Stella - A Story Generation System for Generic Scenarios

Carlos León

Dpto. de Ingeniería de Software e Inteligencia Artificial
Universidad Complutense de Madrid
Madrid, Spain
cleon@fdi.ucm.es

Computational Narrative is a very challenging subfield of Artificial Intelligence. Several story generation systems have been created so far, and each one of them has proposed new approaches and advances in the area. While the state of the art in story generation has clearly evolved, a robust, useful system that can be used in different scenarios as a shared tool has not been created.

The development of STELLA (STORY TELLING ALGORITHM) tries to fill this gap by proposing a large scale generation system in which the domain definition and the main engine are clearly isolated and only linked by a well-defined interface. In this way, different domain-specific knowledge sets can replace or be appended to the one that has been used in the first prototype.

The algorithm that is proposed in STELLA traverses a generative search space in which new stories are generated by the iterative application of domain rules on a initial story, therefore including new parts of the story on each step. Only addition rules are permitted, which eases the rule definition process and allows a better optimization of the generation engine.

This makes STELLA a knowledge intensive system. Domain rules must be included for the system to be able to perform the generation. The software has been designed by following several knowledge engineering principles and techniques. In this way, creating generative rules is easier.

STELLA's rules rely on a knowledge base that describes the domain. Since STELLA is meant to be a generic tool, the architecture of this knowledge base is also modular and extensible.

Since the main intention in the development of STELLA is to create a useful system, some effort has been put on offering several interfaces to the generation algorithm:

- A programmatic access to the set of constraints of the story, which allows full control over the generation. The core of STELLA is itself a library that can be compiled and used in other projects.

- A command line interface that lets the user set the high

level details of the generation. A story can be directly generated by running the main console-based front-end.

- A simple GUI built as a layer on top of the command line version (Figure 1). This version allows the creation of *emotion curves*, which let the user establish constraints on some variables (in Figure 1, the progression of love and danger in the story has been constrained).
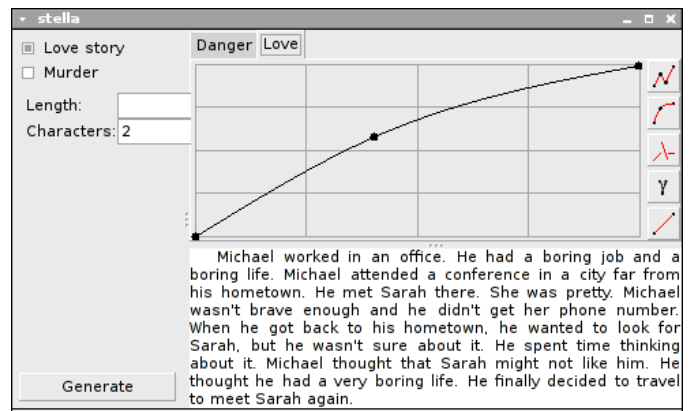

Figure 1: Screenshot of STELLA

While Natural Language Generation is outside the scope of this research, basic template-based generation has been carried out in order to show the output stories. The story shown in Figure 1 has been processed with these templates and slightly corrected by hand.

So far STELLA is able to generate thousands of stories. While most of them probably lack high quality, the knowledge intensive approach ensures acceptability (i.e. all stories can be probably accepted as such by most people). Moreover, the carefully designed software and knowledge engineering process make it possible to improve the system with an affordable effort.

The proposed system puts many recent advances on Computational Narrative together, like character modelling, emotions, story structure and knowledge intensive techniques. The main relative merit of this approach is the development of a tool to be used, with reasonable limits, both in research and in production.