

Say Anything: A Demonstration of Open Domain Interactive Digital Storytelling

Reid Swanson and Andrew S. Gordon

The Institute for Creative Technologies, The University of Southern California
13274 Fiji Way, Marina del Rey, CA, 90292 USA
{swansonr, gordon}@ict.usc.edu

Abstract. Say Anything is a text-based interactive digital storytelling application that differs from other systems in its emphasis on the ability of users to create a narrative in any domain that they wish. The user and computer take turns in writing sentences in an emerging fictional narrative where sentences contributed by the computer are selected from a collection of millions of personal stories extracted from Internet weblogs. In this demonstration, we will present the latest version of the Say Anything application and allow conference participants to author their own original stories using the system.

Key words: interactive storytelling, interactive fiction, collaborative writing, social media, weblogs, information retrieval.

1 Say Anything

Say Anything [4] is an open domain story generation engine that is unlike most other contemporary interactive storytelling systems. Say Anything diverges from current systems in two key areas. First, most state-of-the-art systems rely heavily on rich, 3D graphics to help engage the user with the story and virtual environment. However, similar to the early interactive storytelling systems such as TALE-SPIN [3], the virtual world in Say Anything is brought to life entirely through textual language. Second, in any interactive storytelling system there are two somewhat contradictory objectives: to maintain a strong coherent narrative and to allow the user the freedom to do anything they please. While the ultimate goal is to simultaneously achieve both, most compelling systems have focused on improving the narrative aspect by restricting the user's ability to interact with the world in various ways. Say Anything on the other hand emphasizes the ability for the user to create a narrative in any domain without any restrictions, other than the limits of a keyboard.

Stories are authored with Say Anything in a turn-based fashion where the human user and the computer alternate writing sentences of the developing narrative. The user always begins a story with the first sentence. The system then analyzes what has been written and returns a sentence that tries to continue the story in a coherent and entertaining way. If the user is satisfied with this sentence they may continue by writing another sentence. However, if the sentence does not make sense or is

objectionable for another reason they may click on the returned sentence and be presented with nine other alternative sentences from which they can choose. This alternating writing process continues until the user feels that the story has reached a natural conclusion or has deviated so far off track that recovering a coherent narrative is impossible. Fig. 1 presents a screenshot of the system's primary authoring interface. The following passage is an example of a successful story written in collaboration with the system:

The weather broke, so we sailed out of the harbor. As Victoria grew nearer, the waves grew larger and we furled some foresail and turned to run. We sailed at about 9 knots with good trim, but the storm eventually caught up with us. With its big open cockpit and heavy nose, I didn't like its chances in the kind of sea you get out there almost continuously that time of year. Sure enough the boat was completely inadequate, and we were tossed into the cold ocean. Everyone in our group of seven tourists -- five locals and a Japanese couple -- was pretty excited about the experience. The Japanese couple were the ones that saved us though, with their expert swimming abilities. as far as that goes it was just the four of us. The last tourist was lost at sea, never to be found. Drowned or murdered, the bloated, stinking bodies that turn up by the hundreds will look much the same. Such is the way with storms like that!



Fig. 1. The main writing interface for the *Say Anything* application. The user can write a new sentence, swap one computer-contributed sentence for another, or allow the computer to continue generating sentences.

The system continues a story using a simple three-step architecture. First it analyzes lexical and semantic features of the user's developing story so far, including the most recently contributed sentence. This includes several preprocessing steps that attempt to identify named-entities, remove stop words and perform other transformations to help clean the text. This analysis is used to find a story containing a sufficiently similar sentence to the user's current input from a large corpus of millions of English language stories collected from Internet weblogs [1]. Similarity between sentences is assessed using standard information retrieval techniques implemented in the Apache Lucene search engine toolkit [2]. Several methods for determining similarity have been explored, however, the most successful so far is simply a vector space model with *term frequency-inverse document frequency* feature weights, where terms are both unigram and bigrams. The story retrieved using this technique is treated as a proxy for the user's emerging story and the next sentence in that story is returned as the continuation of the user's narrative.

There are relatively few restrictions on the way the user can interact with our system or in the types of stories they can create with it. The only stated rule is that the user should take turns writing one sentence at a time. However, there are no explicit mechanisms to enforce this rule and users may in fact type as much text as they would like (up to a predefined character limit) into the text box. Whether it is due to the stated rule, the design of the user interface or because it is a natural unit, we have found that users typically only write one sentence per turn. Using the current retrieval model, the average story is just over 14 sentences, or 7 turns in length. When rated on a 5-point scale (1 low; 5 high), by the author and other users, these stories were given a 3.66 coherence rating and a 3.42 rating overall. In comparison a retrieval model that randomly picks sentences from the corpus received a 2.93 coherence rating and a 2.66 overall rating. Although few restrictions are built into the current system, it may be interesting in the future to enforce rules about the length of content, genre and other structural or semantic aspects of the process, as is common in other (non-computational) collaborative writing games.

In addition to supplying the user with a fun improvisational game-like environment to write creative stories through their interaction with the system, Say Anything also provides a research platform for studying various aspects of computational linguistics and narrative coherence. While the information retrieval techniques work surprisingly well, the simple lexical features used to compute similarity do not capture nearly enough knowledge to make the right decision a large percentage of the time. However, the natural interaction with the system (i.e. when users select alternate sentences) provides valuable feedback about the quality of our model. One part of our current research is focused on using this trace data to learn a secondary model using a richer set of syntactic, semantic and discourse features, which can more accurately assess the quality of a candidate sentence to be returned to the user.

This platform also enables us to explore the related issue of narrative adaptation. If we had an infinitely large corpus of stories to draw from, it would always be possible to find a story sufficiently similar to the user's developing narrative. However, even with the vast amount of text written on the web, there will never be a story collection large enough to match every lexical variation in the way that people compose narratives of experiences. One method to address this issue is by recognizing when significant discontinuities would occur and adapting the retrieved content to eliminate

mismatches. For example, it may be possible to recognize when a retrieved story differs from the user's story in the gender of the central characters, and then automatically adapt the personal pronouns of the retrieved content so that it can be included in a coherent way. The user interaction data that we are collecting will help us investigate how to automatically determine which elements in a candidate sentence do not align well with the user's story, and what possible replacements would most improve its quality.

In this demonstration we will allow users to author their own original stories in real time and explore the latest features of the system. In addition the users will be able to explore the other aspects of the system that have been developed to try to foster a web-based community around this type of simple storytelling game.

Acknowledgments. The project or effort described here has been sponsored by the U.S. Army Research, Development, and Engineering Command (RDECOM). Statements and opinions expressed do not necessarily reflect the position or the policy of the United States Government, and no official endorsement should be inferred.

References

1. Gordon, A., Cao, Q., Swanson, R.: Automated Story Capture From Internet Weblogs. In: Proceedings of the Fourth International Conference on Knowledge Capture, Whistler, BC, Canada, pp. 167-168, ACM Press, New York, NY (2007)
2. Gospodnetic, O., Hatcher, E.: Lucene in Action. Greenwich, CT: Manning Publications (2004)
3. Meehan, J.R.: The metanovel: writing stories by computer. PhD thesis, Yale University (1976)
4. Swanson, R., Gordon, A.: Say Anything: A Massively Collaborative Open Domain Story Writing Companion. In: Spierling, U., Szilas, N. (eds.) ICIDS 2008. LNCS, vol 5334, pp. 32-40, Spinger, Heidelberg (2008)