Prototyping the Use of Plot Curves to Guide Story Generation

Carlos León, Pablo Gervás

Escuela Politcnica Superior, Instituto de Tecnologa del Conocimiento
Universidad Alfonso X el Sabio, Universidad Complutense de Madrid
cleonazn@uax.es, pgervas@sip.ucm.es

Abstract

Setting objectives for automatic story generation is needed for a story generation system to produce content. Among the potentially useful methods, curves defining the evolution of specific features of a narrative that evolve along time are particularly appropriate because they focus on the evolution of those features and are easy to create, modify and understand by human users. In this paper we propose a theoretical definition of curve-based story generation, its relation to existing story generation algorithms and how this theory can be applied to new systems.

1. Introduction

Existing computational story generation systems are not autonomous enough to set their own generation objectives. These objectives must be therefore given by the programmer or by the user for the generation system to be able to produce satisfying content. In general, existing systems are able to generate stories in some particular domain which is usually defined in terms of rules or previous stories. Since the concrete definition restricts what can be generated, it can be considered to be a general way of constraining what the automatic story generation system can produce. However, this is far from general and a more specific objective-setting mechanism is more useful.

Among current objective-setting techniques, narrative curve specification is one of the most story-specific for generation. Narrative curves can be defined in general as constraints for some variables or aspects of a story that evolve along the story timeline. The classic *tension arc* is a good example of this definition of a narrative curve (Aristotle, 1974; Tobias, 2012; Zagalo et al., 2004). It very well depicts how the reader's perception of the evolution of the story evolves over time in terms of its conflicts and resolutions.

In this sense, curves offer a powerful and intuitive way of representing how a story should progress. From a human point of view, curves are a very intuitive representation. They are easy to understand, create and adjust to the user's objectives. Additionally, from the perspective of computation many mathematical techniques are available to handle them, which makes it easy to create an implementation for it.

This paper proposes a general definition of narrative curves it terms of their generic aspects and their relation with the computational properties that are interesting for curve-based narrative generation. In particular the application of the proposed model to existing automatic story generation systems has been studied in order to give an example of how this theoretical approach can serve to make a taxonomy of story generation systems that use curves.

1.1. Setting Objectives in Existing Story Generation Systems

Most existing story generation systems are able to set the generation objectives in several ways. Story generation systems based on planning usually define the generation objectives through a set of logic constraints that the planner must use as objective (Lebowitz, 1985; Riedl and Young, 2006). Some systems also take into account the characters' goals. For instance, Meehan's Tale-Spin (Meehan, 1976) performs storytelling by letting virtual characters satisfy their goals. Some systems even go beyond that and consider the duality of both kind of objectives (author's and character's), like MINSTREL (Turner, 1992).

Story grammars have also been used for story generation (Rumelhart, 1975; Lang, 1999). Story grammars do not have an explicit way of setting objectives, however the overall generation objectives can be considered to be hard-coded in the grammar itself. These systems can of course filter out those generated stories that are not wanted by the user (which is a way of establishing objectives) but these do not drive the generation process itself.

Several existing story generation systems also use plot curves to drive the generation process. MEXICA (Pérez y Pérez, 1999) uses the so called *tension* to represent *love*, *emotion* and *danger*. These three features are finally used as a single value (the *tension* itself). This tension is basically a list of discrete values that define a set of restrictions that the *emotional links* between the characters must fulfill.

Barros and Musse address tension arcs in Interactive Fiction (Barros and Musse, 2008). This system uses non-decreasing, discrete curves to represent dramatic tension. The system includes a definition of narrative tension based on the discovery of clues by the player. The interactive play is driven by minimizing the distance between the actual tension and an objective tension curve.

Stella (León and Gervás, 2011) drives the generation of a step-by-step story generation algorithm by trying to adjust the generation to a set of user-defined curves. Stella allows ad-hoc definition of curves and lets the user decide what they represent. In order to make this option possible, Stella does not make any assumption on the semantics of the curves, therefore forwarding the definition of such semantics to the programmer of the domain definition.

2. Defining Objective Curves

Narrative curve-based generation can be carried out in several ways. Since a curve can represent any variable and since the way in which time is handled in stories is not trivial, many aspects of how curves are used must be taken into account. The current model assumes for simplicity that curves have only two dimensions. The x axis represents the value of the variable that the curve is defining and the y axis represents the evolution of the story over time. Section 5. discusses how more dimensions could be added, but for clarity we will be using this simplification.

2.1. Values for Curves

Being a mathematical object, a narrative curve can represent any variable and thus any numerical value. For this explanation we are assuming real numbers. The semantics for the value that the curve takes depends on the computational system that uses it. No constraints are set on these variables because imposing restrictions on the kinds of value that the curve can represent would only limit the expressiveness of the model.

For instance, a curve could represent *danger*. Intuitively this would be representing the amount of perceived danger for the protagonist of the story by the audience. This value could vary in the range $[-1, 1]$, $[0, 1]$ or $(-, +)$, for example (the number of ranges is infinite).

Of course every option has a number of implications. If a finite range is used, the values are limited and there is a notion of *maximum* or *minimum* is the values that the variable can take. It could also be the case that negative values represent the opposite semantics, i.e. the protagonist is not in a dangerous situation.

All these options are not critic by themselves, and this model assumes that the computational system using the curves must define the required semantics and therefore the most appropriate definition of the variables. As can be seen, this model tries to be as general as possible in terms of the types of curves that can be defined.

2.2. Types of Curves

Having chosen a curve as the definition of the evolution of some aspect of the story along time, the generation must try to create a story whose corresponding evolution is similar enough to the objective curve. When a curve is defined and it is meant to drive the generation process in a story generation algorithm, it must be decided how to make the story match the curve.

While the generation is taking place, the story that is being generated must be implicitly or explicitly assigned a curve that must match the objective curve. However this is not straightforward because there is a number of possibilities for comparing curves beyond the pure numerical comparison. This types of matching are semantical and imply an interpretation of the curve.

This model proposes three levels of decision that define 6 types of curve-based generation:

- *Matching type*: the way in which the matching is performed is important. Curves can be matched by a strict comparison of its corresponding sequence or by a relative matching in which only the relative changes are compared.
  - absolute – every point in the objective curve corresponds to a computed value for the story, and it corresponds to one exact point in the story curve. Both points must have the same y value. This kind of curve gives a strict control on the generated curve, but on the other hand the length of the story must be known beforehand or it the generation must be expected to produce a story which is exactly the same length as the objective curve.
  - relative – in a relative matching type only the order of changes matters. The story curve and the objective curve have to match (within some error threshold), but only the relative position of the values in the curve is important. This kind of curves permit a looser definition but require a more complex computation.
- *Time*: handling time in narratives is usually very complex and it is subject to a deep study in the Narratology community. Regarding curves and computational generation, a simple division in which physical and relative changes of time has been made in order to avoid complicated details.
  - physical – this type assumes that the y axis in the curve specifies absolute time (in whatever unit). This means that the physic timeline of the story (the time at which events take place in the story) has to match the values of the curve at its corresponding time. It has to be kept in mind that this kind of time matching does not strictly use real time, but concrete *time units* that can be defined by the story generation system.
  - changes – these curves can be matched against the sequence of variable changes and not against the time in which they happen. Again, this is a loose definition and while it lets the generation produce stories in a less restrictive way, it requires a more sophisticated algorithm.
- *Level*: so far the theoretical definition of curves has assumed the existence of a general definition of story. While this is hard to study, in general the literature on Computational Models of Narrative shows that most systems have some sort of division between *plot* and *discourse*. The discussion of the narratological definitions is way beyond the focus of this paper, but plot can be defined as the list or graph of events taking place in the story world, and discourse as the ordered and filtered representation of events in a linear way, pretty much in a form that somehow resembles a textual narrative.

  According to this, curves can be used for driving the story plot, the discourse, or both. Hence this division:
  - story plot – these curves represent the evolution of the plot or the underlying, logic sequence of actions and events in the story.

- discourse – curves can be matched again the discourse and not the story itself. These curves take the discourse into account.

A more complex model based on the proposed division could be designed. A model in which a curve could be absolute from the beginning and then the matching could switch to relative matching at some point could be defined and implemented, but the model would only become more complex and not necessarily more general.

3. Curve-Based Generation

After the objective curve has been defined, a story generation system must use this information as a partial definition of the objectives. At this point the story generation system has a curve Γ. The system must also have a domain-specific function $\Phi :: S \rightarrow [(\text{partial story}, val)]$ yielding, for every partial story, a value for some particular variable like *love* or *danger* in any comparable domain (real numbers, for example).

The generated content must correspond to a curve that matches the objective curve Γ according to any selection of the matching type (as described in Section 2.2.). This curve is generated by Φ.

At some stage of the generation the story generation algorithm must decide what to include in the partial story. This can be formally described as follows: the partial story is a list of events $S = [e_1, e_2, e_3, \cdots, e_n]$, and the event e_{n+1} can be chosen from a set of potential candidates $C = \{c_1, c_2, c_3, \cdots, c_n\}$.

This selection must be done according to the objective curve Γ. Since given the definition the story generation system includes the domain-specific definition of Φ, it is possible to compute $\Phi(S + c_i)$ for every $i \in [1, |C|]$. This creates a set of partial curves among which it is possible to choose the best candidate.

The assumption of the possibility of the comparison of curves is supported by several mathematical approaches to curve comparison (Buchin et al., 2009; Cui et al., 2009). We can therefore consider that several mathematical solutions for computing the level of matching exist in the literature. A review of these algorithms is obviously beyond the scope of this paper.

According to the proposed definition, we can assume that as long a matching function between two curves is provided, a computational system can use it to generate a story of which the corresponding curve matches the objective curve.

In this way the generation can be automatically driven by the objective curve Γ. The quality of this generation obviously depends partially on the quality of the domain definition of the Φ function. The implications of this are discussed in Section 5..

3.1. Error Threshold

The value of the computed value for a partial story $\Phi(S + c_i)$ will probably be slightly different from the corresponding value of the objective curve Γ at best. This enforces the inclusion of an *error threshold* in the theoretical model.

This error threshold depends on the particular instance of the generation and sets the allowed difference between the objective curve and the value of the Φ function at some particular stage of the generation. The bigger the error threshold the less strict the matching to the objective curve. This additional variable helps to constrain the generation by setting the level of similarity that must exist between the objective curve and the curve corresponding to the story.

4. Application of the Theoretical Model to Existing Story Generation Systems

The proposed theoretical definition of curve-driven story generation can be used to classify existing story generation systems.

MEXICA (Pérez y Pérez, 1999) performs story generation by matching curves following an *absolute* matching type. Regarding time, it uses the *changes* pattern and it takes places at story level.

The study can also be applied to Interactive Fiction. The system by Barros and Musse (Barros and Musse, 2008) uses curves that are matched in an *absolute*, *physical* way and at *discourse* level.

Stella (León and Gervás, 2011) generates stories by using curves that it tries to match by *relative* comparison. The time matching is *physical* and the curve-based generation is carried out at story level. Stella admits curves representing any variable and not only *tension*.

5. Discussion

The proposed theoretical system for driving automatic story generation by the use of plot curves is not enough by itself. Defining objectives is much more complex than the representation of some aspects of the final story. Objectives can be required to be able to enforce a specific action or event to occur in the story. While a careful design could let plot curves be that specific, they are most likely not the best option for that.

In that sense, plot curves can only prototype a certain set of features of the desired story. If a certain event must be present in the resulting story, a logic condition on the objective state of the state space search would be a easier option (assuming that the story is created by searching). In general, the most expressive definition of generation objectives would probably contain several structures: curves, length constraints, a set of events to be included and so on.

Along this paper it has been assumed that curves are one-dimensional. While this can be the case, it makes sense to devise a system in which curves can have more than one dimension. For instance, instead of having a curve that represents a value (like *love*) against time, it would be possible to have that curve that evolves over time and also depending on the number of characters, thus adding a dimension to the Φ function. While this expansion of the current model is beyond the scope of the paper, further work contemplates it.

Section 2.2. describes a classification of curves. The authors are aware that other classifications are possible. This proposal focuses not on the general aspects of curves and their appropriateness for representing data, but on those aspects of automatic story generation that must be represented in most cases. This list has been designed based on an analysis of the literature and on the author's experience, but it is possible to think of a different taxonomy.

The proposed theory is not complete by itself in a specific Artificial Intelligence system. It assumes that the Φ function can be defined in a domain-specific manner. While this is true in general, it must be taken into account that this definition can be very complex. For example, a narrative generation program usually must be able to handle a certain set of semantic symbols that represent some domain. Every implementation of this theoretical model should therefore be able to provide an implementation of Φ able to yield a value for every symbol. Additionally, the impact of some specific symbol will be context dependent. Therefore the implementation of the domain dependent part of the model is complex.

6. Conclusion

In this paper a theoretical definition of plot curves to drive story generation in a automatic story generation system has been presented. The relation of this model with existing story generation algorithms has been explained and the benefits and drawbacks have been discussed.

So far the model has proven to be valid in programmed prototypes and to create a simple taxonomy for story generation systems. More work has to be done in order to demonstrate the general validity of the propose theory.

As previously discussed, the model can be refined. Currently the research on this model is focusing on trying to discover general aspects of curves regarding their semantic properties in order to lessen the amount of effort needed to create a domain-specific implementation of every variable. While some systems base the curve description on the single idea of *narrative tension*, the authors think the generation can be much richer if several variables guiding a narrative are taken into account.

7. References

Aristotle. 1974. *Poética*, volume 8. Colección Biblioteca Románica Hispánica IV.

Leandro Motta Barros and Soraia Raupp Musse. 2008. Towards consistency in interactive storytelling: Tension arcs and dead-ends. *Comput. Entertain.*, 6:43:1–43:17, November.

Kevin Buchin, Maike Buchin, and Yusu Wang. 2009. Exact algorithm for partial curve matching via the Fréchet distance. In *Proc. ACM-SIAM Symposium on Discrete Algorithms (SODA09)*, pages 645–654.

M. Cui, J. Femiani, J. Hu, P. Wonka, and A. Razdan. 2009. Curve matching for open 2d curves. *Pattern Recogn. Lett.*, 30:1–10, January.

R. Raymond Lang. 1999. A declarative model for simple narratives. In *Proceedings of the AAAI Fall Symposium on Narrative Intelligence*, pages 134–141. AAAI Press.

M Lebowitz. 1985. Storytelling as Planning and Learning. *Poetics*, 14:483–502.

C. León and P. Gervás. 2011. A top-down design methodology based on causality and chronology for developing assisted story generation systems. In *Proceedings of the 8th ACM conference on Creativity and cognition*, C&C '11, pages 363–364, New York, NY, USA. ACM.

J. Meehan. 1976. *The Metanovel: Writing Stories by Computer*. Ph.D. thesis, Yale University.

R. Pérez y Pérez. 1999. *MEXICA: A Computer Model of Creativity in Writing*. Ph.D. thesis, The University of Sussex.

M. Riedl and R. Young. 2006. Story Planning as Exploratory Creativity: Techniques for Expanding the Narrative Search Space. *New Generation Computing, Computational Paradigms and Computational Intelligence. Special Issue: Computational Creativity*, 24(3):303–323.

D. E. Rumelhart. 1975. Notes on a schema for stories. *Representation and Understanding: Studies in Cognitive Science*, pages 211–236.

R.B. Tobias. 2012. *20 Master Plots: And How to Build Them*. F+W Media.

Scott Turner. 1992. *MINSTREL: A Computer Model of Creativity and Storytelling*. Ph.D. thesis, University of California at Los Angeles, Los Angeles, CA, USA.

Nelson Zagalo, Anthony Barker, and Vasco Branco. 2004. Story reaction structures to emotion detection. In *Proceedings of the 1st ACM workshop on Story representation, mechanism and context*, SRMC '04, pages 33–38, New York, NY, USA. ACM.