Towards Narrative Information Systems

Philipp Wille, Christoph Lofi, Wolf-Tilo Balke

Institut für Informationssysteme, TU Braunschweig, Germany
{wille, lofi, balke}@ifis.cs.tu-bs.de

Abstract. Narrative interfaces promise to improve the user experience of interacting with information systems by adapting a powerful communication concept which also comes natural in human interaction. In this paper, we outline how such a narrative information system which answers queries using elaborate stories can be realized. We show how to construct query-dependent plot graphs from unstructured data sources. Also, we conduct an extended user study on our prototype implementation in which users adapt and personalize the plots of query-dependent stories. These studies provide further insights into which parts of a story are relevant and can be used as a starting point for either personalizing stories or for crafting better user-independent stories in later works.

Keywords: Narrative Information Systems, Narrative Exploration, Storytelling

1 Introduction

Information systems form the backbone of most Web applications enabling access to a vast amount of digital information. However, interfaces to this information still lack intuitiveness in the sense of how queries and query results are communicated. Especially when faced with complex and multi-facetted queries like “What will diabetes mean for my life?” many state-of-the-art techniques like keyword-based document retrieval are still insufficient as users have to laboriously work through multiple documents. Thus, a core challenge in current information systems research is naturally providing relevant information as opposed to retrieving relevant documents [1].

Stories are particularly well suited to communicate complex and multi-facetted information in an easy-to-understand and easy-to-memorize fashion. They exhibit a plot that covers a series of events in a comprehensible way to help the audience understand their connections and as research in cognitive sciences suggests, this type of packaging information is especially easy for people to digest. The most memorable stories however are interactively tailored for the audience, and a good storyteller takes feedback and directive cues of her audience under consideration.

In this paper, we focus on the core-problem of building a Narrative Information System, namely the construction and interpretation of *query-dependent plot graphs*. To illustrate this concept, figure 1 shows a sample plot graph for a query about *diabetes*. This graph tells an overview story of the topic, briefly touching related topics like *obesity* and *prediabetes* and relating them to the query topic through sub-topics such as *diagnosis* and *causes*. While the shown plot graph looks similar to knowledge

WAIM 2015, p. X, 2015.

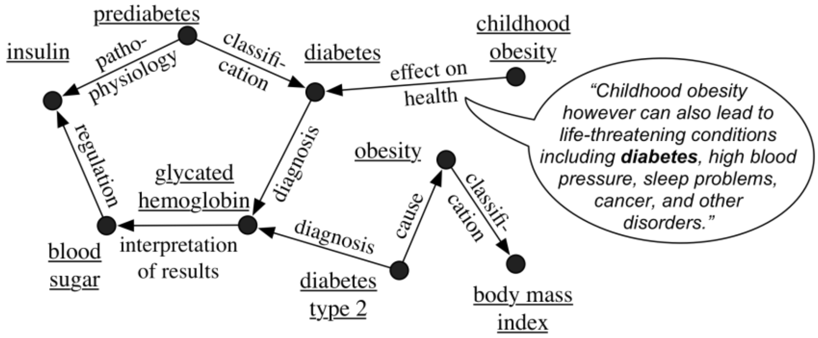

Figure 1: Simplified example plot graph for a general query about *diabetes*.

graphs as for example in [2, 3], the core difference is that the shown plot graph can be transformed into a textual representation (story) via *story synthesis* (the problem of finding a suitable linearization of the plot graph) and has text fragments for each node or edge available for this purpose.

2 Modelling Narrative Queries

Plot graphs have been a common representation for stories in previous works [4, 5], however, in contrast to other recent publications on storytelling ([2, 6]), we do not limit plots to just temporally coherent narrations (e.g. news stories) or causal event chains. Instead, we define them as sets of covered topics that may be connected by sub-topics. To help with *story synthesis*, we annotate each node and edge with one or more *text snippets* representing brief descriptions or explanations.

A plot graph G_P is given by $G_P = (S, T, lb_S, snip_S, lb_T, snip_T)$ with S being a finite set of nodes (representing topics), and T being a finite set of edges $T \subseteq S \times S$ representing transitions (via a sub-topic). Furthermore, we have a finite set L_S of node labels (topic names), and a finite set L_T of edge labels. Then $lb_S \subseteq S \times L_S$ is a relation assigning labels to nodes, and $lb_T \subseteq T \times L_T$ is a relation assigning labels to transitions (therefore allowing multiple edges with different labels between two topics). Finally, we have a set of text snippets TS, and the relations $snip_S \subseteq S \times TS$ and $snip_T \subseteq T \times TS$ assigning snippets to nodes (topics) and transitions (sub-topics).

Our system considers only a finite number of different plot graphs which is restricted by the data sources available to it. This *universe of discourse* containing all topics and transitions which could be used in plots constructed by our system is represented by a *universe graph* G_U that is defined analogously to definition 1 with each possible plot graph G_P being an edge-induced subgraph of G_U. Given a universe graph, a user can search its collection by issuing a query featuring a topic T_C. Based on that central topic, a subgraph of the universe graph is selected as a plot graph that contains topics and transitions *relevant with respect to the query topic*.

Since users try to satisfy a fuzzy information need in facetted-search, our system cannot provide one correct answer for a query, but instead has to generate an individual plot graph for each individual user. This is done by interactively expanding a plot graph based on user feedback. We refer to such an interactively generated plot graph

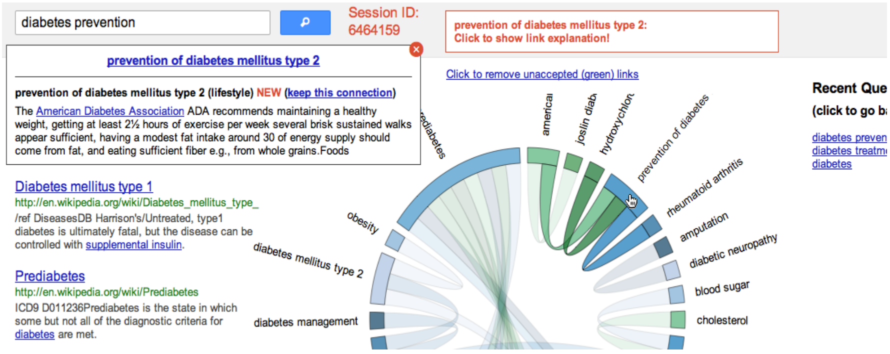

Figure 2: Screenshot of the user interface of *Narrex* as used in our user study

as a query-dependent *discourse graph* G_D that contains only nodes and edges that are *relevant* to a query Q and the user. It is also an edge-induced subgraph of G_U.

Based on discourse graphs, we implemented a Narrative Exploration prototype called *Narrex* (see figure 2). As a representative corpus, we used the full English Wikipedia (as of September 4, 2013). We chose all Wikipedia articles as valid topics and took all sensible first-level headings of each article as sub-topics (e.g. see figure 1). Topic and sub-topic text snippets were selected using the multi-aspect summarization method proposed by Song et.al. [1] that generates topic- as well as sub-topic-specific summaries for a given query topic. Each sub-topic specific summary text contains relevant topics related to the query topic via the sub-topic that can be found with automated link detection [7]. A discourse graph built this way contains the query topic as seed node and all relevant related topics as additional nodes connected by its sub-topics. All topics and sub-topics are annotated with their respective summaries. Using the user interface as shown in figure 2, users of Narrex can interactively and narratively explore the underlying text corpus.

3 User Study & Conclusions

We conducted an extensive user study with 100 native English speakers on the crowdsourcing platform CrowdFlower. Since we aim at improving the intuitiveness of Information Systems dealing with multi-facetted queries, we asked the participants to educate themselves about *the most important life changes a person has to make in case that he/she was diagnosed with diabetes*, a query that stems from the facetted search task of the TREC 2009 web track data set and that is quite representative for what we envision our system to be able to do as there are several different stories from different perspectives that could be told to satisfy a user's information need. The participants were split in two groups with the first using Narrex and the second working on a state-of-the-art keyword-based system. We examined the following issues: (1) what subsets of the discourse graph represent good plot graphs of stories users actually consider interesting and helpful, (2) how would real humans tell the story of a plot graph, i.e. how would they translate a plot graph into a textual representation, (3)

Table 1. Group means (M/SE) for the questions of the third experiment.

	Question a)	Question b)	Question c)	Question d)
Narrex	2.91 / 0.27	2.44 / 0.20	2.71 / 0.26	2.91 / 0.23
Keyword-based	2.87 / 0.26	2.61 / 0.28	2.71 / 0.25	3.03 / 0.28

how informative are the query-dependent story graphs generated in this way for answering a query compared to having verbose information available in form of relevant full documents (i.e. the information obtainable by keyword search).

In the first experiment we studied how our study participants built plot graphs with Narrex to get an understanding which stories are considered good and which topics and transitions are deemed relevant by logging their interactions with Narrex. The generated individual plot graphs differed strongly in the set of contained story items and transitions. However, all these plots share a core set of transitions and attached topics considered relevant by many participants that tells a generic plot providing an overview of the task, whereas the less chosen transitions and items belong to personalized plots satisfying a more specific interest.

In a second experiment, we asked the users to write short texts of 100 to 200 words summarizing their gained knowledge from the search task in order to get an intuition on how a human storyteller would synthesize a story. A group of 6 expert annotators rated the quality and usefulness of each story provided by participants of both groups. In general, stories from Narrex were assigned similar quality (50.46 in average on a scale between 0 and 100) to those resulting from keyword search (51.12 in average) although their authors had access to significantly less texts than the second group that worked on full text documents returned by keyword search.

After completing experiment 2, participants using Narrex as well as those using standard keyword search were asked in a third experiment to rate the following four questions on a scale from 1 (best) to 6 (worst), so we could measure the user satisfaction with either of the systems: (a) how satisfied were you with the information display, (b) to what extend could you answer the main question, (c) how satisfied are you with the knowledge you got in relation to the main question, (d) can you imagine using this type of information display to answer questions like “how does living on my own will affect my life”, “how does my life change if I got pregnant”, or “how does my life change if my wife is diagnosed with cancer”.

The results are shown in table 1. We analyzed the user ratings for each question individually with group means. To compare group means we used an independent sample t-test. The conditions of normal distribution are assumed due to a sample size of at least 30 in each group, homogeneity of variances was tested (Levene’s test). Our results show that even naïve prototypes of Narrative Information Systems can already keep up with keyword search for multi-facetted queries. Also, they dramatically reduce the cognitive burden by directly presenting relevant information in form of relationships between topics and by reducing the texts the user has to read. These findings suggest that further research in Narrative Information Systems is worthwhile.

References

1. Song, W.; Yu, Q.; Xu, Z.; Liu, T; Li, S.; Wen, J.-R.: Multi-aspect Query Summarization by Composite Query. In: 35th Int. ACM SIGIR, ACM, New York, 2012; pp. 325-334.
2. Shahaf, D.; Guestrin, C.; Horvitz, E.: Trains of Thought: Generating Information Maps. In: 21st Int. Conf. on World Wide Web (WWW), Lyon, 2012; pp. 899-908.
3. Shahaf, D.; Guestrin, C.; Horvitz, E.: Metro Maps of Science. In: 18th ACM SIGKDD Int. Conf. on Knowledge Discovery and Data Mining (KDD), Beijing, 2012; pp. 1122-1130.
4. Riedl, M. O.; Thue, D.; Bulitko, V.: Game AI as Storytelling. In: Artificial Intelligence for Computer Games, Springer, New York, 2011; pp. 125-150.
5. Kasch, N.; Oates, T.: Mining Script-like Structures from the Web. In: 1st Int. Workshop on Formalisms and Methodology for Learning by Reading, LA, California, 2010; pp. 34-42.
6. Radinsky, K.; Davidovich, S.; Markovitch, S.: Learning Causality for News Events Prediction. In Proc. 21st Int. Conf. on World Wide Web (WWW), Lyon, 2012; pp. 909-918.
7. Milne, D.; Witten, I. H.: Learning to Link with Wikipedia. In: 17th Conf. on Information and Knowledge Management (CIKM), Napa Valley, California, 2008; pp. 509-518.