bers, M. and J. Cassell. 1998. Interactive Storytelling Systems for Children: Using Technology to Explore Language and Identity. *Journal of Interactive Learning Research* 9: 183-215.

Domeshek, E. 1992. Do the Right Thing: A Component Theory for Indexing Stories as Social Advice. PhD Thesis, Northwestern University. Technical Report #26.

Grice, P. 1989. *Studies in the Way of Words.* Cambridge: Harvard University Press.

Gough, D. 1990. The Principle of Relevance and the Production of Discourse: Evidence from Xhosa Folk Narrative. In Britton, B & Pellegrini, A. (eds.*) Narrative Thought and Narrative Language*. NJ: Lawrence Erlbaum.

Hobbs, J. 1979. Coherence and Coreference. *Cognitive Science* 3: 67-90.

Moorman, K. and Ram, A. 1993. A New Perspective on Story Understanding. In *Proceedings of the 31st Annual ACM Southeast Conference.* Birmingham, AL

Porter, M.F. 1980. An Algorithm For Suffix Stripping. *Program* 14(3): 130-137.

Rhodes, B. and Starner, T. (1996) Remembrance Agent: A Continuously Running Automated Information Retrieval System. In *Proceedings of The First International Conference on the Practical Application of Intelligent Agents and Multi Agent Technology (PAAM '96),* 487-495.

Robertson, S. E., Walker, S., Jones, S., Hancock-Beaulieu, M. M., Gatford, M. 1995. Okapi at TREC-3. In *Proceedings of the Third Text REtrieval Conference (TREC-3),* 500-226. NIST Special Publication.

Shank, R. and Riesbeck, C. 1981. *Inside Computer understanding: Five Programs Plus Miniatures.* NJ: Lawrence Erlbaum.

Smeaton, A.F. 1991. Progress in the Application of Natural Language Processing to Information Retrieval Tasks. *Computer Journal* 35 (3): 268-278.

Strzalkowski, T., Lin, F., Guthrie, J. L., Leistensnider, J., Wilding, J., Karlgren, J., Straszheim, T., Perez-Carballo, J. 1996. Natural Language Information Retrieval: TREC-5 Report. In *Proceedings of the Fifth Text REtrieval Conference (TREC-5),* 500-238. NIST Special Publication.

Strzalkowski, T and Sparck Jones, K. 1996. NLP Track at TREC-5. In *Proceedings of the Fifth Text REtrieval Conference (TREC-5),* 500-238. NIST Special Publication.

Tannen, D. 1984. *Conversational Style: Analyzing Talk Among Friends.* Norwood, NJ: Ablex Publishing Corporation.