Transient Narrative Networks and Information Landscapes for Enhancing Human Understanding

Remi van Trijp¹

¹SONY Computer Science Laboratories, 6, rue Amyot, 75005 Paris, France

Abstract

This demonstration paper presents work-in-progress on a system that explores the use of interactive narratives and information landscapes for helping people to make sense of complex topics, especially those who struggle with data overload, information anxiety and other illnesses of the Digital Age. The current prototype allows people to explore articles from the English Wikipedia in three ways: a news-like web interface for careful reading, a 'transient narrative network' for coherence-building which unfolds based on user interaction, and an information landscape for perceiving which information exists outside of the user's current purview.

1. Introduction

The explosion of data in the Digital Age has so far mainly been treated as a technological problem that can be solved through more sophisticated data management techniques. However, an equally important challenge is how to manage the impact of data explosion on people's well-being. More specifically, problems such as data overload [1] or information anxiety [2] make it increasingly difficult for people to make sense of the events that they experience or perceive on a daily basis. Moreover, data and communication overload also have disastrous effects on interpersonal understanding, leading to more hostile interactions especially on social media [3].

One of the most promising avenues for solving these problems is to use *human-centric* artificial intelligence that can help to enhance the human understanding process. This demonstration paper therefore presents work-in-progress on a system that explores how interactive narratives and information landscapes could be used for this purpose. The core of the prototype is implemented using the Babel Toolkit [4], an open-source software platform that includes an interactive web interface and cognitive language technologies such as Fluid Construction Grammar [5], which will be used in future versions of the prototype for deep linguistic analysis. The core system is designed to be highly modular and extensible through a RESTful architecture that allows the system to request and integrate information from external components.

2. Three Ways of Exploring Information

The first prototype of the system allows users to explore information from wikipedia, which has been selected as the prototype's knowledge base because of its scale, free access, linguistic diversity, and evolving nature. The prototype currently uses the English wikipedia but localization efforts are made to extend the interface to access articles from other languages as well.

The following three subsections explain the system's three ways of exploring information from the perspective of a fictitious person called Emily, a young British woman who is very concerned about the war in Ukraine and who wants to learn more about the major entities that are involved.

2.1. Careful Reading

It is March 24, 2022. Emily hears on the news that British Prime Minister Boris Johnson is speaking at a press conference at the NATO headquarters in Brussels, where he accuses Russia of committing war crimes. Emily doesn't remember well what NATO is, or why the Prime Minister would go there, so she decides to look for more information.

Upon starting the system, Emily is greeted by a traditional-looking, news-like webpage as illustrated in Figure 1. She then types "NATO" in the search bar in the upper right corner, and a wikipedia summary about NATO now appears under the header "focus" in the top left, with an option to read the full article. Below the summary she finds links to related pages, while on the right she sees three more prominent suggestions under the header "related entities."

As their names suggest, both the "related pages"

IJCAI 2022: Workshop on semantic techniques for narrative-based understanding, July 24, 2022, Vienna, Austria

Cremi.vantrijp@sony.com (R. v. Trijp) 2022 Copyright for this paper by its authors. Use permitted under Creative Commons License Attribution 4.0 International (CC BY

CEUR Workshop Proceedings (CEUR-WS.org)



Figure 1: The system's web interface starts with a news-like homepage that invites careful reading, with a search bar in the top right corner that allows the user to query wikipedia. On the left, a summary of the corresponding hit is shown, with the option to read the full article. On the right, the three most similar wikipedia entities are shown based on entity embeddings. Below, related pages are shown that are obtained through the Wikimedia REST API.

and "related entities" section aim to offer more pieces of information that may be relevant to Emily's quest for knowledge. They only differ in where they are derived from: the related pages are obtained directly from Wikipedia's linked data through the Wikimedia REST API, while the related entities were suggested by a model of embeddings of entities (vector representations) trained with the Wikipedia2Vec tool [6], which not only learns from wikipedia's link graph model but which also uses a word-based skip-gram model and an anchor context model (using the neighbouring words of a hyperlink that points to an entity as context).

2.2. Transient Narrative Network

Emily clicks on a number of related pages, as she usually does when browsing the web, but she soon loses track of which topics she has already covered or how they relate to each other. She therefore switches to the *narrative network* view using the navigation bar on top.

Narratives are widely accepted to be a key element of the human sense-making process [7, 8], particularly as a vehicle for creating coherence from otherwise fragmented, disparate and noisy information sources. AI researchers have therefore become increasingly interested in incorporating the concept of a narrative into the design of human-centric AI systems [9, 10], especially now that information is more scattered than ever across billions of connected devices.

Following studies in narratology [11], a narrative is assumed to have a three-layered structure:

- The *fabula* (or *story*) is the set of facts and events; similar to the concept of ground truth. In the current prototype, the set of wikipedia articles is assumed to represent the fabula.
- The *plot* (or *syuzhet*) is a structure that arranges the relevant facts and events into a causal chain or causal network.
- The *narration* (or *narrative presentation*) concerns *how* the narrative is presented.

The prototype has been observing Emily's interactions with the system, and keeps track of which pages she visited (her personal *plot*) and which related pages and entities would make possible excursions or new pathways for exploration (parts of the fabula that may become relevant).¹

¹All the while respecting her privacy: all of her information is stored locally on her own computer.



Figure 2: The system builds transient narrative networks based on how the user interacts with the system. The user can traverse and manipulate the network, and switch back to careful reading mode to learn more about the currently selected node. Here, the larger nodes and the red directed path illustrates the path followed by the user, while smaller nodes and grey edges provide potential pathways.

The resulting narrative, shown in Figure 2, is operationalized as a transient narrative network [12] using vis.js (a dynamic browser-based visualization library). A transient narrative network is a graph that dynamically changes as new nodes and edges are added when there is additional input or when more information is requested from other knowledge sources. Here, the graph shows which nodes represent pages that Emily visited with directed links indicating the order in which she traversed the information space. Additional nodes and edges are automatically added from wikipedia's linked graph model and shown as possible paths for exploration, with more or less prominence depending on where Emily is currently situated in the network (i.e. which node is selected). Emily can now choose to continue exploring the network and expanding it by selecting nodes of interest, or she can return to the careful reading view to learn more details about a topic. She can also manually add, remove and modify her own nodes and edges in other to further personalize or complete the graph.

2.3. Information Landscapes

While Emily now has a more coherent answer to her initial question, she is still somewhat distrustful about the system. How reliable are the system's suggestions? And what information is out there? The narrative network is anchored to her own perspective so she can only see the so-called "adjacentpossible" nodes [13, 14]. What if she is missing out on something?



Figure 3: Information landscapes allow users to get a glimpse of the available information outside of their purview. Here, the 100 most similar entities to "NATO" are shown.

She therefore switches to a third page, which offers a data visualization of all Wikipedia entities, where she can see at which location she currently finds herself. She also gets a glimpse about which information is out there, both close by and remotely located. In the current prototype, the only visualization that is available is a visualization of the embeddings of wikipedia entities by the Wikipedia2Vec tool [6], illustrated in Figure 3, which is not yet connected to the rest of the system.

3. Conclusion and Future Work

This demonstration paper introduced the first prototype of a system under development that aims to enhance human understanding through humancentric artificial intelligence. More specifically, it combines careful reading with narrative networks (for coherence building) and information landscapes (for purposefully navigating the information space) through a web-based interface.

Future work is planned in all three areas. For the careful reading page, additional components need to be integrated that support different reading and comprehension monitoring strategies [15]. These may include a combination of neurostatistical models for keyword extraction (for quickly scanning texts) or entity recognition and linking (for improving the recommendations); and symbolic models for extracting more detailed semantic frames [16, 17].

To enhance the narrative networks, additional knowledge sources need to be integrated such as knowledge graphs (e.g. WikiData). Such knowledge sources can also be employed for offering different narrations of the same plot. Promising work already exists on how knowledge graphs can be exploited for presenting events on a chronological timeline [10], which may further help users to build a coherent picture about topics that they might otherwise explore in a more random fashion. Narrative networks also need to be stored in a Personal Dynamic Network [9].

Finally, more interactive visualization methods are needed that provide users with intuitive ways to understand the information space. This atlas of semantic maps then needs to be connected to the core system so that the system can propose more interesting routes for navigating information.

4. Ethical Statement

There are no ethical issues.

Acknowledgments

The work reported in this paper was carried out in the context of MUHAI (Meaning and Understanding in Human-Centric AI – https://muhai.org/), a project funded by the European Union's Horizon 2020 research and innovation programme under grant No 951846. I wish to thank the workshop organizers Lise Stork, Katrien Beuls, and Luc Steels; the anonymous reviewers; and the MUHAI partners for their invaluable feedback. Many thanks also go to my colleagues Inès Blin and Martina Galletti. Finally, I wish to thank Vittorio Loreto and Hiroaki Kitano for creating such a superb working environment.

References

- A. Toffler, Future Shock, Random House, New York, 1970. Co-authored by Heidi Toffler (uncredited).
- [2] R. S. Wurman, Information Anxiety: Towards Understanding, Doubleday, New York, 1989.
- [3] M. Fan, Y. Huang, S. A. Qalati, S. M. M. Shah, D. Ostic, Z. Pu, Effects of information overload, communication overload, and inequality on digital distrust: A cyber-violence behavior mechanism, Frontiers in Psychology 12 (2021). URL: https://www.frontiersin. org/article/10.3389/fpsyg.2021.643981. doi:10. 3389/fpsyg.2021.643981.
- [4] M. Loetzsch, P. Wellens, J. De Beule, J. Bleys, R. van Trijp, The Babel2 Manual, Technical Report AI-Memo 01-08, AI-Lab VUB, Brussels, 2008.
- [5] L. Steels (Ed.), Design Patterns in Fluid Construction Grammar, volume 11 of Constructional Approaches to Language, John Benjamins, Amsterdam, 2011. URL: https://doi. org/10.1075/cal.11.
- [6] I. Yamada, A. Asai, J. Sakuma, H. Shindo, H. Takeda, Y. Takefuji, Y. Matsumoto, Wikipedia2Vec: An efficient toolkit for learning and visualizing the embeddings of words and entities from Wikipedia, in: Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing: System Demonstrations, Association for Computational Linguistics, 2020, pp. 23–30.
- J. Bruner, The Narrative Construction of Reality, Critical Inquiry 18 (1991) 1–21.
 URL: http://www.jstor.org/stable/1343711, publisher: The University of Chicago Press.
- [8] B. Boyd, On the Origin of Stories: Evolution, Cognition, and Fiction, Harvard University Press, Boston MA, 2009. URL: http://www.jstor.org/stable/j.ctvjf9xvk. doi:10.2307/j.ctvjf9xvk.
- [9] L. Steels, Personal dynamic memories are necessary to deal with meaning and understanding in human-centric AI, in: A. Saffiotti, L. Serafini, P. Lukowicz (Eds.), Proceedings of the First International Workshop on New Foundations for Human-Centered AI (NeHuAI) co-located with 24th European Con-

ference on Artificial Intelligence (ECAI 2020), CEUR Workshop Proceedings, Santiago de Compostela, 2020. URL: http://ceur-ws.org/ Vol-2659/steels.pdf.

- [10] I. Blin, Building narrative structures from knowledge graphs, in: P. Groth, M.-E. Vidal, F. Suchanek, P. Szekley, P. Kapanipathi, C. Pesquita, H. Skaf-Molli, M. Tamper (Eds.), Extended Semantic Web Conference (ESWC) 2022, Springer Nature, Cham, 2022.
- [11] M. Bal, Narratology: Introduction to the Theory of Narrative, University of Toronto Press, Toronto, 2017. Fourth edition, first ed. 1985.
- [12] L. Steels, Narrative art interpretation, in: L. Steels (Ed.), MUHAI Deliverable D1.1 White Paper. Foundations for Incorporating Meaning and Understanding in Human-Centric AI, The MUHAI Consortium, Bremen, 2022.
- [13] S. Kauffman, At Home in the Universe: The Search for Laws of Self-Organization and Complexity, Oxford University Press, Oxford, 1995.
- B. Monechi, A. Ruiz-Serrano, F. Tria,
 V. Loreto, Waves of novelties in the expansion into the adjacent possible, PloS one 12 (2017) e0179303-e0179303. URL: https://pubmed.ncbi.nlm.nih.gov/28594909. doi:10.1371/journal.pone.0179303, publisher: Public Library of Science.
- Y.-F. Yang, Reading Strategies or Comprehension Monitoring Strategies?, Reading Psychology 27 (2006) 313–343. URL: https://doi.org/10.1080/02702710600846852. doi:10.1080/02702710600846852, publisher: Routledge.
- [16] V. Micelli, R. van Trijp, J. De Beule, Framing Fluid Construction Grammar, in: N. Taatgen, H. van Rijn (Eds.), Proceedings of the 31th Annual Conference of the Cognitive Science Society, Cognitive Science Society, 2009, pp. 3023–3027.
- [17] K. Beuls, P. Van Eecke, V. S. Cangalovic, A computational construction grammar approach to semantic frame extraction, Linguistics Vanguard 7 (2021).