

Narrative Discourse as a Target Format for Data to Text Generation

Pablo Gervás

Instituto de Tecnología del Conocimiento
Universidad Complutense de Madrid
28040 Madrid, Spain
pgervas@ucm.es

Abstract

This abstract outlines some issues to be taken into account in data to text generation, regarding the nature of data that might be considered and the type of text that has more chances to be accepted as a successful vehicle for communication. Narrative discourse is proposed as a promising candidate and some ongoing efforts to study this type of discourse are listed.

1 Pitfalls and Opportunities for Data to Text Generation

The concept of data to text generation has been around for some time as a possible application of natural language generation techniques with a very significant potential for impact in terms of applicability to real life problems. In recent times there has clearly been an explosion of data that become available thanks to improvements in computing power, in storage capacity, or in volume of transactions due to the way in which processes of human interaction have progressively become digital. For a discipline like natural language generation (NLG) that has for many years striven to convert into legible texts all sorts of data formats, this wealth of digital data is a treasure trove. An unequalled opportunity arises for NLG: the task of providing text paraphrases of this massive amount of data is beyond the capabilities of human transcribers, so automated solutions for data to text conversions have an opening for proving their worth.

However, there are some issues to take into consideration before assuming that NLG should take on the responsibility of transcribing to text every single item of this wealth of digital information: on the nature of the data that might be so transcribed, and on the kind of text that constitute an

optimal format for humans to understand the nature of the data or to digest its significance to meet their own needs.

First, the large accumulation of data includes data of many kinds: large sets of numerical figures, images, video, audio, temporal series, statistics, operation logs... The radically different nature of these data calls for more than one presentation format. In the past a number of presentation formats – such as graphs, maps, tables, animations, diagrams – have been preferred to text. This is in spite of the fact that the humans that produce them had unquestioned ability to produce text to describe the data sources involved. It is clear that in human experience text is not always the preferred format for presenting data. This should definitely be taken into account when considering where NLG might best be applicable for data to text generation. Regardless of the fact that text may now be produced automatically for a variety of input data, any attempt to compete with the current preferred presentation formats where they have clearly been preferred over text in the past is unlikely to meet with practical success that might give rise to massive uptake.

Second, the fact that humans have for centuries communicated by means of text does not necessarily mean that text in itself is best suited for conveying many of the meanings that it is being used for. It is indeed a universally understood format that can be adapted to convey almost any meaning. Yet in some cases the resulting texts make for difficult reading that only experts in the particular topics can process with ease. This is the case for specialized texts, for instance in legal or scientific fields. Such texts are perfectly understandable to experts, yet remain obscure for non-specialists. If an effort is to be made to convert data into text format, this should be aimed at a type of text that has a high probability of being understandable to the man in the street. In this paper we want to put for-

ward a proposal that narrative might be one such type.

In our everyday lives, a high percentage of the linguistic messages that we exchange with those in our immediate environment occur as part of stories. When we describe something that happened at work, or what we saw on the way there, when we remember how we met someone, when we tell a joke, when we give an example... we are very frequently telling stories. A large number of valuable knowledge management protocols rely on stories as a preferred format: best practices, use cases, examples for medical diagnosis, persuasive argumentation, court defence cases... Yet narrative text has for a long time been side stepped by the natural language generation community in favour of less flexible formats such as weather reports, instructions manuals, or technical documents in general. This is understandable in view of the difficulties and challenges involved in narrative. However, it is also clear that narrative has a high potential for being a preferred textual vehicle for much of the data that surrounds us. Any kind of data that involves evolution over time is likely to benefit from narrative treatment. Data concerning any kind of human dimension, from emotion to motivation, will best find expression in narrative form.

The need for NLG to devote more attention to narrative has already been defended elsewhere (Gervás, 2013a), and some of the existing efforts along those lines were also described. In the present abstract we will focus on summarising our own ongoing work over the past few years to this end.

2 Ongoing Efforts on Narrative Composition

The NIL research group (<http://nil.fdi.ucm.es>) has over the last 15 years worked on natural language generation, with particular emphasis in the generation of narrative and poetry, under the influence of the emerging field of computational creativity. The efforts of our group to develop storytelling technology over the last 10 years have lead to the development of a number of working story telling systems. Our initial attempts focused on achieving full story telling functionality, all the way to a natural language text that told the story. These produced stories that were somewhat stilted. Since then we have tried to partition the storytelling problem into smaller subproblems

(Gervás and León, 2014), with the hope that improving each one of them will lead to a better story teller once the individual solutions are integrated. The problems that we are working on now correspond to the following four stages:

- **INVENTION:** Coming up with an interesting set of events with high potential for being told as a story. These set of events are not necessarily linear in nature and may be more exhaustive than the usual story, including material required to understand why things happen but which may be omitted in the telling (Gervás, 2013b; Gervás, 2014c; León and Gervás, 2014; Laclaustra et al., 2014; Méndez et al., 2014).
- **COMPOSITION:** Coming up with an appropriate linear discourse - described conceptually - for conveying a set of events. This involves filtering out any material that the reader/listener may infer from the rest, but also linearising the material to follow first one character and then another when several relevant events are happening simultaneously at different points in space (Gervás, 2014a; Gervás, 2012; Gervás, 2013c).
- **VALIDATION:** Being able to tell, for a given narrative discourse, whether it is likely to work well as a communication tool (Gervás, 2014b; Gervás, 2013c).
- **TELLING:** Coming up with solutions for converting conceptual linear discourses into text (namely classic tasks of natural language generation).

These modules are being designed to operate together as a pipeline, aiming to produce a text for the story at the end. All the stages described above are ongoing lines of research. Some are more advanced than others but they are all operational to some degree, though we are still working on improving their quality. They are also all currently active within our funded projects, so they are likely to make significant progress in the near future. We believe that the resulting technology has a very strong potential to play a key role in data to text generation, not just as a discourse planning tool for data that is already narrative in its original form, but also to come up with narrative skeletons as may be needed to structure non-narrative data into a story-like presentation format.

Acknowledgments

The projects ConCreTe, WHIM and PROSECCO acknowledges the financial support of the Future and Emerging Technologies (FET) programme within the Seventh Framework Programme for Research of the European Commission, under FET grant numbers 611733, 611560 and 600653.

References

- Pablo Gervás and Carlos León. 2014. Reading and writing as a creative cycle: The need for a computational model. In *5th International Conference on Computational Creativity, ICCV 2014*, Ljubljana, Slovenia, 06/2014.
- Pablo Gervás. 2012. From the fleece of fact to narrative yarns: a computational model of composition. In *Workshop on Computational Models of Narrative, 2012 Language Resources and Evaluation Conference (LREC'2012)*, Istanbul, Turkey, 05/2012.
- Pablo Gervás. 2013a. Narrative composition: Achieving the perceived linearity of narrative. In *Proceedings of the 14th European Workshop on Natural Language Generation*, pages 103–104, Sofia, Bulgaria, August. Association for Computational Linguistics.
- Pablo Gervás. 2013b. Propp's morphology of the folk tale as a grammar for generation. In *Workshop on Computational Models of Narrative, a satellite workshop of CogSci 2013: The 35th meeting of the Cognitive Science Society*, Universität Hamburg Hamburg, Germany, 08/2013. Schloss Dagstuhl – Leibniz-Zentrum für Informatik GmbH, Dagstuhl Publishing, Saarbrücken/Wadern, Germany, Schloss Dagstuhl – Leibniz-Zentrum für Informatik GmbH, Dagstuhl Publishing, Saarbrücken/Wadern, Germany.
- Pablo Gervás. 2013c. Stories from games: Content and focalization selection in narrative composition. In *1 Spanish Symposium on Entertainment Computing*, Universidad Complutense de Madrid, Madrid, Spain, 09/2013.
- Pablo Gervás. 2014a. Composing narrative discourse for stories of many characters: a case study over a chess game. *Literary and Linguistic Computing*, 08/14.
- Pablo Gervás. 2014b. Metrics for desired structural features for narrative renderings of game logs. *Journal of Entertainment Computing*, 08/14.
- Pablo Gervás. 2014c. Reviewing Propp's story generation procedure in the light of computational creativity. In *AISB Symposium on Computational Creativity, AISB-2014, April 1-4 2014*, Goldsmiths, London, UK, 04/2014.
- Iván M. Laclaustra, José L. Ledesma, Gonzalo Méndez, and Pablo Gervás. 2014. Kill the dragon and rescue the princess: Designing a plan-based multi-agent story generator. In *5th International Conference on Computational Creativity, ICCV 2014 (Late Breaking paper)*, Ljubljana, Slovenia, 06/2014.
- Carlos León and Pablo Gervás. 2014. Creativity in story generation from the ground up: Non-deterministic simulation driven by narrative. In *5th International Conference on Computational Creativity, ICCV 2014*, Ljubljana, Slovenia, 06/2014.
- Gonzalo Méndez, Pablo Gervás, and Carlos León. 2014. A model of character affinity for agent-based story generation. In *9th International Conference on Knowledge, Information and Creativity Support Systems*, Limassol, Cyprus, 11/2014. Springer-Verlag, Springer-Verlag.