

Causal networks as the backbone for temporal data-to-text

Pierre-Luc Vaudry and Guy Lapalme

RALI-DIRO – Université de Montréal

C.P. 6128, Succ. Centre-Ville

Montréal, Québec, Canada, H3C 3J8

{vaudrypl, lapalme}@iro.umontreal.ca

1 Introduction

Causal networks have been successfully used to explain the process of narrative comprehension in humans (Trabasso et al., 1989). This motivated their use in the automatic creation of fairy tales (Swartjes and Theune, 2006; Theune et al., 2007). Some have suggested that causal relations also play an important role in improving narrative generation from real-life temporal data (Hunter et al., 2012; Gervás, 2014). Several narrative data-to-text systems already identify and make use of some causal relations (Hallett, 2008; Hunter et al., 2012; Bouayad-Agha et al., 2012; Wanner et al., 2010). Is it possible to go one step further in the identification of causal relations and aim at extracting a causal network that could be used to improve the coherence of generated texts? Given that real-life temporal data could be analysed and interpreted automatically so as to produce an appropriate causal network, how could this information be best used to organize a coherent narrative adapted to the communicative needs of a wide range of data-to-text applications?

We will here address the second question by proposing a bottom-up document planning method for building the rhetorical structure of coherent narratives from a causal network. The causal network represents the plot of the story to be told by the narrative text. The rhetorical structure represents one way, corresponding to a particular viewpoint, of organizing the text to tell that story. The goal is therefore to translate the semantic information contained in a causal network into a rhetorical structure using parameters that characterize this perspective. The proposed approach uses hierarchical clustering and parameterization by adjacency and ordering preferences. It is robust to partial causal networks and does not re-

quire the generated rhetorical structure to be a tree.

The document planning stage would be preceded by signal analysis and data interpretation and followed by microplanning and surface realisation in a complete data-to-text pipeline (Reiter, 2007).

2 Proposed approach

For structuring narrative text content, we propose to use a bottom-up approach. Bottom-up approaches, contrarily to top-down ones, guarantee that all chosen content will be included in the rhetorical structure. This can avoid continuity problems that are due to missing events in the generated text (Portet et al., 2009, pp. 807–808). We retain from the bottom-up approach of Marcu (1997) the use of adjacency and ordering preferences and taking into account all rhetorical relations, not only those appearing in the tree.

Our algorithm takes a causal network similar to the one used by Trabasso et al. (1989) as input. It includes actions, goals and mental reactions, which are attributed to an agent, as well as physical events and states. If distinguishing the mental processes of multiple agents is important, other kinds of events or states must also be included, such as perceptions (Swartjes and Theune, 2006). Moreover, in that case a distinction should be made between facts considered as objective and those existing only in the mind of an agent.

First some additional logico-semantic relations (Bouayad-Agha et al., 2012) are inferred, such as volitional causation, contrast and conjunction. This is done using rules. When a physical event causes a goal and that goal causes an action, we can derive a volitional causation relation between the event and the action. Several types of contrast relations can be inferred from differing mental reactions related to the same goal and

from events identified as unexpected results of other events.

Then, bottom-up hierarchical clustering is used to build a tree structure based on adjacency preferences. Adjacency preferences must be specified for causal relations and other logico-semantic relations inferred from the causal network. For this purpose, causal relations are divided into subtypes depending on the nature of their arguments (event causes goal, goal causes action, etc.). In addition, the temporal adjacency preference specifies how textual distance should relate to relative temporal distance. Adjacency preferences are expressed in terms of how much a given relation prefers to have its arguments appear in the same sentence, in the same paragraph or in another paragraph.

The distance between two events or states is computed from the average of the adjacency preferences of the logico-semantic relations holding between them. The temporal distance relative to the total duration of the period to be narrated is also taken into account, although with a low weight.

The distance at which two clusters are merged, because it reflects the adjacency preferences of their members, can later be used to determine whether they should appear in the same sentence or paragraph.

Next logico-semantic relations are mapped to rhetorical relations in the sense of Rhetorical Structure Theory (RST) (Mann and Thompson, 1987) with respect to communicative constraints. As observed by Bouayad-Agha et al. (2012, p. 9), a causal relation can be mapped to a temporal circumstance rhetorical relation instead of a rhetorical cause, if a neutral perspective is needed. More mappings are possible depending on the types of cause and effect and on which argument is selected as satellite and which as nucleus of the rhetorical relation.

Finally, pairs of sibling clusters are ordered to produce an ordered tree based on ordering preferences. Ordering preferences are specified for each type of rhetorical relation in terms of which of the satellite or the nucleus tends to come first and how strong this tendency is. The temporal ordering preference specifies to what extent chronological or reverse chronological order should be followed. During ordering, sibling clusters in the tree are ordered by averaging the ordering preferences of all the relations holding between them.

Combined with adjacency and ordering preferences, different logico-semantic to rhetorical

relation mappings can be used to generate texts that emphasize different aspects of the relations between the same set of events and states. This could be used to account for differences between literary genres and styles.

To verify this, the method was implemented and a preliminary experiment was conducted. Texts of three different genres (news story, fairy tale and factual report) were successfully generated from a single input.

3 Future work

Future work will consist of further researching the mapping between logico-semantic relations and rhetorical relations under pragmatic constraints, determining whether this mapping should be performed after or before clustering. Additionally, means of acquiring the parameters of this mapping and the adjacency and ordering preferences from a corpus will be investigated. Inferring other logico-semantic relations from the causal network will also be considered.

In addition, setting and final situation of the narrative will be taken into account in document planning. The method will also be tested on longer and more diversified inputs and types of text.

Finally, how causal networks can be best extracted from real-life data with minimal use of handwritten rules will be further researched. In particular, multiple time series models, such as vector autoregressive processes and dynamic Bayesian networks, could possibly be used to discover causal relations types. They could be trained on available time series corpora and/or based on domain knowledge. Low probability events according to the model could be assumed to have external causes and could mark the start of a causal chain. In that case they could be termed trigger events. Incorporating to generation the confidence level with which instances of causal relations were identified in the data is also appealing. It could for example influence the logico-semantic to rhetorical relation mapping.

References

- Nadjet Bouayad-Agha, Gerard Casamayor, Simon Mille, and Leo Wanner. 2012. Perspective-oriented Generation of Football Match Summaries: Old Tasks, New Challenges. *ACM Trans. Speech Lang. Process.*, 9(2):3:1–3:31, August.
- Pablo Gervás. 2014. Composing narrative discourse for stories of many characters: A case study over a

- chess game. *Literary and Linguistic Computing*, August.
- Catalina Hallett. 2008. Multi-modal presentation of medical histories. In *Proceedings of the 13th international conference on Intelligent user interfaces*, pages 80–89.
- James Hunter, Yvonne Freer, Albert Gatt, Ehud Reiter, Somayajulu Sripada, and Cindy Sykes. 2012. Automatic generation of natural language nursing shift summaries in neonatal intensive care: BT-Nurse. *Artificial intelligence in medicine*.
- William C. Mann and Sandra A. Thompson. 1987. *Rhetorical structure theory: A theory of text organization*. University of Southern California, Information Sciences Institute.
- Daniel Marcu. 1997. From Local to Global Coherence: A Bottom-up Approach to Text Planning. In *Proceedings of the Fourteenth National Conference on Artificial Intelligence and Ninth Conference on Innovative Applications of Artificial Intelligence*, pages 629–635, Providence, Rhode Island. AAAI Press.
- François Portet, Ehud Reiter, Albert Gatt, Jim Hunter, Somayajulu Sripada, Yvonne Freer, and Cindy Sykes. 2009. Automatic generation of textual summaries from neonatal intensive care data. *Artificial Intelligence*, 173(7):789–816.
- Ehud Reiter. 2007. An Architecture for Data-to-text Systems. In *Proceedings of the Eleventh European Workshop on Natural Language Generation*, pages 97–104, Stroudsburg, PA, USA. Association for Computational Linguistics.
- Ivo Swartjes and Mariët Theune. 2006. A fabula model for emergent narrative. In *Technologies for Interactive Digital Storytelling and Entertainment*, pages 49–60. Springer.
- Mariët Theune, Nanda Slabbers, and Feikje Hielkema. 2007. The Narrator: NLG for digital storytelling. In *Proceedings of the Eleventh European Workshop on Natural Language Generation*, pages 109–112. Association for Computational Linguistics.
- Tom Trabasso, Paul Van den Broek, and So Young Suh. 1989. Logical necessity and transitivity of causal relations in stories. *Discourse Processes*, 12(1):1–25.
- Leo Wanner, Bernd Bohnet, Nadjat Bouayad-Agha, François Lareau, and Daniel Nicklaß. 2010. Marquis: Generation of User-Tailored Multilingual Air Quality Bulletins. *Applied Artificial Intelligence*, 24(10):914–952.