A Challenge Proposal for Narrative Generation Using CNLs

Eugenio Concepción and **Gonzalo Méndez** and **Pablo Gervás** and **Carlos León**
Facultad de Informática
Instituto de Tecnología del Conocimiento
Universidad Complutense de Madrid
{econcepc,gmendez,pgervas,cleon}@ucm.es

Abstract

We propose a competitive shared evaluation task for Narrative Generation. It would involve the generation of new stories for a given domain from common ground knowledge shared by all systems. A set of source materials will be provided for development, represented in Controlled Natural Language (CNL), which should also be used to phrase the text outputs of participating systems. By having all participating systems operate from the same sources for knowledge and generate in a compatible output format, comparability of the results will be enhanced. Submitted results will be subject to both automatic and human evaluation.

1 Introduction

A story generator algorithm (SGA) refers to a computational procedure resulting in an artefact that can be considered a story (Gervás, 2012). The term `story generation system` can be considered as a system that applies a SGA to construct stories. There is a growing population of such story generation systems that share two significant characteristics: one, they operate from a set of knowledge resources that act as input to the story generation process; two, they rely on elementary text building solutions – usually based on template filling – for producing human-readable versions of their outputs. Comparative evaluations of any kind between these story generation systems are very difficult because: different systems start from different (unrelated) knowledge resources, and text outputs of the different systems are heavily influenced by the (different) sets of templates employed to render them. A common approach to acquiring knowledge resources is to mine a set of reference stories, to obtain from them the required knowledge. These resources usually make explicit two types of information that is implicit in the stories: relation between events in the story and latent variables relevant to it – such as causality, emotion, affinities between characters, narratological concepts... –, and/or information about typical/acceptable sequencing between events – depending on the degree of refinement of the system, sometimes based on the latent variables.

The present proposal revolves around the idea of developing a Controlled Natural Language (CNL) that can be used to specify the *source material* for a story generation task. A CNL is an engineered subset of natural languages whose grammar and vocabulary have been restricted in order to reduce both ambiguity and complexity of full natural languages (Schwitter, 2010). If such a CNL could be used to represent a set of reference stories, while ensuring that any latent variables are made explicit in the representation, it should be possible to automatically extract the relevant knowledge resources from such source material. To make this possible, the type of source material required should include a set of example stories either enriched with explicit mentions of latent variables or accompanied by explicit declaration of the relation between elements in the stories and the latent variables. If textual outputs of story generation systems could be phrased in such a CNL, it should be feasible to compare outputs of different systems on a shared common footing.

171

2 Conceptual Basis

In (Gervás and León, 2014) the authors provided a list of the most relevant classifications of the story generation systems according to the type of knowledge resources that they rely on, and the way these knowledge resources are implemented as specific data structures. That paper proposed a list of aspects of a narrative relevant to story telling systems in this sense: including the discourse produced for the story, the representation of the activity of agents – in terms of actions, interactions, mental states, and movement between locations –, the causal relations between elements in the story, the motivations of agents, the theme of parts of the story, the emotions involved in or produced by the story, the intentions of the author, and the narratological concepts involved in the story structure. These various aspects constitute sources of candidate features for the role of latent features relevant for story telling.

CNLs can be considered as a tradeoff between the expressivity of natural languages and the need for a formal representation that can be handled by computers. The requirements for the definition of a CNL grammar (Kuhn, 2010) relevant for the present purpose are: that it should be fully formalized and interpretable by computers, it should not depend on a concrete algorithm or implementation, it should be easy to implement in different programming languages, and it should be sufficiently expressive (for the task at hand).

3 A Proposal for a Story Generation Shared Task

The feasibility of the shared task relies on the development of two basic resources: a grammar for a CNL capable of representing the various aspects relevant to story telling and the resources required by story generation systems, and a set of source materials that encode the necessary knowledge for generating stories in a specific domain covering a selected set of the relevant aspects.

The challenge as proposed is addressed to existing story generation systems.[1] The task would involve: extracting task-specific instances of the knowledge

---

[1]It may be undertaken by researchers willing to develop a system from scratch if they consider it feasible, but the effort involved would be much higher.

resources required for the candidate system from the source materials provided, adapting the text rendering modules of the story generation system to generate stories as close as possible to the the CNL developed for the task, and submitting the resulting stories for evaluation.

3.1 Development of Resources

The proponents of the challenge intend to enlist the collaboration of authors of existing story generation systems with a two-fold purpose: to ensure that the developed resources provide coverage of as many aspects of narrative deemed relevant from a computational perspective, and to raise interest in the challenge and build a community of candidate participants. The collaboration envisaged would take the form of providing sample instances of the knowledge resources employed by their system for generation in a domain of their choice.

3.1.1 The CNL

Such resources will be used to inform the iterative development of the grammar for the CNL. An initial grammar will be built covering aspects common to all systems and all resources. This grammar will be progressively enriched with any additional aspects covered by some systems and not by others, until all selected aspects are covered. Depending on what aspects are covered by the compiled resources and how easy they are to embed into a story, decisions will need to be made on how to represent the relevant latent variables, either as explicit enrichment of stories or as separate declaration of their relation to story elements. Some progress has already been made along these lines (Concepción et al., 2016).

A parser will be developed for the CNL, capable of building actual data structures for the various aspects represented. Both the grammar for the CNL and the code for the parser will be made available to participants. The parser will be designed so that it has a specific module for saving the data structures to disk. Such module may be reinstantiated by participants to select which part of the knowledge in the data structure is saved onto what particular representation format for a particular system.

3.1.2 The Source Materials

The CNL will be used to develop source materials for a particular domain chosen as focus for the challenge. Source materials may consist of a set of enriched stories and/or a set of definitions of relations between story elements and latent variables. Additional knowledge relevant to the domain may also need to be encoded – using the CNL – in the source materials. The basic scope and structure of such additional material will be based on the concept of a *story bible* or *show bible* as considered by screenwriters for information on a television series' characters, settings, and other elements.

3.1.3 Evaluation Procedures

Textual outputs produced by participating systems will be processed using the parser described in 3.1.1. Outputs will be rated automatically on the following parameters: grammaticality – based on conformance to the grammar –, novelty with respect to reference stories in the source materials – data structures built by the parser from the outputs will be compared with those arising from the reference stories according to existing metrics for narrative similarity (Peinado et al., 2010; Hervás et al., 2015) – , and additional rating schemes developed for any relevant features – as the data structures generated by the parser will include explicit representation of these aspects, development for specif metrics is possible for features like degree of causal connectivity, rise and fall of emotion or affinity between characters over a story, or any others explicitly represented.

For the parameters chosen, judgements from human evaluators will also be compiled.

3.1.4 Expected Timeline

A tentative timeline is proposed which would involve: requesting contributions – as samples of knowledge resources – from interested researchers by the end of September 2016, publish source materials in March 2017, outputs to be submitted by participants by July 2017, final results presented at INLG 2017. However, in view of the various uncertainties existing in the proposal, it may be necessary to comtemplate the need to postpone the submission deadline to 2018, in which case the tentative timeline may be re-distributed accordingly over the intervening period.

4 Expected Benefits

The development of agreed versions of source materials from which story generation resources can be extracted, a grammar for outputs of story systems, and procedures for quantitative measurement of relevant features would constitute significant benefits.

Acknowledgments

References

E. Concepción, P. Gervás, G. Méndez, and C. León. 2016. Using CNL for knowledge elicitation and exchange across story generation systems. In *5th Workshop on Controlled Natural Language (CNL 2016)*, Aberdeen, Scotland, 07/2016. Springer, Springer.

P. Gervás and C. León. 2014. The need for multi-aspectual representation of narratives in modelling their creative process. In *5th Workshop on Computational Models of Narrative*, OASIcs-OpenAccess Series in Informatics.

P. Gervás. 2012. Story generator algorithms. In *The Living Handbook of Narratology*. Hamburg University Press.

R. Hervás, A. Sánchez-Ruiz, P. Gervás, and C. León. 2015. Calibrating a metric for similarity of stories against human judgment. In *Creativity and Experience Workshop, International Conference on Case-Based Reasoning*, Bad Homburg, Frankfurt, Germany, 09/2015.

T. Kuhn. 2010. Codeco: A practical notation for controlled english grammars in predictive editors. In *Controlled Natural Language*, pages 95–114. Springer.

F. Peinado, V. Francisco, R. Hervás, and P. Gervás. 2010. Assessing the novelty of computer-generated narratives using empirical metrics. *MINDS AND MACHINES*, 20(4):588, 10/2010.

R. Schwitter. 2010. Controlled natural languages for knowledge representation. In *Proc. of the 23rd International Conference on Computational Linguistics: Posters*, COLING '10, pages 1113–1121, Stroudsburg, PA, USA. Association for Computational Linguistics.