The AIIDE-17 Workshop on
Intelligent Narrative Technologies
WS-17-20

Towards Intelligent Narrative-Based Interfaces for Sensemaking

James Niehaus, Scott Neal Reilly, Peter Weyhrauch, James Tittle, R. Michael Young
Charles River Analytics, Inc., University of Utah
jniehaus@cra.com, snealreilly@cra.com, young@cs.utah.edu, jtittle@cra.com, pweyhrauch@cra.com

Abstract

Professionals such as military analysts, medical providers, and business administrators have a significant need to understand and make sense of an increasing quantity of data. We propose intelligent narrative-based interfaces as a means of organizing, presenting, and interacting with this data to improve user comprehension. This paper identifies critical questions and potential research paths in intelligent narrative-based interfaces for sensemaking. We propose a general architecture of intelligent narrative-based interfaces and identify central research questions in this area. We believe this new area of research has significant potential to use narrative technologies to impact a wide range of real world problems.

Introduction

The wealth of data available to modern professionals far exceeds their ability to effectively process it and make decisions. This data comes from an increasingly massive collection of widely available digital content, such as social media, blogs, and professional publications; new arrays of sensors that collect and report audio, video, environmental and other data; and increasingly sophisticated analysis systems that transform this data into new information. Professionals such as military analysts, medical providers, and business administrators have a significant need to understand and make sense of this data through the help of automated analytics, visualizations, and interfaces.

Narrative is a natural and prevalent form of communication that is used daily by the vast majority of society to relate, summarize, and emphasize information to make sense of the world (Bruner, 1991). For the purposes of this work, we define narrative broadly as the presentation of a causally related set of events, including supporting information. They can convey complex data and facts in straightforward structures that are easily comprehended. Narratives provide temporal, spatial, and intentional information about events in a format often more easily understood and better recalled than flat text or data (Graesser, Singer, & Trabasso, 1994), and this information is exactly the type of information that professionals need to make sound decisions from data. Because of these advantages, narratives have significant potential to provide an organizing structure to aid professionals in processing and understanding data, through narrative-based interfaces for sensemaking. Tools that can automatically collect, organize, and present information in narrative formats are extremely advantageous, providing the benefits of narrative structure at a low per-analysis cost. In this paper, we propose a general architecture for such tools and identify some of the research challenges posed in the development of such an architecture.

Related Work

Some research has addressed the creation of narratives from sports and video game play data. The StatsMonkey system (Allen et al., 2014) automatically generated narratives from raw baseball game numerical data, and some of the authors continue to offer intelligent narrative generation services from data (seemingly, in text outputs) with the Quill product and patent (Birbaum et al.) by the company Narrative Science. Bouayad-Agha et al. (2011; 2012) select content and generate narrative summaries of football matches from game information scraped from websites, providing the ability to include perspectives in the summaries. Lareau et al. (2011) provide a method for detecting event sequences that are likely to be of high interest from Australian Football League game statistics. Gervas (2013) has completed similar work in selecting content for narratives in chess games. Summarization has also been applied to video gameplay

This material is based upon work supported by the United States Air Force (USAF) under Contract No. FA8750-17-C-0018. Any opinions, findings and conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of the United States Air Force. Predictive System for Cyber Hostility using Integrated Computational Models (PSyCHIC) is supported by the Office of the Director of National Intelligence (ODNI) and the Intelligence Advanced Research Projects Activity (IARPA) via the Air Force Research Laboratory (AFRL) contract number FA8750-16-C-0108. The U.S. Government is authorized to reproduce and distribute reprints for Governmental purposes notwithstanding any copyright annotation thereon. Disclaimer: The views and conclusions contained herein are those of the authors and should not be interpreted as necessarily representing the official policies or endorsements, either expressed or implied, of ODNI, IARPA, AFRL, or the U.S. Government.

logs in Antoun et al. (2015), Ryan et al. (2015), Pizzi et al. (2010), and Barot et al. (2017). Cardier (2007) has explored the exploration of conceptual spaces for narrative-based knowledge systems.

Work has been done on the use of planning-based models to characterize structural properties of stories, including of causal relationships (Young, 2010), intentional relationships (Riedl and Young, 2010; Teutenberg and Porteous, 2013), inter-character conflict (Swanson and Jhala 2012; Ware and Young, 2014), belief dynamics (Teutenberg and Porteous, 2015; Thorne and Young, 2017) and intention dynamics (Young, 2017). These planning-based models have been leveraged as the core knowledge representation use in a mixed analytical/generative system called Bardic (Barot et al., 2017) to generate narrative summarizing the large amounts of low-level event log data generated by the online multiplayer games (Valve, 2013).

Other work has used differing models of narrative structure to map from underlying data to story structure, including hierarchical, qualitative models (Molina and Flores, 2013), task models (Miranda et al., 2015) and automatically and manually written association rules (Vaudry and Lapalme, 2016). Recently, manual creation of narrative-based interfaces has become an effective means to communicate complex information (Segel and Heer, 2010), and this technique has been suggested by earlier work (Gershon and Page, 2001). This paper argues for the use of ongoing work in computational narrative to develop systems that automatically collect and process data and generate these narratives.

Intelligent Narrative-Based Interfaces for Sensemaking

As a motivating example, consider a hospital administrator who is attempting to determine the best staffing policies for physicians, surgeons, and nurses in her hospital based on past data. The administrator may pull electronic medical record (EMR) data, staffing reports, and feedback from employees. One method to view this data may be large tables of information with descriptive statistics such as mean patient treatment time and mean patient outcome measures. Another may be charts and graphs that relate data elements visually. These organizations may work well for answering some questions, but may miss elements of the larger context such as how hospital size, administration leadership, technological capabilities, and types of patients seen interact at different periods of the hospital's history. For this analysis, a narrative-based interface that relates the motivations of the administration and staff to causal interactions among events

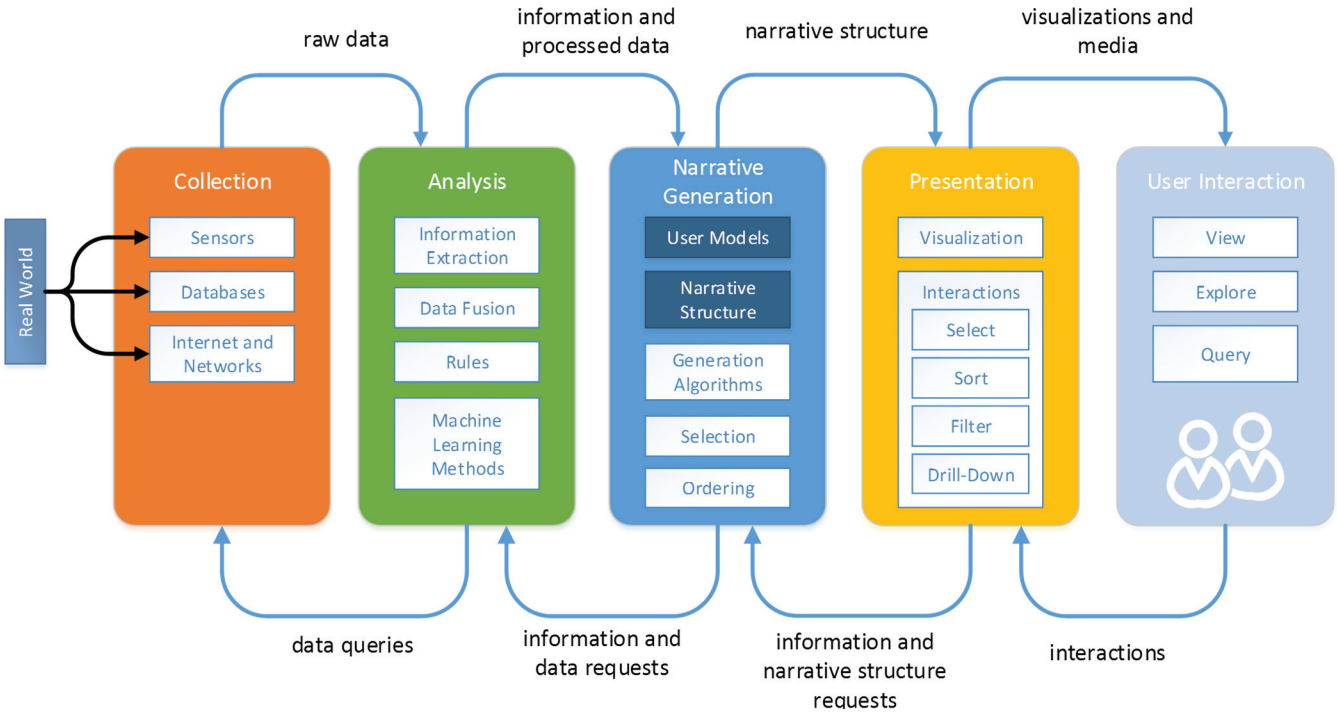

Figure 1: A general architecture of intelligent narrative-based interfaces for sensemaking. Each process step (e.g., Collection) includes examples of the types of subsystems it may include (e.g., Sensors, Databases, and Internet and Networks), but these examples are not meant to be exhaustive. Real world data goes through the processes of collection, analysis, narrative generation, and presentation before reaching the user. Interactions with the user generate requests for narrative structure, information, and data. These elements are computed and returned up the chain of processes in a continuous loop of interaction.

and patient outcomes could provide essential insight into and ease understanding of this complex information.

General Architecture

This subsection defines a general architecture for intelligent narrative-based interfaces. We have developed this architecture through multiple efforts designing and building intelligent narrative-based interfaces. Other researchers may have different concepts or architectures that are better tuned to their research challenges, and our purpose is not to prescribe a solution, but to generate discussion and provide suggestions that we believe to be generally helpful to the community. This architecture provides a basis for generating our research questions in the next-subsection.

Figure 1 proposes a general architecture for intelligent narrative-based interfaces for sensemaking. On the left of the figure, data from the *Real World* goes through the *Collection* process to gather raw data with sensors, into databases, and across the Internet and other networks (e.g., the input of patient data into the EMR system). This raw data is then processed in *Analysis* to create structured data and information using techniques such as information extraction, data fusion, rules, and machine learning (e.g., the computation of patient treatment time statistics correlated with different hospital policies and treatment team makeups). This information and structured data is the plot of the narrative, and it is provided to *Narrative Generation* to select and organize the data according to user models, characters, intentions, and narrative structures, creating the narrative discourse (e.g., the evolution of the hospital as it grows, takes on new responsibilities, and evolves with new technologies and treatments). The user models and narrative structures are highlighted as central components of the intelligent-narrative based interfaces, and they are given special consideration in the following Research Questions subsection. The narrative structure is passed to *Presentation* to create the narrative medium using visualizations and other media (e.g., the organized set of text, labels, and graphs that depict the hospital's history). The presentation places the narrative in an interface and provides affordances for interactions such as enabling the user to select, sort, filter, and drill-down on aspects of the narrative. In *User Interaction*, the user views, explores, and queries the interface to interact with the narrative (e.g., requesting more information about particular time periods or types of staff).

These user actions generate requests that travel down the chain of processes until they can be fulfilled. Interactions may be satisfied by the presentation process, or it may generate a request for additional narrative structure from narrative generation. This additional structure may require more data from the analysis process, which may require more raw data from the collection process, potentially to be gathered from the real world. Once the request can be satisfied the data and processing flows back up the chain as before, presenting an updated narrative-based interface to the user.

Research Questions

This section presents the central research questions we identify from our key concepts and general architecture for intelligent narrative-based interfaces.

Causal and Intentional Analysis. To explain the underlying data in a narrative, the narrative generator must identify and select causal and intentional relationships from the underlying data. These relationships determine a critical part of the ground truth of the narrative. Key questions for this analysis include: How are the events identified? How are the causal and intentional relationships determined, and how are uncertainties in these relationships represented and conveyed to the users?

Narrative Structure. The narrative structure organizes the information and processed data of the analysis into a form that takes advantage of the user's existing understanding of events and information. Key questions for narrative structure include: How are features of the underlying domain, user questions, and context of use best matched with potential narrative structures? What narrative structures are most widely applicable to intelligent narrative-based interfaces, and to what extent can these structures be composed? An evidence-based guide and library of narrative structures for information presentation would be an invaluable resource for future systems.

User Models. User models work with narrative structures to ensure that the narratives are well comprehended by the user. User models include representations of user knowledge, information goals, and narrative processing. Key questions for user models include: What aspects of the user are the most important to track? How do they represent narrative elements such as time, space, causality, objects, characters, and intentions? How can user comprehension or other key aspects of the narrative experience be assessed or estimated? What are guidelines for customization of the narrative to individual users? How can the interface discover and track users' goals? How can the user be assessed during interaction? What are the best uses of cognitive and emotional models for modeling users?

Narrative Visualizations and User Interaction. The visualization of and user interaction with the narrative elements are critical components of the architecture because they are the means by which the user views the narrative and supporting information. Visualizations can be highly domain specific—such as providing maps for geo-political news stories, trends in patient data over time, or box scores for summarizing sporting events—or general as in generic charts and graphs. Key questions for visualizations and user interaction include: How are different dimensions of the nar-

rative best visualized for particular users? Which information is shown initially, and how does the user query the system to find out more? What are the most useful visualization and interaction flows to relate the narrative, and what mediums best support this flow?

Evaluation. The evaluation of intelligent narrative-based interfaces is central to future research and ultimate successes or failures. Key evaluation questions include: What are the most useful metrics and measures for these interfaces, e.g., can usability be separated from overall information communication? How can approaches be compared, especially across domains? How can individual components, such as elements of the narrative generation, be evaluated independently of the whole system? What representative user bases and tasks are needed for informative evaluation?

Conclusions

This paper identifies critical questions and potential research paths in intelligent narrative-based interfaces for sensemaking. We believe this new area of research has significant potential to use narrative technologies to impact a wide range of real world problems.

References

Allen, N., Templon, J., McNally, P., Birnbaum, L., & Hammond, K. 2010. Statsmonkey: A data-driven sports narrative writer. In AAAI Fall Symposium: Computational Models of Narrative.

Antoun, C., Antoun, M., Ryan, J., Samuel, B., Swanson, R., and Walker, M. 2015. Generating natural language retellings from Prom Week play traces. Procedural Content Generation in Games

Barot, C., Cardona-Rivera, R., Eger,M., Glatz, M., Potts, C. M., Robertson, J., Shukunobe, M., Tateosian, L., Thorne, B. R., and Young, R. M., 2017. Bardic: Generating Multimedia Narrative Reports for Game Logs, Technical report, University of Utah. EAE TR 2017-003.

Birnbaum, L., Hammond, K., Allen, N., & Templon, J. 2014 System and method for using data to automatically generate a narrative story. US Patent 8,688,434.

Bouayad-Agha, N., Casamayor, G., & Wanner, L. 2011. Content selection from an ontology-based knowledge base for the generation of football summaries. In Proceedings of the 13th European Workshop on Natural Language Generation, Association for Computational Linguistics, 72-81.

Bouayad-Agha, N., Casamayor, G., Mille, S., and Wanner, L. 2012 Perspective-oriented generation of football match summaries: Old tasks, new challenges. ACM Transactions on Speech and Language Processing (TSLP), 9(2):3.

Bruner, J. (1991). The narrative construction of reality. *Critical inquiry*, *18*(1), 1-21.

Engelmore, R., and Morgan, A. eds. 1986. *Blackboard Systems.* Reading, Mass.: Addison-Wesley.

Gershon, N., & Page, W. (2001). What storytelling can do for information visualization. Communications of the ACM, 44(8), 31-37.

Gervas, P. 2013. Stories from games: Content and focalization selection in narrative composition. In Actas del Primer Simposio Espanol de Entretenimiento Digital, page 25, 2013.

Lareau, F., Dras, M., and Dale, R. 2011 Detecting interesting event sequences for sports reporting. In Proceedings of the 13th European Workshop on Natural Language Generation. Association for Computational Linguistics 200-205.

Miranda, B., S. C., Garbay, C., and Portet, F. 2015. "Generating Récit from Sensor Data: Evaluation of a Task Model for Story Planning and Preliminary Experiments with GPS Data." In *ENLG*, pp. 86-89.

Molina, M., and Flores, V. 2012. Generating multimedia presentations that summarize the behavior of dynamic systems using a model-based approach. *Expert Systems with Applications*, 39(3):2759–2770.

Riedl, M., and Young, R. M. 2010. Narrative planning: balancing plot and character, *Journal of AI Research,* 39(1):217–268.

Ryan, J., Mateas, M., and Wardrip-Fruin N. 2015. Open Design Challenges for Interactive Emergent Narrative. In Proc. International Conference on Interactive Digital Storytelling.

Segel, E., & Heer, J. (2010). Narrative visualization: Telling stories with data. *IEEE transactions on visualization and computer graphics*, *16*(6), 1139-1148.

Singer, M., Graesser, A. C., & Trabasso, T. (1994). Minimal or global inference during reading. *Journal of Memory and Language*, *33*(4), 421-441.

Swanson, R., and Jhala, A., 2012. “Rich computational model of conflict for virtual characters”, in *Intelligent Virtual Agents*, Springer Berlin Heidelberg, pp. 502–504.

Teutenberg, J., and Porteous, J. 2013. Efficient intent-based narrative generation using multiple planning agents.In *Proceedings of the 2013 International Conference on Autonomous Agents and Multi-Agent Systems*, 603–610. Inter-national Foundation for Autonomous Agents and Multiagent Systems

Teutenberg, J., and Porteous, J. 2015. Incorporating global and local knowledge in intentional narrative planning. In *Proceedings of the Conference on Autonomous Agents and Multi-Agent Systems*, 1539–1546.

Thorne, B., and Young, R. M. 2017. Generating stories that include failed actions by modeling false character beliefs. Technical report, University of Utah. EAE TR 2017-001.

Valve Corporation. 2013. DoTA 2. (Video game, PC).

Vaudry, P-L, and Lapalme, G. 2016. "Assembling Narratives with Associative Threads." In *The INLG 2016 Workshop on Computational Creativity in Natural Language Generation*, p. 1.

Ware, S. G., and Young, R. M. 2014. Glaive: a state-spacenarrative planner supporting intentionality and conflict. In *Tenth Artificial Intelligence and Interactive Digital Entertainment Conference*

Young, R. M. 2017. Sketching a generative model of intention management for characters in stories: Adding intention management to a belief-driven story planning algorithm. Technical report, University of Utah. EAE TR 2017-002.

Young, R. M., 2010. Story and Discourse: A Bipartite Model of Narrative Generation in Virtual Worlds, in *Interaction Studies: Social Behaviour and Communication in Biological and Artificial Systems*, no. 2, vol 18, pages 177-208.