

Towards a Constraints Approach to Generating Personalised Horror Film Variants

Item type	Meetings and Proceedings
Authors	McNeely, M. (Malcolm); Porteous, J. (Julie); Ferreira, J. F. (João); Lindsay, A. (Alan)
Citation	McNeely, M., Porteous, J., Ferreira, J. F., Lindsay, A. (2017) 'Towards a Constraints Approach to Generating Personalised Horror Film Variants' Workshop on Constraint Satisfaction Techniques for Planning and Scheduling Problems (COPLAS) 2017, 19-20 June 2017, Pittsburgh, USA
Eprint Version	Post-print
Publisher	AAAI
Additional Link	https://www.aaai.org/
Rights	Authors can post items submitted on their own personal website or their institution or company's website prior to publication. Copyright AAAI 2016 http://www.aaai.org/Organization/organization.php For full details see http://www.aaai.org/ojs/index.php/aimagazine/about/editorialPolicies#authorSelfArchivePolicy [Accessed: 25/05/2017].
Downloaded	20-Sep-2018 08:03:27
Link to item	http://hdl.handle.net/10149/621163

Towards a Constraints Approach to Generating Personalised Horror Film Variants

Malcolm McNeely¹ and Julie Porteous¹ and João F. Ferreira^{1,2} and Alan Lindsay¹

¹School of Computing, Teesside University, UK

²HASLab/INESC TEC, Universidade do Minho, Braga, Portugal
firstinitial.lastname@tees.ac.uk

Abstract

In this work we explore the use of a Constraint Programming (CP) model for the selection of suitable filmic content to use to output a horror narrative in a way that is personalised to a particular viewer type. We adopt a bipartite representation of narrative where planning is used to generate an outline narrative and then CP is used to build the discourse presentation: the filmic variant that is presented to the viewer. In the paper we overview our model and illustrate it with example horror variants generated by our prototype system. We also present an analysis of sample filmic sequences output by the system along with the results of a small user evaluation which demonstrate the potential of the approach.

Introduction

Interactive Multimedia Storytelling (IS) systems allow users to modify the unfolding of a narrative as it is presented to them. Such systems feature automated generation, and re-generation, of narrative as needed in response to user interaction. It is useful in the context of IS to use a bipartite representation where, after (Young, 2007), narrative is seen as being made up of the following parts: the *fabula* which is an abstraction of the story world seen as a sequence of events; and the *discourse* which is the way in which the story is presented to the audience.

For the generation of the *fabula* a range of different technologies have been used, with the dominant approach in research prototypes being AI planning (for example, see the work of (Aylett, Dias, and Paiva, 2006; Riedl and Young, 2010; Porteous, Cavazza, and Charles, 2010a; Haslum, 2012)) as it: provides a natural fit for representation with narratives as plans; it ensures causality which is important for the generation of meaningful and comprehensible narratives; and it provides flexibility and generative power. However other techniques have been used including Monte-Carlo Tree Search (Kartal, Koenig, and Guy, 2014) and Linear Logic (Martens et al., 2014).

For the discourse presentation of narrative a range of different output media have been used including computer graphics (Mateas and Stern, 2005), computer games engines (Porteous, Cavazza, and Charles, 2010b) using cinematic staging techniques and text (Orkin and Roy, 2012): all of which enable dynamic generation of visual content at run-time. In contrast some work has also used pre-recorded filmic content, such as (Piacenza et al., 2011), as there is great appeal in using film since, despite recent progress in graphics rendering, the visual quality of film still surpasses that of 3D generated graphics. However the use of filmic content is challenging as it can not be dynamically generated on-the-fly. Rather, IS systems have to make use of pre-existing shots¹ which, when suitably tagged using a small number of semantic categories (e.g. characters, location, mood, activity), can be used as flexible units that can be recombined in order to present the narrative content to an audience. This approach capitalises on the “Kuleshov effect” whereby shots are interpreted differently depending on the context (Piacenza et al., 2011).

The motivation for the work we describe in this paper comes from our industrial partner² who is keen to exploit the visual quality of video by using filmic content in an IS System set in the horror genre. The specific problem we have addressed is how to select appropriate content for the presentation of a narrative *fabula* in a way that is (i) tailored to the preferences of different types of user i.e. a personalised horror film variant; and (ii) fits within a fixed-length time format (e.g. full length movie, window between adverts, and so on). In this work we focus on two types of viewers, *gore* watchers and *thrill* watchers, as identified by Johnston (1995) and attempt to best-fit the presentation within an input duration.

Whilst AI planning is well suited to narrative *fabula* generation we observed that the task of discourse generation using pre-recorded filmic content is more naturally modelled as a Constraint Programming (CP) problem. The rationale for this is because the focus of

¹Shots are sequences of continuous still images (frames) as filmed through a single, uninterrupted camera take (Cotsaces, Nikolaidis, and Pitas, 2006).

²Portal Entertainment

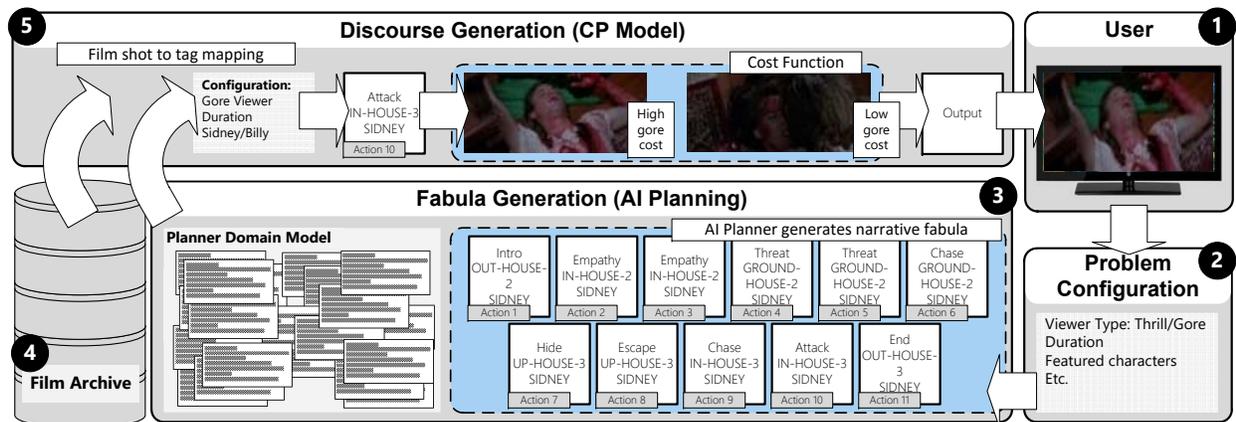


Figure 1: System Overview. For a given user (1) a bespoke problem is configured (2) and an output narrative Fabula is generated using an AI Planning approach (3). The Fabula and User Settings are communicated to the Discourse Generator (5) which uses a CP approach to select appropriate content from the film archive (4) to generate an output film that is optimised towards the user’s preferences. This is then presented to the user. For more detail see text.

discourse presentation generation is on expressing relevant constraints and on finding a feasible solution that seeks to maximise certain properties, rather than the sequence of state changes required for the fabula. A further advantage to this decoupling of the discourse generation task from the narrative generator is because it allows the use of non-planning based generators that do not handle constraints (e.g. approaches based on Linear Logic (Bossler, Cavazza, and Champagnat, 2010; Martens et al., 2013, 2014)).

In the paper we show how this problem can be formulated as a CP model and overview our approach. We illustrate it with examples taking from our prototype IS which uses planning for narrative fabula generation and our CP solution to select appropriate filmic content to form the output presentation. For this prototype, we have built a film archive populated with tagged content taken from the film *Scream* (Craven, 1996) as illustration. We also present the results of an initial user evaluation which demonstrate the potential of the approach.

Related Work

There have been a number of attempts at interactive cinema over the years, from Činčera’s (Kinoautomat, 1967) to more examples of interactive films such as (Last Call, 2010) and (Accidental Lovers, 2008). There is also ongoing research interest in the area of multimedia systems research with a number of prototype systems aimed at storyfication: assisting in the construction of narratives around existing collections of video content. For example, Shen, Lieberman, and Davenport (2009) introduced a video editing system that helped authors compose a sequence of scenes that tell a story, by selecting from a corpus of annotated clips.

Zsombori et al. (2008) presented an approach to interactive television which could adapt during delivery to the preferences of viewers which relied on a number of authoring and delivery tools to configure the system. A limitation of all these approaches is that they rely on manual coupling of content with pre-defined plot branches, which restricts them to simple rearrangements of an input narrative and associated branching structure. In order to support the generative potential afforded by plan-based planning systems, it is necessary to be able to exploit shot polysemy which allows it to be reused in a much wider range of valid semantic contexts.

An approach which comes much closer to this, and is more closely related to our work, is the Video Based Storytelling system of Piacenza et al. (2011), which uses an archive of video content tagged with a small number of semantic categories such as characters, location, and mood. Their system features plan based narrative fabula generation and a semantic mapping from actions to appropriate semantic tag categories in attempt to find suitable content to stage given narrative actions. The planner sends requests for video content for all actions in the fabula and if none can be found, the story can be replanned incrementally until appropriate content is found. In contrast, within our approach the fabula is viewed more flexibly as providing an outline of the plot and the CP solver finds the best content with which to present that outline in order to satisfy viewer preferences.

An important aspect of our approach is tailoring output presentations towards different viewer types. For the purposes of this initial work and motivated by our industrial partners we focus on two archetypes: thrill

watchers and gore watchers. These types were first identified in a study by Johnston (1995), who looked at the motivations of viewers watching horror (for further detail see the next section) and are still useful classifications in film studies and criticism (for example, (King and Hourani, 2007) and (Hess, 2010)).

System Overview

Our CP model is embedded in a prototype system that outputs horror film variants and here we briefly describe the different system components, as shown in Figure 1.

Within the system, a bespoke problem instance is created for a given user ①② on the basis of such things as the user type, feature characters, the type of ending and so on. This is communicated to the Fabula Generation Module ③ which is implemented using an AI plan-based approach although we note that other approaches which output fabula as a sequence of narrative actions could also be used here. Within our approach the fabula is used to give an outline order of story content, including details of characters, locations, their roles and the ending of the narrative (e.g. whether the protagonist lives or dies). However this outline ordering is flexible and not all of these elements of story content are required. This is necessary in order to ensure constraints can be satisfied at run time, as we discuss later (see section Scene Generation as a Constraint Problem below).

A shared vocabulary is used to communicate this outline fabula to the Discourse Generation module: the narrative action names and parameters are the same tags that are used to tag content in the film archive.

The tagged film content is stored in an archive which is accessed by the CP model, as shown in part ④. For the purposes of our demonstrator system we used the film *Scream* (Craven, 1996) and broke this down to the shot level using the Video Shot and Scene Segmentation tool (Multimedia Knowledge and Social Media Analytics Laboratory, 2014). This content was tagged according to the identified categories as part of the domain model construction for the system (in total this represented about 5% of the original film).

The Discourse Generation module ⑤ is responsible for the generation of the output sequence, satisfying the constraints for the input viewer type and duration. This sequence of shots is then presented as a film, using a VLC media player (VideoLAN, 2016), with sound overlaid using royalty-free and horror themed music (MacLeod, 2015).

Horror Film Content Tagging

We have formulated a vocabulary which we use to mark up film content for use in our prototype system and which is used as a medium for communication between the Fabula and Discourse Generator and also a set of Discourse tags which are used internally within the Discourse Generator as discussed below.

Shared Semantic Vocabulary

These are semantic categories which can be used to describe the sorts of things that a shot can convey when presented to the viewer. We refer to these as *shot types*. They were formulated on the basis of advice from screenwriting manuals such as (McKee, 1997) along with input from our industrial partner. We use these shot types to categorise content that can be used as follows:

- Empathy: content that can be used to create feelings of empathy from the viewer for the protagonist. If it conveys “positive aspect ... the audience can recognise as like themselves and having recognised that they then identify with the protagonist” (McKee, 1997). For example, a teenager, home alone, making popcorn or a teenager coming home waiting on their own perhaps for a parent to return.
- Threaten: content that can show a character in danger or at risk (with different intensity levels).
- Chase: show a character being chased by another (agent/patient and differing intensity levels).
- Hide: show a character hiding from another (agent and patient).
- Kill: show one character killing another (agent and patient)
- Attack: show one character attack another (agent and patient, and differing intensity levels).
- Escape: one character escaping from another (agent and patient, and differing intensity levels).
- Introduction: show some introduction to the film (e.g. the location, time of day, and so on)
- End: show the end of the narrative (e.g. whether the protagonist is alive or dead at the end of the film).

These shot types are used to tag content in the film archive. They are also the action names used in the AI planning model that is used by the Fabula Generator and which are communicated from the Fabula Generator to the Discourse Generator (see Figure 1).

The following are also used, in combination with the shot types listed above, to tag content and to ensure the selection of appropriate content and continuity:

- Character information: (i) A unique name, used as an identifier for the character; and (ii) their status, whether they are alive or dead.
- Location: a unique name to identify where the shot is located and to ensure continuity.

Discourse Generation tags

A small number of tags are also used to mark up the film content to be used locally within the Discourse Generation module as they relate solely to the presentation of the narrative and the preferences of different types of user. In particular they are used in the cost function within the CP model as part of the shot selection problem. These tags are:

Example Shot	Content Tagging
	① Empathy In-house1 Casey Billy ② Duration: 8 seconds Gore rating: 0 Thrill rating: 0
	① Threaten In-house1 Casey Billy 1 ② Duration: 8 seconds Gore rating: 0 Thrill rating: 1
	① End Out-house2 Sidney alive ② Duration: 6 seconds Gore rating: 0 Thrill rating: 0

Figure 2: Example Shots shown with their associated tags: ① the Semantic Vocabulary that is shared between the Fabula and Discourse Generators; and ② tags used solely for Discourse (see text for detail).

- Duration of the shot (measured in seconds).
- Gore Rating: a numeric rating of how gory a shot is which is based on subjective aspects such as the amount of graphic violence, blood and so on.
- Thrill Rating: a numeric rating of the shots appeal to the thrill watcher based on aspects such as tension, empathy, identification with the protagonist etc.

As illustration, Figure 2 shows some shots from the *Scream* film archive that was used in our experiments and their associated tags (both shared semantic vocabulary and discourse tags).

Scene Generation as a Constraint Problem

We consider the problem of generating *filmic output* (films for short) for the horror genre where a film is a sequence of shots, which themselves are sequences of continuous still images (frames) as filmed through a single, uninterrupted camera take (Cotsaces, Nikolaidis, and Pitas, 2006). We want this sequence of shots to follow an outline order of story content generated by the Fabula Generation Module. For a given *viewer profile* (viewer for short), we also want this sequence to satisfy the user’s preferences for the horror genre as closely as possible. Finally, we also want the duration of the sequence of shots to be within some specified target duration. Within this the different elements are:

Shots We have a set of *shots* S . Following (Cotsaces, Nikolaidis, and Pitas, 2006), a shot is defined as a sequence of continuous still images (frames) as filmed through a single, uninterrupted camera take. Shots are of varying duration (number of seconds) with

$$\forall s \in S. \text{duration}(s) \in \mathbb{N}$$

We also have a special shot, denoted by `no_shot`, that represents the empty shot. We assume that this shot is not an element of S and we define its duration to be zero:

$$\text{duration}(\text{no_shot}) = 0$$

Shot Types We classify shots using a subset of shot types, based on the semantic vocabulary described in the previous section. For all shots s , we have

$$\text{shot_type}(s) \subseteq \{ \text{introduction}, \text{empathy}, \text{threaten}, \text{chase}, \text{attack}, \text{kill}, \text{escape}, \text{hide}, \text{end} \}$$

In cases where $|\text{shot_types}(s)| > 1$ only one element of $\text{shot_types}(s)$ is required to fulfil the requirement of the story. We also define for each shot an intensity level associated with some shot types:

$$\begin{aligned} \text{attack_level}(s) &\in \mathbb{N} \\ \text{kill_level}(s) &\in \mathbb{N} \\ \text{empathy_level}(s) &\in \mathbb{N} \\ \text{escape_level}(s) &\in \mathbb{N} \\ \text{hide_level}(s) &\in \mathbb{N} \end{aligned}$$

As an example, the higher the value of $\text{attack_level}(s)$, the higher the intensity of the attack is.

Viewer, Viewer Profile and Variation Based on the analysis of Johnston (1995), we focus on the following types of graphic horror viewers: thrill watchers and gore watchers. As outlined below, these two types of viewers have different characteristics which influence selection of content that best fits their preferences.

Thrill Watchers: these viewers have a disposition towards (a preference for) content that establishes empathy with the protagonist, such as shots s where $\text{shot_type}(s) = \text{empathy}$. This type of viewer is also argued to be high sensation seeking which means they are motivated by the suspense of the film and have more identification with the victims than other types of viewer such as gore. An important element in creating suspense is temporal distension and a slow progression of events has been noted as a common feature in suspense scenes (de Wied, 1995). Hence for this type of watcher we assign a higher preference for build up of threat and chase than for gore watchers. It has also been noted that the intensity of suspense for a viewer increases as they feel disposition towards protagonists or outcomes (Cheong and Young, 2008), so it is not surprising that viewers with higher empathy and sensation seeking prefer suspenseful films.

We tag shots with a *thrill rating* which measures how thrilling the shot is. This category of viewer is interested in high thrill ratings. We use *thrill_rating(s)* to denote the thrill level of a shot s .

Gore Watchers: in contrast to the thrill watchers, these viewers are characterised by low empathy, high sensation seeking, and a strong identification with the killer. In terms of content, we conclude that the preference would be for graphic violence and shots which feature this type of content. In terms of the shot types introduced earlier this includes Attack and Kill. In addition, shots are tagged with a *gore rating* which denotes the amount of graphic violence which also denotes what shots would be more preferable to this category of viewer. We use *gore_rating(s)* to denote the gore level of a shot s .

Thus for our prototype system we used the following content preference for different type of viewer:

Viewer type	Preferred content
Thrill	Empathy, Threaten, Escape, Hide, Chase
Gore	Kill, Attack

Viewer Relevance We assign a relevance value (i.e. cost function), $relevance(s, v) \in \mathbb{N}$, to each shot $s \in S \cup \{\text{no_shot}\}$ and viewer type v . The higher this value is, the more relevant the shot is. Following the discussion above, we define the relevance of a shot for gore viewers as follows:

$$\begin{aligned}
relevance(s, \text{gore}) = & \\
& \text{gore_rating}(s) \times (1 + \text{attack_level}(s) + \\
& \quad \text{kill_level}(s)) \\
& - \text{thrill_rating}(s) \times (1 + \text{empathy_level}(s) + \\
& \quad \text{escape_level}(s) + \text{hide_level}(s) + \\
& \quad \text{threaten_level}(s) + \text{chase_level}(s))
\end{aligned}$$

For thrill viewers, relevance is defined as:

$$relevance(s, \text{thrill}) = -relevance(s, \text{gore})$$

The shot *no_shot* has the same (neutral) relevance value for both types of viewer:

$$relevance(\text{no_shot}, _) = 0$$

Narrative Transition The planner outputs a baseline story, consisting of a sequence of m shot types. We represent this sequence as BL and we index it using square brackets (e.g. $BL[0] = \text{introduction}$ or $BL[m-1] = \text{end}$). We want to generate sequences of shots that respect this ordering. We have represented shot ordering within each plan step in this model, which allows the model to be extended with ordering constraints over these individual shots.

Constraint Programming Model

The goal is to find a scene that satisfies the constraints shown below.

Variables

$s_{ij} \in S \cup \{\text{no_shot}\}$ This is the j^{th} shot of the i^{th} category (i.e. the i^{th} element of BL), where $0 \leq i < m$ and $0 \leq j < |S|$.

Constraints Our objective is to maximise the relevance of scenes for a given viewer type v , where v is either *gore* or *thrill*.

$$\max \sum_{\substack{0 \leq i < m \\ 0 \leq j < |S|}} relevance(s_{ij}, v) \quad (1)$$

such that

$$\forall i, j. s_{ij} \neq \text{no_shot} \Rightarrow (\text{shot_type}(s_{ij}) = BL[i]) \quad (2)$$

$$\forall i \forall j, j'. s_{ij} \neq \text{no_shot} \wedge s_{ij'} \neq \text{no_shot} \wedge j \neq j' \Rightarrow s_{ij} \neq s_{ij'} \quad (3)$$

$$\forall i, j. (j < m-1 \wedge s_{ij} = \text{no_shot}) \Rightarrow (s_{i(j+1)} = \text{no_shot}) \quad (4)$$

$$\min D \leq \sum_{\substack{0 \leq i < m \\ 0 \leq j < |S|}} duration(s_{ij}) \leq \max D \quad (5)$$

Constraint (2) guarantees that the baseline story ordering is respected. Constraint (3) excludes scenes with repeated shots (the only shot that can be repeated is the special shot *no_shot*). Constraint (4) guarantees if at any point a *no_shot* is selected for the sequence of shots of a certain type, then all the subsequent shots for the same type must be *no_shot*. Finally, constraint (5) ensures that the total duration of the scenes is within the range of acceptable durations.

Example 1 Let us assume that the planner creates the following baseline story

$$BL = [\text{introduction}, \text{chase}, \text{escape}, \text{end}]$$

and that we have a total of 10 shots

$$S = \{\text{s0}, \text{s1}, \text{s2}, \text{s3}, \text{s4}, \text{s5}, \text{s6}, \text{s7}, \text{s8}, \text{s9}\}$$

with the following properties:

- Shots tagged with *introduction*: s0 and s2
- Shots tagged with *chase*: s3, s4, s5, and s8
- Shots tagged with *escape*: s6 and s9
- Shots tagged with *end*: s1, s7, and s8

Then, a possible solution is:

$$\begin{aligned}
\text{solution} = & [\text{s0}, \text{s2}, 8*\text{no_shot}, \\
& \text{s3}, \text{s4}, \text{s5}, \text{s8}, 6*\text{no_shot}, \\
& \text{s6}, \text{s9}, 8*\text{no_shot}, \\
& \text{s1}, \text{s7}, \text{s8}, 7*\text{no_shot}]
\end{aligned}$$

where $k*\text{no_shot}$ denotes a sequence consisting of k occurrences of *no_shot*.

A different way of visualising the solution is in terms of the variables used in the model:

- $s_{00}=s0, s_{01}=s2$, and $\forall_{2 \leq j < |S|} \cdot s_{0j}=\text{no_shot}$
- $s_{10}=s3, s_{11}=s4, s_{12}=s5, s_{13}=s8$, and $\forall_{4 \leq j < |S|} \cdot s_{1j}=\text{no_shot}$
- $s_{20}=s6, s_{21}=s9$, and $\forall_{2 \leq j < |S|} \cdot s_{2j}=\text{no_shot}$
- $s_{30}=s1, s_{31}=s7, s_{32}=s8$, and $\forall_{3 \leq j < |S|} \cdot s_{3j}=\text{no_shot}$

Gecode Implementation

Our model was implemented using an open source generic constraints solver called Gecode (Schulte, Tack, and Lagerkvist, 2010). We used a matrix representation m_{ij} where the columns i represented a film shot while the rows j represented the individual properties of each shot (shot type, shot duration, etc.). Sequences of values in i were constrained using table constraints so that they can only match the data given by our tags. Additionally, our model contained a single variable c which holds the output of result of Constraint (1) according to the state of the matrix m .

A branch and bound search strategy was used to solve the constraint satisfaction problem, optimising over the variable c .

We tested our implementation using the fabula generated in figure 3, altering the settings as shown in figure 1 to output solutions for both thrill and gore viewers. The average solving time over 100 attempts was 109ms using both plans. This was run on a 2.26-GHz Intel Xeon dual-core.

Evaluation

For the evaluation we used our implemented prototype system with components as shown in Figure 1.

For the Fabula generator we built a small planning model, represented using PDDL with action names corresponding to the shot types which characterise the horror genre. As an example, consider the following action:

```
(:action threaten
:parameters (?c2 ?c1 - character ?l - location)
:precondition (and
  (at ?c1 ?l) (alive ?c1) (antagonist ?c2) (alive ?c2)
:effect (and
  (increase (peril ?c1) 2))))
```

which increments the perceived *peril* of the situation of the character being threatened i.e. the patient of the action. In this example the character doing the threatening (the agent) is the antagonist. Narrative Fabula goals were specified in terms of functions such as *peril* or predicates such as *alive* or *dead*. Note that the set of action names used in the Planning Model are the same as the *shot-types* introduced in the CP model and used to tag filmic content: thus providing a shared semantic vocabulary for communication between the plan-based Fabula Generator and the Discourse Generator.

For the purposes of this evaluation output narratives were generated using METRIC-FF (Hoffmann and Nebel,

2001). The Film Archive was populated with content taken from the film *Scream* (Craven, 1996) which was automatically segmented and then tagged as described earlier (see the section Horror Film Content Tagging). For our initial experiments, as the focus was on content selection, the original film dialogue and soundtrack were omitted and output filmic sequences were overlaid with appropriate horror themed music.

Example: Generating Filmic Output

Our expectation is that the use of a cost function based on viewer preference to different types of content will enable the selection and sequencing of appropriate content. In this section we present, through analysis of two system generated filmic sequences, results that support our expectations for the CP approach. The system generated filmic sequences are shown in storyboard format in Figure 3. Across the centre of the figure is a single output Narrative Fabula, consisting of a sequence of 13 narrative actions. Along the top is the output discourse presentation generated for the Thrill Viewer type ① and at the bottom the filmic sequence for the Gore Viewer type ②. For each filmic sequence the cost functions, in terms of “value” for gore and thrill are plotted. We make the following observations:

- Different shot types (i.e. actions from the Fabula) have been selected, subject to the constraints, which are more appropriate for the different types of viewer. For example: the thrill presentation features more *Escape*, *Hide* and *Chase* shot types: the gore presentation features more *Attack*.
- Visual inspection of the storyboards shows that the system has selected very different looking film content for the different types of viewer. For example, the gore sequence features more graphic violence.

User Evaluation

We conducted a small user evaluation to evaluate how users responded to the output horror films generated by our CP approach. A total of 33 adults participated in the evaluation. Each participant was asked to view a total of 4 films: 2 generated fabulas with a gore and thrill oriented discourse presentation for each³. All films and questions were delivered via an online questionnaire.

1) Content Variation We were interested to explore whether participants could detect any noticeable difference in the level of gore or thrill in the system generated films. Two Fabula’s were generated: Fabula-1 where the protagonist survives (ie a goal of (*alive girl1*)); and Fabula-2 where the protagonist dies (ie with the goal (*dead girl2*)). For each Fabula the Discourse Generator output a Gore variant and a Thrill variant. Thus participants viewed the following videos:

³For reference the output films are available to view: (Please be advised that this film received a rating for audiences aged 15 and up as it contains blood and violence.) <https://goo.gl/LdGVlH>

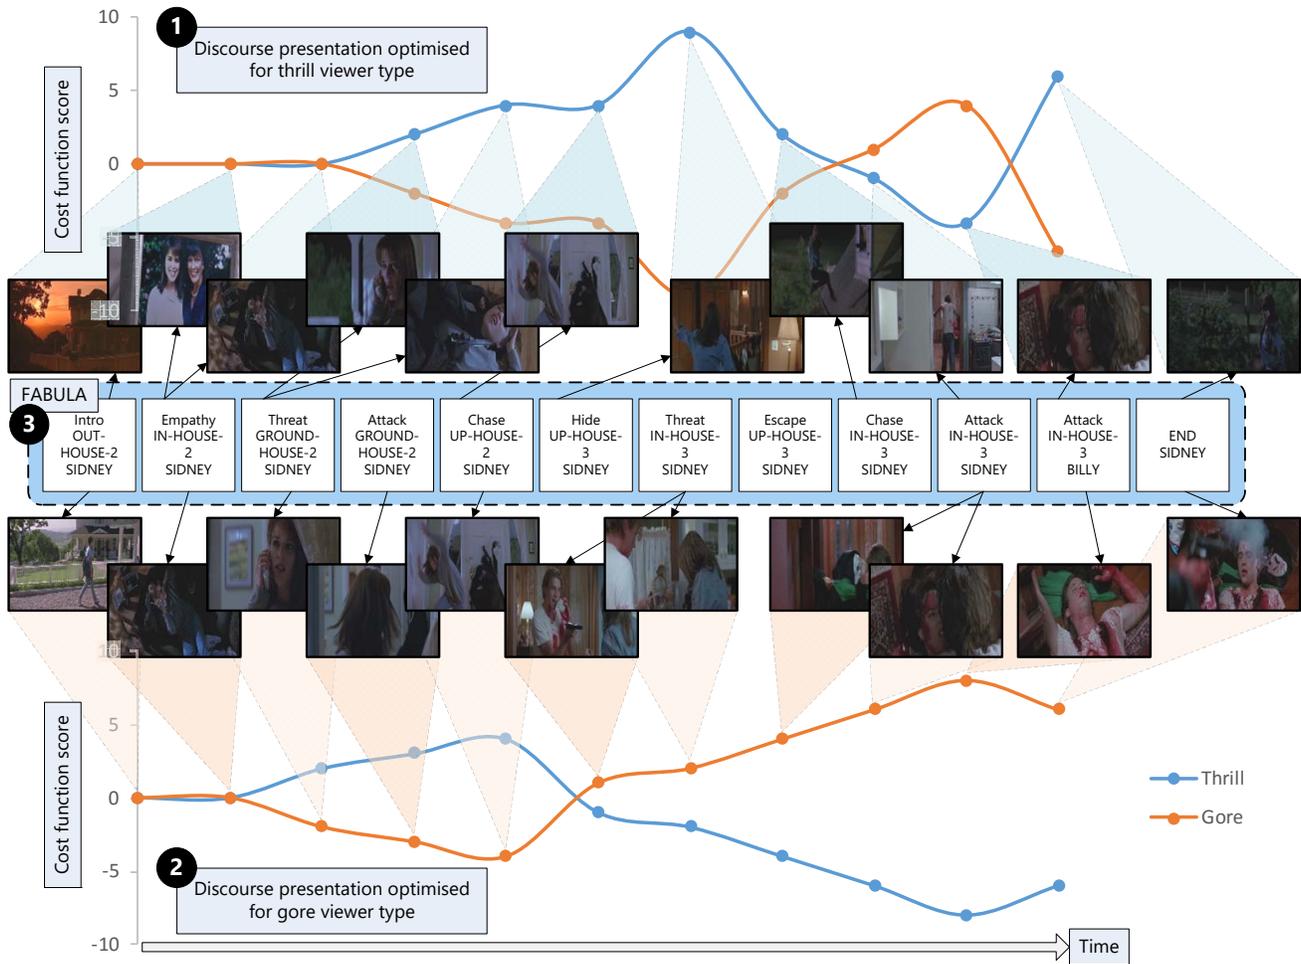


Figure 3: Illustrations of Output Filmic Variants for viewer types: Thrill ① and Gore ②. A single Fabula which has been generated by the system is shown across the centre of the figure ③. From this the system has selected: different narrative actions from the baseline narrative to present depending on Thrill or Gore viewer type; and inspection of the storyboards shows very different content selected for the discourse for the same baseline narrative.

- Video 1: Fabula-1 Thrill
- Video 2: Fabula-1 Gore
- Video 3: Fabula-2 Thrill
- Video 4: Fabula-2 Gore

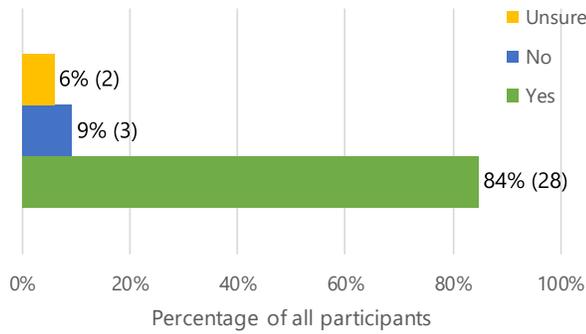
and then asked the following questions:

1. “Q: Did one video contain more graphic violence?” with options *yes, no, don't know*. If *yes* was selected they were prompted to select between the 2 videos.
2. “Q: Did one video display more suspense?” with options *yes, no, don't know*. If *yes* was selected they were prompted to select between the 2 videos.

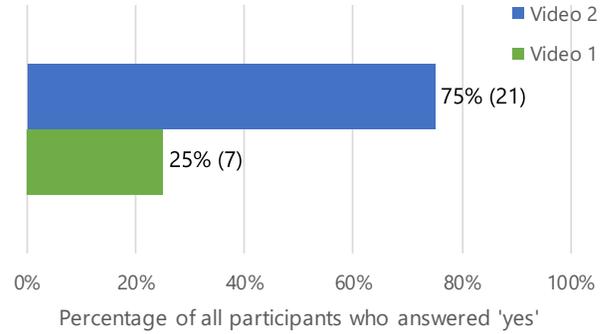
The results of these questions are shown in Figures 4, 5 (for videos 1 and 2) and 6, 7 (for videos 3 and 4). For all of the videos generated by the system we observe that in all cases the users responses were in line with

our expectations: we expected them to observe more graphic violence (suspense) in those sequences generated for the gore (thrill) viewer type and in all cases this was shown. However we also observe that this was more convincing for the gore sequences than for thrill. For example, for videos 3 and 4, 31 out of the 33 participants correctly selected video 4 as displaying more graphic violence, whereas only 15 out of the 33 felt that there was much difference in the amount of suspense between the 2 videos (although 12 of those correctly identified video 3 as being more suspenseful).

There are a number of factors which might explain this: in particular one might be the relatively small amount of content which we had marked up and which was available for use in the system, especially for the thrill viewer type. Another factor is the relatively short duration of output which we used for discourse gener-

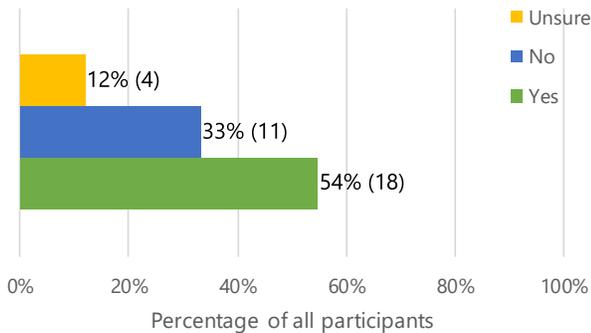


(a) User responses to the question: “*Did one video display more GRAPHIC VIOLENCE than the other?*”

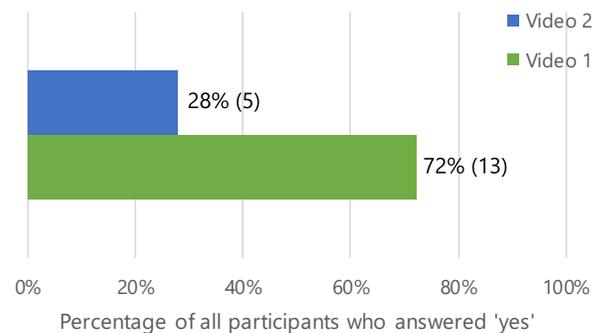


(b) Users who answered “yes” to (a) were also asked “*Which video displayed more GRAPHIC VIOLENCE?*”

Figure 4: User Results for Videos 1, 2: (a) the majority of users thought one video contained more graphic violence; and (b) of those, the majority selected video 2 as displaying more graphic violence. As video 2 was of type *Gore* this result supports our expectation.



(a) User responses to the question: “*Did one video display more SUSPENSE than the other?*”



(b) For users who answered “yes” to (a) they were then asked “*Which video displayed more SUSPENSE?*”

Figure 5: User Results for Videos 1 and 2: (a) shows the majority of users thought one of the videos contained more suspense; and (b) shows the majority of those judged video 1 to be more suspenseful. Video 1 was type *Thrill* so the result is consistent with our expectation.

ation for this initial study (on average this was 72 seconds) and it may be that longer duration is required to build up suspense/thrill whereas graphic violence is more easily observed.

2) Narrative Understanding We were also keen to see whether viewers could comprehend aspects of the narrative that was presented to them, in particular the outcome of the story for the protagonist (feature character). As a measure of this participants were asked the following question for both videos 1/2 and 3/4:

“*Q: What was the ending for the character [name]?*”

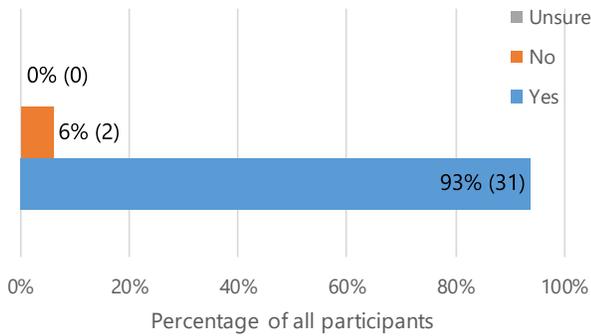
(and where the [name] of the feature character in video 1/2 was Casey and for video 3/4 was Sidney). We report that 100% of participants were able to answer correctly about whether the character was alive or dead

(irrespective of whether they were familiar with the original movie). This result was consistent with our expectation.

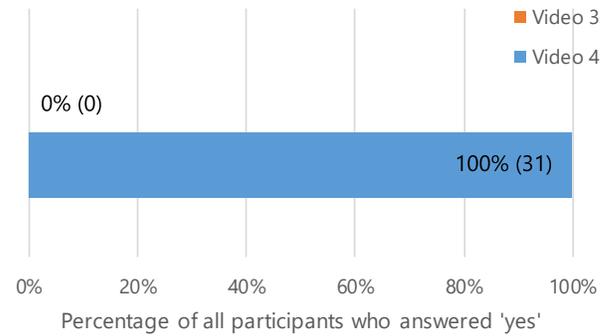
Conclusion

In this paper we have set out a bipartite approach to narrative variant generation which utilises AI planning to generate the fabula and Constraint Programming to generate the discourse. Our CP model relies on appropriately tagged content to maximise specific attributes based on the type of viewer it is catering for. We developed a prototype system and, using material from an existing horror film, have shown that this approach is effective based on a small user study.

In future work we intend to extend the size of the film archive in order to increase the space of different filmic sequences that can be generated e.g. in order to carry

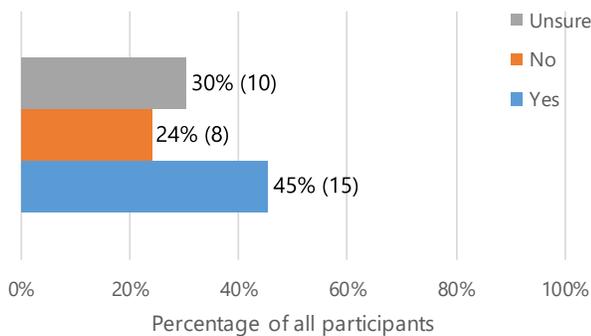


(a) User responses to the question: “Did one video display more GRAPHIC VIOLENCE than the other?”

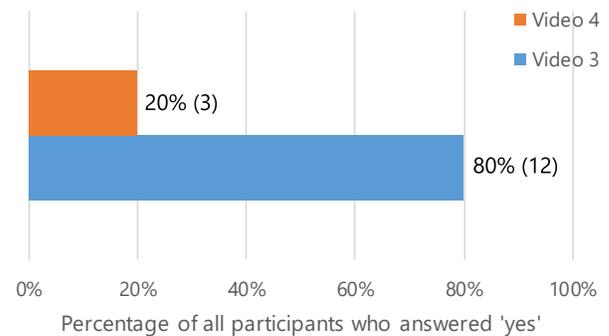


(b) Users who answered “yes” to (a) were also asked “Which video displayed more GRAPHIC VIOLENCE?”

Figure 6: User Results for Videos 3 and 4: (a) the majority of users thought one video contained more graphic violence; and (b) of those, the majority selected video 4 as displaying more graphic violence. As video 4 was of type *Gore* this result supports our expectation.



(a) User responses to the question: “Did one video display more SUSPENSE than the other?”



(b) For users who answered “yes” to (a) they were then asked “Which video displayed more SUSPENSE?”

Figure 7: User Results for Videos 3 and 4: (a) shows most users thought one of the videos contained more suspense; and (b) shows the majority of those judged video 3 to be more suspenseful. Video 3 was type *Thrill* so the result is consistent with our expectation.

out further user evaluation. We also intend to increase the CP model to include a wider range of features such as lighting, camera angle and so on.

References

- Accidental Lovers. 2008. <http://mlab.uiah.fi/~isaarine/accidental lovers.html>.
- Aylett, R.; Dias, J.; and Paiva, A. 2006. An Affectively Driven Planner for Synthetic Characters. In *Proc. 16th Int. Conference on Automated Planning and Scheduling (ICAPS)*.
- Bosser, A.; Cavazza, M.; and Champagnat, R. 2010. Linear Logic for Non-Linear Storytelling. In *Proc. 19th European Conference on Artificial Intelligence (ECAI)*.
- Cheong, Y.-G., and Young, R. M. 2008. Narrative generation for suspense: Modeling and evaluation. In *Joint International Conference on Interactive Digital Storytelling*, 144–155. Springer.
- Cotsaces, C.; Nikolaidis, N.; and Pitas, I. 2006. Video shot boundary detection and condensed representation: a review. *IEEE signal processing magazine* 23(2):28–37.
- Craven, W. 1996. *Scream*. Dimension Films.
- de Wied, M. 1995. The role of temporal expectancies in the production of film suspense. *Poetics* 23(1-2):107–123.
- Haslum, P. 2012. Narrative Planning: Compilations to Classical Planning. *Journal of AI Research* 44:383–395.
- Hess, J. P. 2010. The psychology of scary movies. *Filmmaker IQ*. Located at:

- <http://filmmakeriq.com/lessons/the-psychology-of-scary-movies/>.
- Hoffmann, J., and Nebel, B. 2001. The FF Planning System: Fast Plan Generation through Heuristic Search. *Journal of AI Research* 14:253–302.
- Johnston, D. D. 1995. Adolescents’ motivations for viewing graphic horror. *Human Communication Research* 21(4):522–552.
- Kartal, B.; Koenig, J.; and Guy, S. J. 2014. User-Driven Narrative Variation in Large Story Domains using Monte Carlo Tree Search. In *Proc. of the 13th Int. Conference on Autonomous Agents and MultiAgent Systems (AAMAS)*.
- King, C. M., and Hourani, N. 2007. Don’t Tease Me: Effects of Ending Type on Horror Film Enjoyment. *Media Psychology* 9(3):473–492.
- Kinoautomat. 1967. <http://www.kinoautomat.cz/index.htm?lang=gbr>.
- Last Call. 2010. <https://www.youtube.com/watch?v=qe9CiKnrS1w>.
- MacLeod, K. 2015. This house. <http://incompetech.com/music/royalty-free/?genre=horror>. Accessed: 2017-03-06.
- Martens, C.; Bossler, A.-G.; Ferreira, J. F.; and Cavazza, M. 2013. Linear logic programming for narrative generation. In *International Conference on Logic Programming and Nonmonotonic Reasoning*, 427–432. Springer Berlin Heidelberg.
- Martens, C.; Ferreira, J. F.; Bossler, A.-G.; and Cavazza, M. 2014. Generative story worlds as linear logic programs. In *Seventh Intelligent Narrative Technologies Workshop*.
- Mateas, M., and Stern, A. 2005. Structuring Content in the Façade Interactive Drama Architecture. In *Proc. of 1st Conference on Artificial Intelligence and Interactive Digital Entertainment (AIIDE)*.
- McKee, R. 1997. *Story: substance, structure, style, and the principles of screenwriting*. NY: ReganBooks.
- Multimedia Knowledge and Social Media Analytics Laboratory. 2014. Video shot and scene segmentation tool. <http://mklab.itl.gr/project/video-shot-segm>. Accessed: 2017-03-06.
- Orkin, J., and Roy, D. K. 2012. Understanding Speech in Interactive Narratives with Crowdsourced Data. In *Proc. of 8th Conference on Artificial Intelligence and Interactive Digital Entertainment (AIIDE)*.
- Piacenza, A.; Guerrini, F.; Adami, N.; Leonardi, R.; Teutenberg, J.; Porteous, J.; and Cavazza, M. 2011. Changing video arrangement for constructing alternative stories. In *Proc. of the 19th ACM International Conference on Multimedia*.
- Porteous, J.; Cavazza, M.; and Charles, F. 2010a. Narrative Generation through Characters’ Point of View. In *Proc. of 9th Int. Conference on Autonomous Agents and MultiAgent Systems (AAMAS)*.
- Porteous, J.; Cavazza, M.; and Charles, F. 2010b. Applying Planning to Interactive Storytelling: Narrative Control using State Constraints. *ACM Transactions on Intelligent Systems and Technology (ACM TIST)* 1(2):1–21.
- Riedl, M. O., and Young, R. M. 2010. Narrative Planning: Balancing Plot and Character. *Journal of Artificial Intelligence Research* 39:217–267.
- Schulte, C.; Tack, G.; and Lagerkvist, M. Z. 2010. Modeling and programming with gecode. *Schulte, Christian and Tack, Guido and Lagerkvist, Mikael* (2015).
- Shen, E. Y.-T.; Lieberman, H.; and Davenport, G. 2009. What’s Next?: Emergent Storytelling from Video Collection. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, 809–818. ACM.
- VideoLAN. 2016. <http://www.videolan.org/vlc/index.en-GB.html>.
- Young, R. M. 2007. Story and Discourse: A Bipartite Model of Narrative Generation in Virtual Worlds. *Interaction Studies: Social Behaviour and Communication in Biological and Artificial Systems* 8:177–208.
- Zsombori, V.; Ursu, M.; Wyver, J.; Kegel, I.; and Williams, D. 2008. *ShapeShifting documentary: A golden age*. Lecture Notes in Computer Science. Springer. 40–50.