

3 The Craft of Stories

3.1. Chapter Preface

As IS is such a wide and interdisciplinary area, it is important to focus not only on which approaches have been tried in research, but also on what exactly constitutes the research field and what its goals should be. Based on the previously review of the area, this chapter presents what we feel is the main problem in story craft, along with some considerations on its breakdown on more manageable chunks/sub-problems.

The text contained here is the second part of the article co-authored with Fábio Guerra and Antonio Furtado, and published in 2009 in the Monografias em Ciência da Computação Series (MCC 36/09), at PUC-Rio.

3.2. Introduction

In this chapter we will discuss the central research topic of this work. In Sections 3.3 and 3.4 some related work is mentioned with an emphasis on the attendant terminology, and in Section 3.5 we introduce and motivate the term that we chose to designate the topic. Then, in Section 3.6, the fundamental problem of this study is presented, and in Section 3.7 we discuss a sub-division of the problem into smaller problems. In the last section we offer a few remarks resulting from this discussion.

3.3. Interactive Storytelling

Numerous terms are used to define the field of study in which the generation of stories fits in, the most popular and widely used being *Interactive Storytelling* [Glassner 2004, Crawford 2005].

This term is usually applied in the context of digital games, where, in most cases, the story is already defined before the game begins. In this case, the only role for the computer to play is telling the story and its variations to the player.

Even though the term is adequate to digital games, its use is arguably correct in other applications, especially in applications where narration is not the most important aspect, the act of telling stories being part of a larger process.

Most of the work published in the area is aimed at applications to digital games. In this context, the stronger emphasis is on narrating stories in ways that augment player immersion in the virtual world [Murray 1997, Glassner 2004, Crawford 2005]. But this work sees narration as just part of the whole process,

highlighting important conceptual aspects affecting not only the narration but also the generation of stories.

3.4. Narrative Intelligence

Another term that is widely used is *Narrative Intelligence*. According to Mateas and Sengers [1999], the study of Narrative Intelligence tries to relate the usage of narratives in human experience with their applications in Artificial Intelligence.

This is a very wide area, encompassing from the application of narratives to the design of user interfaces to complex systems to interpret stories, including interactive fiction.

The subject of this report can be categorized as part of Narrative Intelligence, which in turn can be categorized within Artificial Intelligence. Even so, the present work concerns a much more specific problem, albeit considerably wide and complex, demanding a great deal of research effort from academia until its main problems can be well understood.

3.5. Story Craft

We propose the use of the term *Story Craft*, defined as: the art and science of applying scientific and common knowledge and techniques to the conception, generation, and narration of stories. Another option would be *Story Engineering*, but, since the approach adopted so far is not so formal, and in view of the consideration given to individual tastes, we feel “craft” would be more fitting than “engineering”.

On using this nomenclature, we posit that the generation of stories must be seen as a skilled craftsmanship process, whose methods and techniques to find solution to the problems involved must be precisely (if not formally) specified. In order to tackle this objective, it is necessary to have well defined models concerning what is a story and what is a good story.

As previously discussed [Borje et al. 2009b], it is not possible to define an ideal or sound model for a good story, as the quality of a story depends heavily on tastes and expectations of those watching or participating in it. One of the goals of story craft, therefore, is to find models that might be used in specific situations, fulfilling the expectations of a good percentage of the interested audience.

Different techniques must be used for the generation of stories according to which model is utilized. The difference between these techniques should then reflect the differences in ‘style’ among the story generation systems based on each model.

3.6.The Fundamental Problem

The fundamental problem in Story Craft, as described by Guerra [2008], can be summarized as follows:

Problem 1: Given a knowledge base subjected to a set of constraints, a computer system should be able to generate and tell stories to a given audience, with or without interaction, in a way that satisfies the constraints.

The generated stories may be original or not, and must obey whatever well defined constraints that may have been prescribed. The necessary data that must be in the knowledge base can, of course, vary from application to application. How the data is gathered, stored, queried, and processed is another issue that must also be addressed.

In order to help understanding the problem, some examples are shown of specifications that one might wish to pass to an ideal process of generation of stories:

- A story with emotion and suspense, that takes place in Rio de Janeiro during the 1960s, featuring a schizophrenic character and leading to a surprising ending;
- A story geared towards seven year old children, that helps them comprehend the geography of Brazil's centre-west region;
- An interactive story where the user plays the role of the protagonist and must try to explain a mysterious series of murders.

Currently, there is no technology or know-how to, on the basis of just these descriptions, automatically generate quality stories. This is a very hard problem and constitutes one of the great challenges for Artificial Intelligence investigation in this century.

3.7.Sub-problems

As in any process, the first step to take is to get to know better the problem and its inputs and outputs and, only then, look for a strategy for its resolution. Many considerations can be made starting from the problem definition. A first observation is that there is a big difference between the story generation and narration phases.

In the generation phase a story description is produced, containing enough information for starting its future narration. In the narration phase, stories are cast in the form they will be told, more often than not already considering the peculiarities of the specific medium in which they will be presented. It is also at this stage that possible interactions with the audience must be anticipated.

During the rest of this chapter the term (story) *generator* is used to refer to the module responsible for generating stories, and the term (story) *narrator* for the module responsible for the narration of the stories. The term storyteller is not used here to avoid possible misunderstandings, as it is already used with different meanings in different contexts.

Note that any change to the story is performed by the generator. If any unplanned change to the story is made, the narrator must interact with the generator. This way, the narrator can be seen as a mediator between the generator and the audience, which allows for greater flexibility in handling different approaches for audience interactivity.

An example of the passage from generation to narration can be observed in the movie industry. In this kind of production there is a neat separation between the stages of *scripting*, where it is defined in detail how the story develops, and *execution*, where the story is then produced in its final version. In this context, each stage is usually performed by totally different teams.

Another problem in the craft of stories resides in the transfer of knowledge between human authors and the computer. A SGS needs access to some knowledge base that describes how stories may be generated and told. What is this data and what is the best way to build such knowledge base are some of the issues discussed in Section 3.7.3.

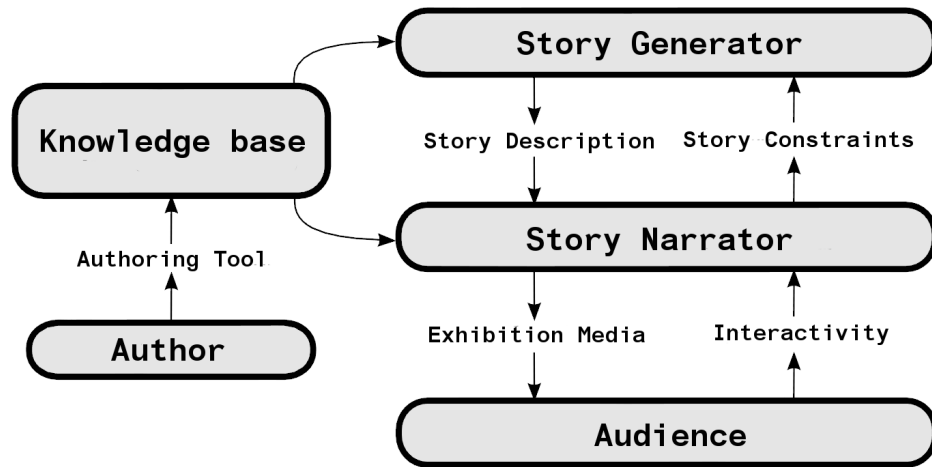


Figure 3.1: Guerra's [2008] schema for the generation and narration of stories

We propose a conceptual separation of the fundamental problem of automatic story craft into four basic sub-problems:

- Story generation;
- Story narration;
- Knowledge base representation;
- User experience evaluation/mediation.

These four problems will be discussed in more detail in the next sections. This separation and the inter dependencies among the separate problems can be more easily visualized in Figure 3.1.

It must be clear that Figure 3.1 refers to the fundamental problem of story craft in an abstract form. It does not refer to a specific implementation.

The fourth problem, user experience evaluation (or mediation), was not contemplated in Guerra's original formulation; in the figure, it relates to the audience box. As briefly mentioned in [Karlsson et al. 2009b], when describing some systems that try to model user experience, this problem covers a wide range of issues: user modelling; behaviour analysis; plan recognition; user/player/audience points of view; psychology profiling; motivational feedback; emotion and moods; player classification; evaluation of dramatic qualities; etc. Even though determining what makes a story (or its parts) more interesting and rewarding to the audience is an important issue and interesting in itself, further detailing what is needed to properly treat those items is beyond our present scope (except for the brief description already mentioned).

The knowledge base concerns all the available knowledge data, not necessarily stored in the same place or available to every program module. Ideally it should be properly assisted by authoring tools to help prospective authors in creating the story space with less effort.

The generation of stories can be performed in a sequence of steps, the same being true for story narration. Also, an alternation scheme can be adopted, with each generation step followed by a corresponding narration step. The best way to organize this process depends entirely on the needs of specific implementations.

Still, it is important that generation and narration be handled separately. Thus research efforts will be able to focus only on sub-problems pertaining to one or to the other, without the need to address the whole problem at the same time. It is important to note that separating narration and generation does not mean that there can't be an interchange of information between the modules responsible for each sub-problem.

Nowadays it is very common to see the creation of story generation systems whose performance is limited to very specific subjects or scenarios. This makes it difficult, among other inconveniences, to compare different works in story generation. It is important that the research community develop tools that allow reuse across different story generation systems (Guerra [2008] provides a case study on a framework with this issue in mind).

3.7.1. Story Generator

A story generator can be seen as a tool for the creation of stories, at the fabula level. The generator must guarantee that stories satisfy any specified constraints.

A simplified representation of the story generator comprises the generator method itself, the knowledge base to be accessed and the constraints that must be obeyed. The output must be a (simplified) representation of the generated story. For convenience, it is assumed that all necessary knowledge data is available in a single knowledge base.

It is important to stress that the description of the story at the fabula level will always be simplified, since the amount of information necessary for a description fully matching a human author's intuition would be too big for representation in digital media. The degree of simplification, however, varies according to the needs of each application.

Depending on the application it may be possible for the generator to create only story pieces, instead of complete stories. A partial description can be particularly useful to allow interactive generation, whereby the full story is only completed after multiple calls to the generator combined with user interventions.

In addition to generating the sequence of events (fabula level), restrictions can also be set to prepare for the narration. However, the final presentation order and timing of the events can be stipulated by the story narrator, as long as it does not violate any of the constraints imposed by the generator.

3.7.1.1.Plot Manager

Plot managers are interactive story generators, which receive as input an already executed piece of story and present as output a suggestion on how to continue the story. Its main function is to effect corrections so that the story plot can develop correctly. If some inconsistency is detected in the input, the manager intervenes in order to reach the original goal, satisfying all the necessary constraints [Mateas and Stern 2003, Roberts and Isbell 2007].

These corrections can be made by, for example, forcing the execution of some events or preventing others from happening. A frequent source of inconsistencies in plots is user interaction, especially when the story narrator allows a high interactivity level.

Plot managers are widely used in a special genre of digital games called interactive drama. In this context, they are called drama managers. A good description of the role of drama managers can be found in the work of Roberts and Isbell [2007].

3.7.2.Story Narrator

A story narrator can be seen as a mediator between the story generator and the audience. It must be able to receive one or more stories from the generator and, possibly, some constraints on how these stories must be told. Having access to a

knowledge base, it should manage to tell these stories to an audience, through a pre-determined exhibition medium.

By definition, the function of the story narrator is simply to narrate stories delivered to it. But if by any reason (e.g. user interaction) the input stories become non-valid, the narrator must autonomously ask the story generator to try and fix it or create a new valid story.

Conversely, it is important to recall that the generator/author role does not need to be at all times played by automatic software facilities. The user himself shall always have the option to assume this role.

3.7.2.1. Story Representation

There are multiple ways in which stories can be represented. For example, the story narrator may receive the story under the form of a tree or forest data structure, in which each node represents a possible variation of the story and each edge represents user interaction.

A story can also be represented through a simple sequence of actions and a partial or total ordering. These actions can be grouped into sequences of actions or scenes, which in turn can be grouped into acts. Thus the story can ultimately be viewed as a hierarchy of actions.

Regardless of the chosen representation, the narration should be coherent with the input data, and, if interactive, the narrator must know how to deal with audience interaction in the exhibition medium.

3.7.2.2. Exhibition Media

Every story narration happens through some exhibition medium. Stories can be told, for example, under the guise of comics, text, animation, or digital games, which in turn, according to Apperley [2006], can utilize different “video-game genres”, such as simulation, strategy, action, or RPG.

In the case of comics or film, the interaction is limited to moving a little bit forward or backward, pausing, or resuming. We can say that this is a passive type of interaction where the user cannot change the story generated. In contrast, for example in some digital games, the user can change the story being told, thereby performing active interaction. Active interaction is particularly interesting, because it increases public participation in the story.

On the other hand, this higher degree of interaction can generate inconsistencies in the plot. There are two ways of dealing with the problem. The first is to install a plot manager (as explained in Section 3.7.1.1) to remedy the inconsistencies and the second is limiting the power of user interaction.

There are games that restrict the interaction of a player to a few actions (limiting interaction opportunities and possibly affecting user engagement), and other games that give greater freedom to player interaction (thus potentially diminishing the dramatic properties of the experience). It is also part of the role of the story narrator to define the degree of freedom that the player will have when interacting with the story.

Story narrators need to answer two main questions:

1. What actions are allowed to be taken by the audience?
2. At what time can these actions be executed?

A typical “story-focused” digital game allows few actions to be performed by the user, and only at specific moments of the game. The higher the level of interactivity, together with a wider variety of actions, the more difficult it is to guarantee a coherent and interesting narrative. But this is no drawback to more open-ended games which do not adopt a plot-based approach.

3.7.2.3. Interactive Narratives

Mateas and Stern [2003] indicate that there are two main approaches to the interactive narration of stories: structured narratives and procedural simulations.

In the first approach, only a small set of actions is allowed when the audience is interacting with the story, which makes it easier to avoid inconsistencies due to interactions. Thus the stories generated this way tend to have a more coherent plot, but the creation of believable characters becomes more difficult. This is due to the fact that instead of having the characters “striving” for their personal goals, they are often obliged to take “forced” actions in order for the story to have a satisfactory closing.

The procedural simulation approach is very popular among current digital games. It consists of simulations of a virtual world with numerous agents interacting with the player. In this approach it is much easier to have credible characters, as their behaviour is not restricted by plot-related constraints, thus possessing much higher autonomy. The generated “emerging narratives” [Aylett 1999] are the simple result of the player's interaction with the characters of the story, and as a consequence may be poorly structured.

Indeed the distinction between plot-centred and character-centred narratives, which is widely recognized, appears to be similar in nature to the distinction between structured narratives and procedural simulations [Riedl 2004].

There are several studies that seek to establish systems centred both on plot and on characters. One way of achieving this double requirement is through the use of plot handlers configured to correct the possible failures resulting from characters’ actions in the story. An example of such systems is Façade, which was

described in [Karlsson et al. 2009b] and, in more detail, in the work of Mateas and Stern [2003]. A possible alternative is to include elements in the story generator that make characters more convincing and credible. An example where such approach is successfully utilized is Mimesis, reported in detail in Riedl [2004].

3.7.2.4. Adaptation

Stories can be adequately adapted to the different media where they may be presented, and also to try and satisfy some audience-defined restrictions and constraints.

For example, one may wish to exclude from a presentation any scenes of violence or nudity that may figure in the description of a story. Or yet, one may prefer to emphasize fight scenes and comic situations. In general, it is possible to show stories in ways that may please a larger number of people.

The adaptation of stories is performed by a special kind of story generator, here called *adaptor*. The main difference between an adaptor and other story generators is that it receives as input a previously created story together with a new set of restrictions. As output, it presents a possibly much modified story that satisfies the additional restrictions.

3.7.3. Knowledge Base

The knowledge base represents the set of all available data for the generation and narration of stories. A story craft application must be equipped with means for the representation, acquisition, storage, and access to this knowledge.

There are multiple ways in which stories may be represented. For example, as mentioned before, the narrator can receive the story as a multi-level hierarchy of actions.

Also, one can represent characters via behavioural rules or some sort of character models. Information on the virtual world, restrictions of the chosen literary genre, user models for interaction with the audience, amongst many other types of information, perhaps even graphical assets, can also be part of the knowledge base.

Having authorial tools in place to make the underlying complexity of the knowledge base transparent to authors is a very desirable goal, especially if the representation uses complex formal models or specialized programming languages. There is not much perspective for a system that places too great a burden on potential authors.

For example, although Façade has been a successful experience, its architecture requires a great effort from the prospective authors. It uses four different content languages; and two years were spent just for authoring a game

that has only one scene, two characters, and takes about 20 minutes to complete [Mateas and Stern 2003].

3.7.3.1.Types of Information

Many further details are important for the generation and narration of stories concerning, for instance, the models and techniques presented in [Karlsson et al. 2009b].

The amount of information depends on how specific or generic is the literary genre of interest. For example, in stories of chivalry, the characters can wear armour and fight with swords, but the use of motor vehicles and firearms would be incongruous. Moreover, to generate or adapt a story that includes passages through different times and wide spaces it is necessary to have available information on their diverse technological and cultural characteristics.

A simple knowledge base can store information about the specific actions that each character can perform. Axioms with logical propositions (what must be valid within the story world) can also be stored, besides the facts that must hold at the beginning of the story.

Pozzer [2005] proposes to model a literary genre through the use of three conceptual schemas: static, dynamic, and behavioural. The first determines what states are valid in view of the conventions of the literary genre. The dynamic scheme indicates the actions that can be performed by the characters, and defines their pre-conditions and their effects in the world. The behavioural scheme tells how the characters are supposed to react to certain situations, which motivate them to pursue goals compatible with their assumed roles.

3.7.3.2.Story Repository

Studies on the creation of stories teach that reading literary works is a great help in writing books or film scripts [Field 1982, Vogler 2007, McKee 1997, Howard and Mabley 1995, Aarne and Thompson 1961, Uther 2004]. It is convenient then that the knowledge base be able to include a repository of stories.

Some pieces of past stories can be used and adapted to generate new stories. Examples of such usage are systems that apply case-based reasoning over recorded plots [Fairclough and Cunningham 2003] and the use of a library (or hierarchy) of typical plans inspired in motifs and story functions to support the generation of stories [Karlsson et al. 2006a, Karlsson et al. 2006b].

Borrowing from existing stories is a standard practice, leading to stories that share common elements (intertextuality). But it is also desirable that the methods for drawing from a repository can manage, as much as possible, to avoid the generation of stories that are too predictable.

3.7.3.3.Reuse

Almost every SGS discussed in this study uses very restricted knowledge bases that can only operate with specific genres. In addition, each SGS usually comes with its own knowledge base implementation, both in terms of structure and data contents, which are often not described in any level of detail. That practically makes any direct comparison impossible between the various existing story generation algorithms.

Although there are similarities between SGSs, no effort was found in the surveyed references to draw up a common knowledge base, which would be rich and flexible enough to be shared with other SGSs. One should also notice that the quality of the generated stories is still very dependent on the quality of the knowledge base used.

3.7.3.4.Common Sense Knowledge

Some stories involve lots of events from people's daily lives. The generation of stories like these will only be possible with the use of common sense knowledge. This kind of knowledge is based on "obvious" information that people have been learning through experience [Minsky 2006]. For example, we all know that walking with untied shoelaces can cause one to trip and fall.

Few studies in the available references make use of common sense knowledge reasoning techniques in the generation of stories [Minsky 2000, Liu and Singh 2002]. But for the creation of more comprehensive SGSs, it is an inevitable necessary step to have access to such knowledge.

3.8.Some Remarks

In this chapter the term Story Craft was suggested to define an area of study that is of great importance to digital entertainment.

The fundamental problem of crafting stories was presented, and the four sub-problems that must be addressed were discussed: the generation of stories, the narration of stories, the manipulation of the knowledge base, and the problem of evaluating user experience (although the latter was only briefly treated). Some of the issues, approaches, and possible developments of the main problem were also presented.

It became clear to us that this is a difficult problem to tackle and that the current attempts towards viable solutions are still far from satisfactory results, raising a strong demand for research projects in the area.

Interactive storytelling (as a general area of research), although studied for some time, has regained considerable interest recently. Even though studies on

the subject have yielded a number of promising approaches, ideas, and experiments, it is still a problem far from settled.

Most of the research on IS has focused on developing systems for specific “experiences”, i.e. to a simple specific story experience; important examples are Façade [Mateas and Stern 2005] and Mirage [El-Nasr 2007]. The somewhat narrow scope of those efforts confirms that IS is a complex problem, and that it may be advisable to invest on more controlled experiments. Also, most of the approaches presented here fail to break down the problem into manageable sub-parts, which would have made it less difficult to reuse components, or at least to compare different initiatives.

Except for a few soap-opera-inspired systems [Lebowitz 1985, Barber and Kudenko 2007, Thue et al. 2007, Weyhrauch 1997], there are few generative approaches and few attempts to present a conceptual model for a given genre (one such ongoing effort is the LogTell-R prototype [Furtado 2008, Karlsson et al. 2010a], featuring a conceptual model and a model-based generative system).

Another interesting area that has been amassing efforts is the management of user experience in an IS environment. As the presentation of the different approaches to user experience in [Karlsson et al. 2009b] suggests, this includes evaluating pieces of stories or continuing episodic storylines (“story arcs”), but – perhaps more importantly – modelling the user intentions or motivations while the system is running in order to improve the “interestingness” of the story and users’ engagement.

Even after defining and breaking down the problem (proposing possible structures to represent story knowledge and user models, how to create/populate such models, and story generation per se), a lot of effort remains to be done in the area of dramatization or visualization of the stories. The final objective should be to provide to the target audiences some sort of engaging and immersive environment.

Efforts towards automatic background music generation [Casella and Paiva 2001], light control for dramatic intensity [El-Nasr et al. 2006], animation systems that convey emotions of the characters [Perlin and Goldberg 1996], and intelligent camera placement [Passos et al. 2008, Lima et al. 2009] are additional indications of the broad scope of the area.

Pervading all the above research lines is a concern with the development of tools, especially for non-technical users. Little attention has been given to the creation of tools for designing, populating and exploring the supporting knowledge bases and, in general, to help prospective authors at the various stages of story composition and adaptation. It is unrealistic to expect that the field of IS will effectively take shape and gain popularity while there is scarce

user-friendly support to the creation of interactive content by prospective authors.