Semantic Storytelling: Towards Identifying Storylines in Large Amounts of Text Content

Georg Rehm Karolina Zaczynska Julián Moreno-Schneider

Speech and Language Technology Lab, DFKI GmbH
Alt-Moabit 91c, 10559 Berlin, Germany
Corresponding author: georg.rehm@dfki.de

Abstract

In this position paper we present an approach and vision we call Semantic Storytelling. The idea is to develop a system that, given an incoming document collection, is able to (semi-)automatically extract or generate different story paths or plot lines towards the goal of supporting knowledge workers (journalists, authors, scholars, politicians, business analysts etc.) in their daily work of processing huge amounts of incoming content. We outline the different components needed, which can be summarised as preprocessing, semantic analysis and content enrichment, as well as generating storylines. Our idea is to take into account the specificities of different text genres, which, we believe, will help us to generate better results according to the needs and characteristics of the respective text genre. We give a brief example where Semantic Storytelling can be applied and try to pinpoint the main conceptual, scientific and technical gaps that still need to be addressed fully to realise our vision of a Semantic Storytelling system.

1 Introduction

The ever increasing amount of information available online is posing an enormous challenge for information and content curation professionals whose professional job profile includes analysing or making use of the information including understanding the, in a wider sense, storyline of an event or series of events. The problem of identifying or generating narrative structures in a robust way is yet to be solved. There is a big need for tools that support users to automatically identify, interpret and relate the elements of a (fictional or real) narrative.

In that sense, many professionals have to cope with huge amounts of incoming information and content that need to be processed (scanned, skimmed, contextualised, evaluated and, eventually, further processed) in a short amount of time in order to produce a new piece of information, for example a news article, a social media post, a longread or a press statement – let us call this group "knowledge workers" or "content curators". Generally, they still need to do a great amount of work intellectually. If automatic tools are available, these are often restricted to specific tasks, for example, keyword-based alerting or named entity tagging.

This position paper describes the conceptual design and partial implementation of a Semantic Storytelling prototype [RMSB+18]. Our aim is that the system will be able to process an incoming set of textual data

In: A. Jorge, R. Campos, A. Jatowt, S. Bhatia (eds.): Proceedings of the Text2StoryIR '19 Workshop, Cologne, Germany, 14-April-2019, published at http://ceur-ws.org

(e. g., a document collection) using several semantic analysis technologies in order to generate a large variety of semantic annotations that can be exploited for the purpose of semantic storytelling. This involves supporting users through various interactive and dynamic data and content exploration methods that rely on abstract story knowledge.

This article is structured as follows. Section 2 covers related work. In Section 3 we describe our vision regarding Semantic Storytelling and Section 4 concludes the article and discusses the most relevant conceptual, scientific and technical gaps.

2 Background and Related Work

Several approaches are closely related to our Semantic Storytelling concept and vision, all of them concentrating on their own specific objectives and providing solutions for their respective challenges. Ours is to enable, ideally with limited or no human intervention at all, the identification of plots or storylines based on text collections for which we generate rich and deep semantic annotations.

Some approaches focus on extracting information using NLP techniques in order to use it, in later stages, for generating stories. An unsupervised approach for clustering news articles based on identified event instances is presented by Ribeiro et al. [RFT17], while Li et al. [LZY17] present a supervised prediction model to analyse different strength levels of claims in science news as a fact-checker. The NewSum Toolkit [VM15] is a combination of NLP and Machine Learning technologies supporting a number of steps for news article writing like gathering data, automatic classification and summarisation of large amounts of incoming articles. More complex approaches are used by Yarlott et al. [YCGF18] based on the hierarchical theory of discourse by van Dijk [vD88], or by Dai et al. [DTH18], where a content representation structure of the documents is used to build a first predictive model using these indicative structures as features. News recommendation has gained attention in works such as Cucchiarelli et al. [CMSV18], where journalists get recommendations by taking an event and checking if they got a greater echo in Twitter or Wikipedia postings, or in Bois et al. [BGJ+17], where newspaper articles are recommended based on lexical similarity, linked through a graph representation of relations.

A different class of systems is mainly oriented on providing content or applications for entertainment purposes. For example, Wood [Woo08] uses a collection of pictures and other media to generate albums, Gervás [Ger13] focuses on gaming. Other groups use story structuring methods as part of therapy programs [KBE14], while other approaches focus on "storytelling" or, rather, text generation, in a particular domain, typically recipes [CLNU13, Dal89] or weather reports [Bel08, RSHD05, TSRD06], requiring knowledge about characters, actions, locations, events, or objects [GDAPH05, RY10, Tur14]. A notable exception to this approach, where domain knowledge is a prerequisite, is [LLUJR13], who attempt to construct plot graphs from a set of stories annotated using crowd sourcing. Some authors include the order of events [Cha11].

For the demanding question of how to generate a story grammar which orders events detected in a previous step into storylines, many approaches are based on theories taken from literature studies, more precisely, narratology. For example, Caselli and Vossen [CV17] use the plot structure as described by Bal [Bal97] for a chronological and logical ordering of events, Yarlott et al. [YF16] and others used Propps morphology of Russian hero tales [Pro68] as theoretical background for story detection and generation systems. We plan to experiment with the concept of text genres, specifically text-structural conventions, to get a better understanding of structure in texts and to better extract the main events inside these structures that often exhibit specific communicative functions. One approach describing text genres according to their communicative functions can be found in [Sha18].

Another important component of our Semantic Storytelling vision is the graphical user interface, which will enable users to interact not only with the information that has been analysed but also with the generated storylines. Examples of final story visualisations are Ma et al. [MLF+12] or Segel and Heer [SH10], who present several options. Novel visual interactive methods are presented in Kybartas and Bidarra [KB15]. In contrast, an approach focused on the content management is presented by Mulholland et al. [MWC12]. Storyteller visualises complex relations between events found in newspaper articles [vMVvdZ+17], while user interactions are limited to filtering the data set.

3 Towards Semantic Storytelling

This section presents our Semantic Storytelling vision including the concept, an indicative use case and the different components needed for the development of a complete system.

3.1 Semantic Storytelling – Brief Overview

Storytelling is a human technique to order a series of events in the world and find meaningful patterns in it [Bru91]. By telling a story we relate events into a schematic structure, for example, in terms of topic, locality or causal relationships, and construct explanatory models of the world and events. Semantic Storytelling can be seen as the attempt to translate the theories of storytelling into a formal, and machine-processable scheme.

Storytellers dynamically adjust their narratives and tell their stories differently depending on who the listener is [RLEW13]. The most simplistic goal of storytelling is the automatic (or semi-automatic) generation of stories, where a story is considered a natural language text containing a complete, correct and unambiguous story. The definition by Rishes et al. [RLEW13] splits storytelling into semantic content generation and natural language generation.

We rather see a storyline as a set of building blocks, which depending on their combination (temporal, geographical, semantic, causal) form a story, which allows us to provide a wider range and more flexibility for suggesting storylines. Our goal is to develop a suggestion- or recommender system that allows the, ideally, automatic arrangement of (named) entities, i. e., conceptual instances, and events, within a storyline, where users benefit from a recommender system and a controlled context and navigation tool.

3.2 Indicative Use Case Illustration

The following example is meant to illustrate the functionality of an "ideal implementation" of the Semantic Storytelling system we have in mind. Tens of thousands of books on the Second World War provide detailed information on events that involve different persons, places, alliances etc. A historian, journalist or author working on the topic needs to be able to order and arrange this vast amount of content in an intelligent way to create new content. An ideal system can support the understanding of historical interactions and relationships. The goal is to identify all persons, places and events, to position events on a timeline, also to identify the causal, temporal etc. relationship between different events. While Natural Language Processing is not yet able to perform these tasks without any errors, we firmly believe that the application of state of the art methods can provide a benefit to the user, for example, by following the storylines of individual persons, exploring their relationships with others, focusing upon specific events, scrolling backward or forward in time.

3.3 Architecture and Components

The abstract architecture of our system is composed of three main building blocks: Semantic Analysis, Text Genre-specific Story Knowledge and Semantic Generation (Figure 1). In the following, we briefly describe the three sets of components, especially concentrating on the conceptual and technological gaps.

3.3.1 Semantic Analysis

This building block involves various processing steps that relate to the annotation, extraction and classification of certain parts of the incoming content in order to enrich the documents, for example, by adding semantics and information taken from external sources. Named Entity Recognition, Named Entity Linking and Time Expression Analysis are needed to identify named entities of various types and classes (Persons, Locations, Organization, Others) and to anchor the content to a timeline. Extracted entities, topics etc. will be linked to external knowledge graphs (e. g., DBPedia,[1] Wikidata,[2] Geonames,[3]). A robust approach at Topic Detection is needed in order to assign abstract topics to, say, individual sentences, paragraphs, chapters and documents. Annotated topics will enable yet another different layer of accessing and recombining the processed content. Managing the linguistic annotations in a Linked Data format (we use NIF in our current prototype [HLAB13]) allows the exploitation of Linked Open Data for storyline generation. While robust Event Detection with a high coverage, carried out at the same level of semantic abstraction, is still beyond the state of the art of Natural Language Processing, such a module is crucial to enable the re-composition of storylines out of a large and heterogeneous set of identified events. Automatically anchoring events to a timeline is also beyond what is possible right now fully automatically. To analyse a wide variety of incoming documents, we need to be able to process different classes or genres of documents, we need to identify and work with Discourse Structure, we need to identify the genre or type of a document, we need to be able to distinguish fact from fiction. While

[1] `https://dbpedia.org`
[2] `https://www.wikidata.org`
[3] `https://www.geonames.org`

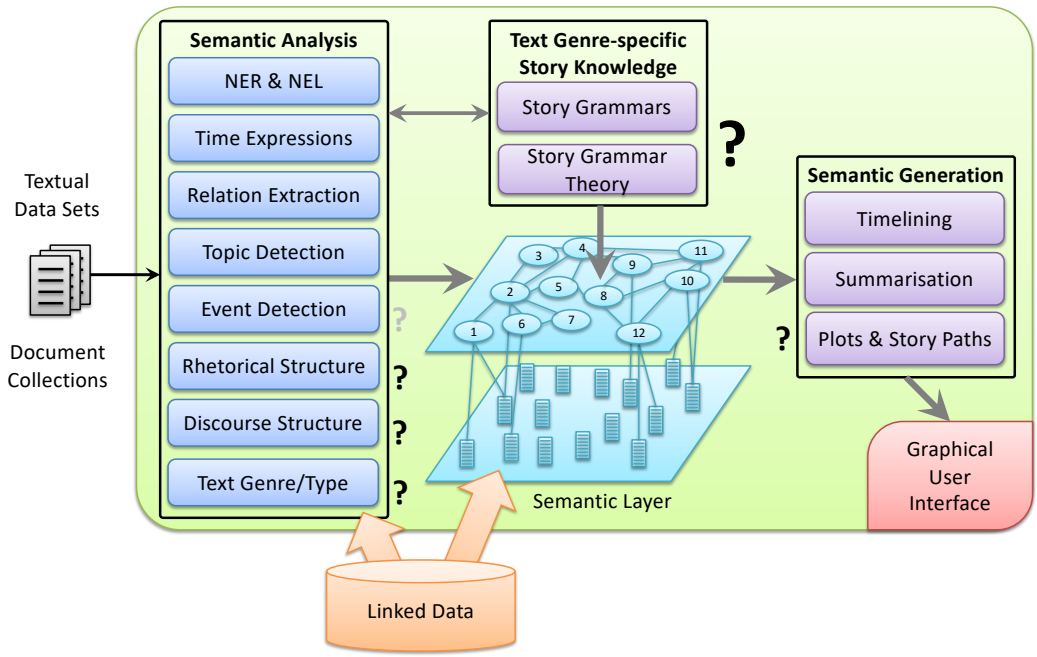

Figure 1: The abstract architecture of our Semantic Storytelling approach

components such as these are beyond what is technically feasible or possible currently, we believe that discourse- and genre-informed processing is a crucial component of Semantic Storytelling [Reh07].

3.3.2 Semantic Generation

As previously mentioned, Semantic Generation involves the dynamic and interactive recomposition and visualisation of extracted information based on the information extracted from the Semantic Analysis step. This especially involves arranging content elements (documents, paragraphs, sentences, claims or events) on a dynamic timeline. Summarisation techniques can be used to compress larger pieces of content into bites that can be easily digested, moved around on the screen and maybe expanded back into longer or their original versions. The principles by which the actual construction of storylines based on the recomposition of previously extracted information will be performed, is still an open question. In contrast to template-filling approaches [MSBR17], we will focus on approaches that are based on computational narratology to generate narrative structures of story lines, while using automatically extracted information and external knowledge provided as Linked Open Data [RMSB+18].

3.3.3 Story Knowledge

The most crucial missing conceptual piece of our Semantic Storytelling vision is, critically, what we call Text Genre-specific Story Knowledge for advanced text and discourse-informed document processing. This includes technologies and approaches for representing and identifying the structure, patterns, sequences and abstract entities of different types of stories and to make this explicit knowledge available to corresponding analysis and, later, generation components. Therefore, we will build generic processing pipelines for different text types, which will allow us to handle typical features found within these. The idea is to apply a generic processing pipeline optimised by the characteristics of the respective text genre. While for news, for example, a timeline-based event ordering would be useful, for journal articles discourse structure aspects like discourse parsing and claim extraction would be the key aspects a knowledge worker would be interested in.

Text genre-informed information retrieval will make it possible to process each text according to its typical text genre-specific structure and communicative function. Using this knowledge we can more precisely extract the most important lexicogrammatical realisations which express these main commmunicative aims, for example to inform about an event, to discuss a scientific claim or to present a story etc. (cf. Sharoff's classification of web text genres [Reh02], where text genres are defined by generalised communicative aims and not predetermined

by lexicogrammatical realisations [Sha18]). This approach will allow us to extract events and order them in a flexible way, addressing the needs of the respective use case and document collection.

4 Conclusions

Semantic Storytelling can be conceptualised as the automatic (or semi-automatic) generation of different storylines based on information extracted, classified and annotated within extensive textual data sets or document collections [BMSN+16]. We have developed a number of initial prototypes that demonstrate part of the functionality needed [RHS+17, MSBR17, SBN+16, RMSB+18].

Our goal is the development of a prototype platform that will support knowledge workers in the complex and time-consuming task of handling, evaluating, processing, sorting and processing document collections in order to generate new pieces of content. One goal of the platform is to enable users to identify interesting stories as efficiently as possible based on the (extracted) information available.

We are trying to pinpoint the key open questions in order to suggest a roadmap for Semantic Storytelling for the next years. While technologies such as Named Entity Recognition and Linking, Time Expression Analysis, Topic Detection and Text Classification have been in production use in many different applications for years, important components such as Event Detection but especially more advanced discourse analysis tools including Rhetorical Structure Analysis and Text Genre detection must still be considered avantgarde and not ready for production use yet. While research on such technologies is making progress, the wider field of Natural Language Understanding and Language Technology still needs to fully discover and embrace the relevance and importance of what we call Text Genre-specific Story Knowledge for truly advanced text and discourse-informed document processing. In our future work we will concentrate on the development of corresponding technologies and knowledge representation approaches.

Acknowledgements

The project QURATOR is supported by the German Federal Ministry of Education and Research (BMBF), "Unternehmen Region", instrument "Wachstumskern" (grant no. 03WKDA1A).

References

[Bal97] Mieke Bal. *Narratology: Introduction to the Theory of Narrative*. University of Toronto Press, 1997.

[Bel08] Anja Belz. Automatic Generation of Weather Forecast Texts Using Comprehensive Probabilistic Generation-space Models. *Nat. Lang. Eng.*, 14(4):431–455, October 2008.

[BGJ+17] Rémi Bois, Guillaume Gravier, Eric Jamet, Emmanuel Morin, Pascale Sébillot, and Maxime Robert. Language-based construction of explorable news graphs for journalists. In *NLPmJ@EMNLP*, pages 31–36. Association for Computational Linguistics, 2017.

[BMSN+16] Peter Bourgonje, Julian Moreno-Schneider, Jan Nehring, Georg Rehm, Felix Sasaki, and Ankit Srivastava. Towards a Platform for Curation Technologies: Enriching Text Collections with a Semantic-Web Layer. In H. Sack, G. Rizzo, N. Steinmetz, D. Mladenic, S. Auer, and C. Lange, editors, *The Semantic Web: ESWC 2016 Satellite Events, Heraklion, Crete, Greece, May 29 – June 2, 2016, Revised Selected Papers*, pages 65–68. Springer International Publishing, June 2016.

[Bru91] Jerome Bruner. The narrative construction of reality. *Critical Inquiry*, 18(1):1–21, 1991.

[Cha11] Nathanael William Chambers. *Inducing Event Schemas and Their Participants from Unlabeled Text.* dissertation, Stanford University, 2011.

[CLNU13] Philipp Cimiano, Janna Lüker, David Nagel, and Christina Unger. Exploiting Ontology Lexica for Generating Natural Language Texts from RDF Data. In *Proceedings of the 14th European Workshop on Natural Language Generation*, pages 10–19, Sofia, Bulgaria, August 2013. Association for Computational Linguistics.

[CMSV18] Alessandro Cucchiarelli, Christian Morbidoni, Giovanni Stilo, and Paola Velardi. What to write and why: a recommender for news media. In *SAC*, pages 1321–1330. ACM, 2018.

[CV17] Tommaso Caselli and Piek Vossen. The event storyline corpus: A new benchmark for causal and temporal relation extraction. In *Proceedings of the Events and Stories in the News Workshop*, pages 77–86. Association for Computational Linguistics, 2017.

[Dal89] Robert Dale. Cooking Up Referring Expressions. In *Proceedings of the 27th Annual Meeting on Association for Computational Linguistics*, ACL '89, pages 68–75, Stroudsburg, PA, USA, 1989. Association for Computational Linguistics.

[DTH18] Zeyu Dai, Himanshu Taneja, and Ruihong Huang. Fine-grained structure-based news genre categorization. In *Proceedings of the Workshop Events and Stories in the News 2018*, pages 61–67, Santa Fe, New Mexico, U.S.A, August 2018. Association for Computational Linguistics.

[GDAPH05] Pablo Gervás, Belén Díaz-Agudo, Federico Peinado, and Raquel Hervás. Story Plot Generation based on CBR. In Ann Macintosh, Richard Ellis, and Tony Allen, editors, *Applications and Innovations in Intelligent Systems XII: Proceedings of AI-2004, the Twenty-fourth SGAI International Conference on Innovative Techniques and Applications of Artificial Intelligence*, pages 33–46, London, 12/2004 2005. Springer London.

[Ger13] Pablo Gervás. Stories from Games: Content and Focalization Selection in Narrative Composition. In *I Spanish Symposium on Entertainment Computing*, Universidad Complutense de Madrid, Madrid, Spain, 09/2013 2013.

[HLAB13] Sebastian Hellmann, Jens Lehmann, Sören Auer, and Martin Brümmer. Integrating NLP using Linked Data. In *Proceedings of the 12th International Semantic Web Conference*, 2013. 21-25 October 2013.

[KB15] Ben Kybartas and Rafael Bidarra. A Semantic Foundation for Mixed-Initiative Computational Storytelling. In Henrik Schoenau-Fog, Luis Emilio Bruni, Sandy Louchart, and Sarune Baceviciute, editors, *Interactive Storytelling - 8th International Conference on Interactive Digital Storytelling, ICIDS 2015, Copenhagen, Denmark, November 30 - December 4, 2015, Proceedings*, volume 9445 of *Lecture Notes in Computer Science*, pages 162–169. Springer, 2015.

[KBE14] Ben Kybartas, Rafael Bidarra, and Elmar Eisemann. Integrating semantics and narrative world generation. In *Proceedings of FDG 2014 - Ninth International Conference on the Foundations of Digital Games*, apr 2014. Fort Lauderdale, FL.

[LLUJR13] Boyang Li, Stephen Lee-Urban, George Johnston, and Mark O. Riedl. Story Generation with Crowdsourced Plot Graphs. In *Proceedings of the Twenty-Seventh AAAI Conference on Artificial Intelligence*, AAAI'13, pages 598–604. AAAI Press, 2013.

[LZY17] Yingya Li, Jieke Zhang, and Bei Yu. An nlp analysis of exaggerated claims in science news. In *Proceedings of the 2017 EMNLP Workshop: Natural Language Processing meets Journalism*, pages 106–111. Association for Computational Linguistics, 2017.

[MLF+12] Kwan-Liu Ma, Isaac Liao, Jennifer Frazier, Helwig Hauser, and Helen-Nicole Kostis. Scientific Storytelling using Visualization. *IEEE Computer Graphics and Applications*, 32(1):12–19, 2012.

[MSBR17] Julian Moreno-Schneider, Peter Bourgonje, and Georg Rehm. Towards User Interfaces for Semantic Storytelling. In Sakae Yamamoto, editor, *Human Interface and the Management of Information: Information, Knowledge and Interaction Design, 19th International Conference, HCI International 2017 (Vancouver, Canada)*, number 10274 in Lecture Notes in Computer Science (LNCS), pages 403–421, Cham, Switzerland, July 2017. Springer. Part II.

[MWC12] Paul Mulholland, Annika Wolff, and Trevor Collins. Curate and Storyspace: An Ontology and Web-Based Environment for Describing Curatorial Narratives. In Elena Simperl, Philipp Cimiano, Axel Polleres, Oscar Corcho, and Valentina Presutti, editors, *The Semantic Web: Research and Applications: 9th Extended Semantic Web Conference, ESWC 2012, Heraklion, Crete, Greece, May 27-31, 2012. Proceedings*, pages 748–762, Berlin, Heidelberg, 2012. Springer Berlin Heidelberg.

[Pro68] Vladimir Y. Propp. *Morphology of the folktale.* Publication ... of the Indiana University Research Center in Anthropology, Folklore, and Linguistics. University of Texas Press, 1968.

[Reh02] G. Rehm. Towards Automatic Web Genre Identification – A Corpus-Based Approach in the Domain of Academia by Example of the Academic's Personal Homepage. In *Proceedings of the 35th Hawaii International Conference on System Sciences (HICSS-35)*, Big Island, Hawaii, January 2002.

[Reh07] Georg Rehm. *Hypertextsorten: Definition – Struktur – Klassifikation.* Books on Demand, Norderstedt, 2007. PhD thesis in Applied and Computational Linguistics, Justus-Liebig-Universität Giessen, 2005.

[RFT17] Swen Ribeiro, Olivier Ferret, and Xavier Tannier. Unsupervised event clustering and aggregation from newswire and web articles. In *Proceedings of the 2017 EMNLP Workshop: Natural Language Processing meets Journalism*, pages 62–67. Association for Computational Linguistics, 2017.

[RHS+17] Georg Rehm, Jing He, Julian Moreno Schneider, Jan Nehring, and Joachim Quantz. Designing User Interfaces for Curation Technologies. In *19th International Conference on Human-Computer Interaction – HCI International 2017*, Vancouver, Canada, July 2017.

[RLEW13] Elena Rishes, Stephanie M. Lukin, David K. Elson, and Marilyn A. Walker. *Generating Different Story Tellings from Semantic Representations of Narrative*, pages 192–204. Springer International Publishing, Cham, 2013.

[RMSB+18] Georg Rehm, Julian Moreno-Schneider, Peter Bourgonje, Ankit Srivastava, Rolf Fricke, Jan Thomsen, Jing He, Joachim Quantz, Armin Berger, Luca König, Sören Räuchle, Jens Gerth, and David Wabnitz. Different Types of Automated and Semi-Automated Semantic Storytelling: Curation Technologies for Different Sectors. In Georg Rehm and Thierry Declerck, editors, *Language Technologies for the Challenges of the Digital Age: 27th International Conference, GSCL 2017, Berlin, Germany, September 13-14, 2017, Proceedings*, number 10713 in Lecture Notes in Artificial Intelligence (LNAI), pages 232–247, Cham, Switzerland, January 2018. Gesellschaft für Sprachtechnologie und Computerlinguistik e.V., Springer. 13/14 September 2017.

[RSHD05] Ehud Reiter, Somayajulu Sripada, Jim Hunter, and Ian Davy. Choosing words in computer-generated weather forecasts. *Artificial Intelligence*, 167:137–169, 2005.

[RY10] Mark Owen Riedl and Robert Michael Young. Narrative Planning: Balancing Plot and Character. *J. Artif. Int. Res.*, 39(1):217–268, September 2010.

[SBN+16] Julin Moreno Schneider, Peter Bourgonje, Jan Nehring, Georg Rehm, Felix Sasaki, and Ankit Srivastava. Towards Semantic Story Telling with Digital Curation Technologies. In Larry Birnbaum, Octavian Popescuk, and Carlo Strapparava, editors, *Proceedings of Natural Language Processing meets Journalism – IJCAI-16 Workshop (NLPMJ 2016)*, New York, July 2016.

[SH10] Edward Segel and Jeffrey Heer. Narrative Visualization: Telling Stories with Data. volume 16, pages 1139–1148, Piscataway, NJ, USA, November 2010. IEEE Educational Activities Department.

[Sha18] Serge Sharoff. Functional text dimensions for the annotation of web corpora. *Corpora*, 13(1):65–95, 2018.

[TSRD06] Ross Turner, Somayajulu Sripada, Ehud Reiter, and Ian P. Davy. Generating Spatio-temporal Descriptions in Pollen Forecasts. In *Proceedings of the Eleventh Conference of the European Chapter of the Association for Computational Linguistics: Posters & Demonstrations*, EACL '06, pages 163–166, Stroudsburg, PA, USA, 2006. Association for Computational Linguistics.

[Tur14] Scott R. Turner. *The Creative Process: A Computer Model of Storytelling and Creativity.* Taylor & Francis, 2014.

[vD88] Teun van van Dijk. *News as discourse.* Communication Series. L. Erlbaum Associates, 1988.

[VM15] Ivan Vulić and Marie-Francine Moens. Monolingual and cross-lingual information retrieval models based on (bilingual) word embeddings. In *Proceedings of the 38th International ACM SIGIR Conference on Research and Development in Information Retrieval*, SIGIR '15, pages 363–372, New York, NY, USA, 2015. ACM.

[vMVvdZ+17] Maarten van Meersbergen, Piek Vossen, Janneke van der Zwaan, Antske Fokkens, Willem van Hage, Inger Leemans, and Isa Maks. Storyteller: Visual analytics of perspectives on rich text interpretations. In *Proceedings of the 2017 EMNLP Workshop: Natural Language Processing meets Journalism*, pages 37–45. Association for Computational Linguistics, 2017.

[Woo08] Mark D. Wood. Exploiting Semantics for Personalized Story Creation. In *Proceedings of the 2008 IEEE International Conference on Semantic Computing*, ICSC '08, pages 402–409, Washington, DC, USA, 2008. IEEE Computer Society.

[YCGF18] W. Victor H. Yarlott, Cristina Cornelio, Tian Gao, and Mark Finlayson. Identifying the discourse function of news article paragraphs. In *Proceedings of the Workshop Events and Stories in the News 2018*, pages 25–33. Association for Computational Linguistics, 2018.

[YF16] W. Victor H. Yarlott and Mark A. Finlayson. Proppml: A complete annotation scheme for proppian morphologies. In *CMN*, volume 53 of *OASICS*, pages 8:1–8:19. Schloss Dagstuhl - Leibniz-Zentrum fuer Informatik, 2016.