

Integrating Semantics and Narrative World Generation

Ben Kybartas, Rafael Bidarra, Elmar Eisemann
Computer Graphics and Visualization Group
Delft University of Technology
Mekelweg 4, 2628 CD Delft, Netherlands
{B.A.Kybartas|R.Bidarra|E.Eisemann}@tudelft.nl

ABSTRACT

Implementation of a narrative world is often laborious and remains fairly independent from working directly with the narrative itself. Currently, there are few tools able to simplify this process. The aim of our research is to support this process by developing a declarative approach for generating narrative worlds. A set of narrative semantics will be used to discover and further generate appropriate content. The declarative approach and semantics will be initially implemented and validated in a virtual therapy tool for patients with post-traumatic stress disorder. The project will evolve towards the creation of a more general narrative creation tool.

1. PROBLEM INTRODUCTION

Computers are often used to assist with the creation and representation of narrative. From serious medical software to more artistic endeavors such as games and film, there are a broad range of ways to express, interact with, or convey narratives to an intended audience. Unfortunately, it is difficult and often tedious to create these narrative worlds. Frequently, professional techniques such as modeling, animation and physics simulations are required, limiting the potential of non-professional users to implement their desired narrative without a significant level of training.

We are interested in exploring the potential of procedural content generation methods combined with a declarative approach for modeling narratives to simplify the process of representing narratives within virtual worlds. Procedural systems have been successfully used to generate a variety of content from trees, to cities, to full game worlds [13]. While simplifying various design processes, there are still limits to the capabilities of many of these systems. For instance, there is often still a significant level of training required to understand and effectively use most generative tools. Within our research, we are aiming at methods which have a high expressive power, and yet are very accessible and easy to use.

2. RESEARCH PROPOSAL

In this project, we propose investigating the possibilities for a declarative approach to modeling narrative worlds that is simple and powerful to use. The focus will be upon coupling a declarative narrative with procedural content generation, with the goal of taking the declared narrative and generating it within a virtual environment. Within this generation process, we must have an understanding of what is required to create the declared narrative. Here, we explore the notion of creating a semantic model relating specifically to narrative. This leads to three main research phases, which we present below.

- *Declarative techniques for defining narratives.* This involves formalizing narrative as a sequence of events relating to objects in the virtual world. The goal is to create an expressive, yet simple, means for specifying and representing narratives. This representation should be intuitive and intelligible to people without experience in the field of declarative modeling.
- *Story Semantics.* In order to generate the final narrative, there must be a low-level understanding of what is required to create the given narrative, and how different pieces of the narrative interact with each other. We propose using condition analysis to understand what is needed by, or changed in, each part of the narrative.
- *Narrative World Generation.* A goal of this research is to develop a tool wherein the user may iteratively declare their desired narrative, and see the generated result. The initial intention is to create a tool for patients suffering from post-traumatic stress disorder to recreate traumatic instances in their past as part of an online therapy program.

Our research differs itself from much of the existing applications of procedural content generation in the field of narrative technology. Much of the existing research has focused on plot generation and direction, as surveyed by Gervás in [4] and Riedl and Bulitko in [10]. Our research instead lies in the generation of narratives within three-dimensional worlds, similar to the *text-to-scene* tools [3] and PASTe's director controlled animation environment [12], but with more focus on the mapping of the narrative onto the virtual environment.

In the remainder of this section, we present a preliminary exploration into the terminology that may be used to define

the declarative narrative. We conclude with a basic example to clarify some of the concepts presented in the previous sections.

2.1 Terminology for Declarative Narratives

Narrative research within the field of computer science formally models narrative using techniques from narratology [1]. While our research lies within generating narrative worlds, as opposed to the narrative plots, maintaining a similar terminology would allow us, for example, to combine approaches to develop a more encompassing narrative generation technique which generates worlds and plot.

We are considering using a similar definition of narrative in our own work, based upon the formal model of narrative frequently used in plot generators [4], as well as some preliminary research into the field of narratology and specifically narrative *spaces* [11] and *plot* [7]. Using this definition we will integrate our definition of narrative with the semantic model provided by the ENTIKA framework [6]. ENTIKA, is an in-house framework that facilitates the specification and application of semantics to a virtual world. Within ENTIKA, *entities* are described as anything which is known to exist in a virtual world. An entity is described as having multiple *attributes*, which characterize the behaviour of the entity. Entities are related to one another using *relations*. These relations may group entities which satisfy a set of conditions into *families*, or by *predicates* relating to how one entity may perceive another. Lastly, *services* indicate an entity’s ability to perform an action given a context. A *context* is defined as being a set of conditions, where each service can only be performed if these conditions are met. The result of performing a service is an *event*, which is an *action* performed by an *actor* upon a *target*. The event, in turn, causes an *effect* on the virtual world.

Below, we discuss our proposal for a mapping of the formal definition of narrative onto the terminology used within the ENTIKA framework. In the first phase of this research, this model can likely provide the terminology needed for a declarative approach for defining narratives.

Entities could be used to represent the narrative definition of *existents*, entities which occupy the narrative space and may be influenced by events [11]. They would be known as *physical entities* within the ENTIKA framework as they refer to tangible objects, and are accompanied by a 3D model when represented in a virtual environment. We may also wish to distinguish between unthinking entities such as a “car” or a “light”, and thinking entities such as “people” or “dogs”. Here it may be beneficial to adopt the terms “prop” for an unthinking entity, and “actor” for a thinking entity, following the style of Aristotelian-based plot generators [1]. The **attributes** may be defined as key-value pairs which relate to properties of these entities. For example, when declaring attributes for a car, we can state that the car *colour* is *white*. Here, the *colour* is the key and *white* is the value.

Relations may be of use when determining how to place entities within the virtual environment. For example, *placement relations* refer to where specific entities need to be located. If we have a “car” entity being used within the narrative, we may want to have a placement relation stat-

ing that the *car* must be on the *street*. We may likewise use relations to define attitudes between actors, so that, for example, two enemy actors frown upon seeing each other.

Services can be used to declare the events which will compose the resulting narratives. The narrative itself, following the definition provided, would involve performing these events in a user-defined order. One extension of the original semantic model could be to add attributes to these events as well. For example, if we have a service which involves a *drive* action, we may wish to state that the driving is being performed *recklessly*. This would specify how the action is performed in our virtual environment.

The **Context** of a service is also of importance. We must provide the *conditions* which specify what must be true within the virtual environment for a particular event to occur. Referring back to the service which provides the “drive” action, we may want to state that only entities with a *drivable* attribute are allowed to be the target entity within the event. Conditions can, of course, be more complex; for example, we may state that, in order to *drive*, the actor must be *in* the target, and the target must be *activated* (ex. the driver is sitting in the car, and the car is turned on).

The **Effects** specify what changes will be made by performing each service. For example, a “drive” event may reduce the energy available to the drivable entity, and also move both the driver and the drivable entity from one position to another.

2.2 Declarative Narrative Generation

The final phase of our research is to have the declared narrative generated within a virtual environment. One of the main issues surrounding the generation step is how to determine the missing content from a user’s narrative. We intend to use the conditions described above to allow for the automatic completion of any content needed to generate the narrative. For example, imagine we are simply given the one line narrative that *the man drives the car*. We can then examine the two conditions defined in the previous section, and see that there must be two prior events wherein the man first enters the car, fulfilling the first condition, followed by an action where the man starts the car, thus satisfying the second condition. Likewise, we may use random values to fill in missing attributes, such as selecting a random colour for the car if none is specified by the user.

Subsequently, the user would be able to explore the resulting visualization of this narrative and refine it as desired. Thus if the random colour of the car given above is undesired, the user may set that, and the virtual environment will regenerate an updated narrative. This allows the user to continually refine and add to their declarations, and let the resulting narrative converge to their desired result. This results in a *mixed-initiative* approach, where the computer and user co-operate in order to achieve the high-level desires of the user, with the computer handling the low-level work. Mixed-initiative approaches have been proposed for plot generation by León and Gervás’ [9] and Charles et al. [2] created a timeline-based navigation tool which allows users to modify plot and see the updates within their game world. Again, our approach is unique from this past research in that

we are focused on the generation of the narrative world.

We are also interested in using conditions and effects to discover conflicting events. For example, if the narrative states that “the man destroys the car”, and then “the man drives the car”, then there is a conflict here since the car that is expected to be driven in the second action is now destroyed. We are also exploring the possible approaches to take when discovering such a conflict. This may include generating missing events, or more appropriately signaling to the user that a conflict has been discovered.

One might also wish to allow the user to declare *filters* in order to modify the way the narrative is generated. The purpose of this is to, for example, simulate themes or emotions within our narrative. Applying these emotions to the representation of virtual narrative, for example by using pathetic fallacy, remains to our knowledge unexplored, although features such as suspense have been examined at the level of plot [14]. We may, for instance, implement a *suspense* filter which applies to the narrative on a whole. We could then use this filter in certain conditions as a conditional statement in order to control the way the scene is generated. If we believe that suspenseful stories should have dimmer lighting, then we may add a condition for the “light” entity which states that *light-intensity* should only be at 50% if the *suspense* flag is set to true. This would mean that if the suspense filter is applied, then all lights would have their *light-intensity* attribute set at half the regular intensity. With this technique, we may define multiple filters, and switch between these, or even combine them, to change the overall feel of the scene.

The end goals of this research involve the creation of a visual narrative generation tool, and thus examining potential issues with the implementation of such a tool is also important. While much of this remain to be explored there are some notable challenges. First, the declarative potential must be fairly expansible, which should result from having a well defined narrative model. Second, the method by which users may declare their desired narrative is also important, allowing for a simple yet expressive way to declare events and introduce entities. Finding approaches to these concerns will be an important focus later in the research. One such approach may be to represent narratives as directed acyclic graphs, where each node within the graph relates to an event. The user could use the directional arrows to indicate the ordering of events, and select the entities from a set list they define beforehand.

2.3 Example Narrative

We now present a sample narrative and discuss how we expect to be able to process and generate the narrative in a virtual environment. We construct the narrative in such a way that it may relate to a traumatic experience of the lead character, possibly resulting in symptoms of post-traumatic stress disorder. The narrative is given below.

The street is crowded. The man drives down the street. The car collides with another car. The man crawls out of the car. The man is placed in the ambulance. The other driver is dead.

We see that the first line is simply setting an attribute of

the street entity, and not relating to an actual event. Thus we can set one of the street’s attributes as being crowded when processing the first line.

Within the “first” event, we are given a driving action with the “man” entity as the *actor* and being on the street as a condition. However, when observing the drive event, we may note that the *target* entity required for the event (here, a car, bike, etc.) is not provided. Within the generation process, we may choose to signal a missing information here, or to simply generate an entity with a “drivable” attribute to be used as the target entity. This may have conflicts within the second event, since the second event uses both a car entity as the actor, and a second car entity as the target. For example assume the system chose a bicycle for the target in the first event. In this case, the second event will cause a conflict, since the man was supposedly driving a bike but now the narrative refers to a car instead. Here we are exploring two possibilities: we may go back to the first event and switch the generated bicycle with a car, or we can fail the generation process and force the user to declare what the man is driving in the first line.

While failing is a simpler option for the system, it is still important in certain cases to head back and modify earlier events to allow for later ones. Take for instance the next line, “the man crawls out of the car”. For this line to be true the car must be in a position that the man is able to crawl out of it. The car cannot, for example, be knocked into a canal by the collision, even though that approximation was possible one event earlier. Thus if the car is in a bad position, we must modify the previous event to allow the later event to happen.

In the second last event, being placed in an ambulance, we would have to generate an ambulance and the necessary events for the man to be “placed in” the ambulance. This shows how a single event in a narrative may actually require many events to be generated.

The final line after this shows another potential conflict. At this point, the writer states that “the other driver is dead”. However if the driver was killed during the collision, then there is a conflict since, when arranging the events by following the narrative, the second driver will only die after the man is picked up by the ambulance. Here it is important to understand the refinement process, since the user will be able to move the “the other driver is dead” line to before the “man crawls out of the car” line and witness the updated narrative, if such a change is desired. Having this dynamic refinement process allows the user to see mistakes and update the narrative in real-time, getting immediate feedback.

3. RESEARCH PLANNING

The first step of this research, currently underway, is to perform a literature review on previous research surrounding the use of declarative techniques and utilizing semantics to assist in procedural content generation.

Following this, in the second step we will focus on extending the ENTIKA semantic model to include a narrative semantic model. Once this model is in place, we can examine the process of generating the content required to represent the

declared narrative in a virtual environment.

The third step of this research is to integrate the implemented system as part of a larger project, the VESP project, which aims to create an online virtual exposure therapy tool for treatment of *Post-Traumatic Stress Disorder* (PTSD). Virtual therapy for PTSD has been shown to be effective in various forms, such as using exposure therapy [8] and even short interactive games [5] as treatment tools for war veterans. The system being developed will allow the patient to recreate traumatic incidents, under supervision, using a declarative approach and a specific model for story semantics. Concurrently to this project, another colleague is investigating an *e-coach*, which will provide support and guidance to the patient, and supervise the patient's interaction with the virtual environment tool.

Following the development of this system, several clinical trials will be performed. These tests will be executed by a third colleague to evaluate the efficacy of our system in contrast to more traditional treatment methods. We will use these analyses to understand where our declarative approach may be improved, and where its limitations lie.

In the final phases of this research, we will explore further extensions into this work. This may include examining creating a more general story development tool, where users can add new services and entities manually, allowing for a much broader range of potential narratives to be generated.

4. CONCLUSION

Computers are a powerful and versatile means of conveying narratives through various forms of media. Using interactivity and simulation, a computer narrative can engage, entertain and even challenge the user in ways which are impossible for previous forms of media to achieve. Creating these virtual worlds, however, may require a deep understanding of everything from animation to physics, depending on what is desired for the narrative. Despite many existing systems to generate or automate features in virtual worlds, there still remains a lack of research into providing high-level tools for directly creating and manipulating narratives.

We are interested in the potential of a declarative approach for creating narratives within a virtual environment. The approach would view narratives in terms of high-level components, based around the terminology provided by the ENTIKA semantic model. We will then explore the implementation of a story semantic model, to analyze and determine what is desired by the user, and what needs to be generated in order for this narrative to fulfill their intent. Lastly, we presented our research plan, in which we will start with a basic implementation of the proposed research and then apply it by implementing a virtual exposure therapy tool for the online treatment of PTSD. In the final phases of the research, the tool would be extended to more general purpose uses, allowing the user to modify the declarative approach to suit their needs and customize the semantic model to allow for tailored narratives.

5. ACKNOWLEDGMENTS

This research has been supported by the Netherlands Organisation for Scientific Research, under project no. 314-99-

104.

6. REFERENCES

- [1] M. Cavazza and D. Pizzi. Narratology for interactive storytelling: A critical introduction. In S. Göbel, R. Malkewitz, and I. Iurgel, editors, *Technologies for Interactive Digital Storytelling and Entertainment*, volume 4326 of *Lecture Notes in Computer Science*, pages 72–83. Springer Berlin Heidelberg, 2006.
- [2] F. Charles, J. Porteous, M. Cavazza, and J. Teutenberg. Timeline-based navigation for interactive narratives. In *Proceedings of the 8th International Conference on Advances in Computer Entertainment Technology*, ACE '11, pages 37:1–37:8, New York, NY, USA, 2011. ACM.
- [3] B. Coyne and R. Sproat. Wordseye: An automatic text-to-scene conversion system. In *Proceedings of the 28th Annual Conference on Computer Graphics and Interactive Techniques*, SIGGRAPH '01, pages 487–496, New York, NY, USA, 2001. ACM.
- [4] P. Gervás. Computational approaches to storytelling and creativity. *AI Magazine*, 30(3):49–62, 2009.
- [5] C. Holmgård, G. N. Yannakakis, and R. Khaled. The games for health prototype. In *Foundations of Digital Games*, pages 396–399, 2013.
- [6] J. Kessing, T. Tutenel, and R. Bidarra. Designing semantic game worlds. In *PCG 2012 - Workshop on Procedural Content Generation for Games*, 2012.
- [7] K. Kukkonen. Plot. In P. H. et al., editor, *The Living Handbook of Narratology*. Hamburg: Hamburg University, 2014.
- [8] S. Leaman, B. O. Rothbaum, J. Difede, J. Cukor, M. Gerardi, and A. Rizzo. Treating combat-related PTSD with virtual reality exposure therapy. In E. Weiss and J. Rubin, editors, *Handbook of Military Social Work*, pages 113–140. John Wiley & Sons, Inc., Hoboken, NJ, 2013.
- [9] C. León and P. Gervás. A top-down design methodology based on causality and chronology for developing assisted story generation systems. In *8th ACM Conference on Creativity and Cognition*, Atlanta, 11/2012 2011.
- [10] M. O. Riedl and V. Bulitko. Interactive narrative: An intelligent systems approach. *AI Magazine*, 34(1):67–77, 2013.
- [11] M.-L. Ryan. Space. In P. H. et al., editor, *The Living Handbook of Narratology*. Hamburg: Hamburg University, 2014.
- [12] A. Shoulson, M. Kapadia, and N. I. Badler. Paste: A platform for adaptive storytelling with events. In *Intelligent Narrative Technologies 6*, 2013.
- [13] R. M. Smelik, T. Tutenel, K. J. de Kraker, and R. Bidarra. A declarative approach to procedural modeling of virtual worlds. *Computers and Graphics*, 35(2):352–363, Apr. 2011.
- [14] S. Ware, R. Young, B. Harrison, and D. Roberts. A computational model of plan-based narrative conflict at the fabula level. *Computational Intelligence and AI in Games*, *IEEE Transactions on*, PP(99):1–1, 2013.