Generating Personalized Narrative Experiences in Interactive Storytelling through Automated Planning

Simone Agostinelli
Sapienza Università di Roma
Rome Italy
agostinelli@diag.uniroma1.it

Federica Battaglini
Sapienza Università di Roma
Rome Italy
federica.b90@gmail.com

Tiziana Catarci
Sapienza Università di Roma
Rome Italy
catarci@diag.uniroma1.it

Federica Dal Falco
Sapienza Università di Roma
Rome Italy
federica.dalfalco@uniroma1.it

Andrea Marrella
Sapienza Università di Roma
Rome Italy
marrella@diag.uniroma1.it

ABSTRACT

Interactive storytelling is a form of narrative in which the storyline is not predetermined. In this paper, we discuss how the use of automated planning techniques in Artificial Intelligence can be employed to generate personalized narrative experiences in interactive storytelling. We show the feasibility of our approach through a mobile application for cultural heritage based on mini games, whose order of presentation is dynamically determined to increase the user engagement in museum-like spaces.

KEYWORDS

Interactive Storytelling, Automated Planning, Artificial Intelligence, Mobile Application, Cultural Heritage

ACM Reference format:

Simone Agostinelli, Federica Battaglini, Tiziana Catarci, Federica Dal Falco and Andrea Marrella. 2019. Generating Personalized Narrative Experiences in Interactive Storytelling through Automated Planning. In CHITALY'19, Sept. 23–25, 2019, Padova, Italy. 5 pages.

1 INTRODUCTION AND BACKGROUND

For many years, it was a common belief in Human-Computer Interaction (HCI) that storytelling and interactivity were diametrically opposed, meaning that one can have story or one can have interactivity, but not both simultaneously [11]. The purpose of *interactive storytelling* is to contradict this belief, providing a bridge between these two concepts. Interactive storytelling is a form of narrative in which the storyline is not predetermined. The author of the story creates the setting (i.e., the world), characters, and situations that the story must address, in a way that the user (called *player* in this context) experiences a unique story based on her/his interactions with the story world [2, 6]. A key ingredient of interactive storytelling is that the player is enabled to make decisions that may affect the direction and the outcome of the narrative experience, ensuring that new narrative events unfold comprehensibly.

Approaches for interactive storytelling have been recently employed in many mainstream video games and have found application as a technique to improve user engagement in different activities, including visits to cultural sites and cultural tourism in general [7, 10]. According to [9], they can be classified in those ones that provide a *strong autonomy* or a *strong story*. On the one hand, the strong autonomy approach advocates that interactive narratives are generated procedurally by simulating a virtual environment populated by autonomous agents that play the roles of characters [1] (e.g., open-ended games like "The Sims" or "Minecraft"). On the other hand, strong story approaches are those in which some form of *global story control* is implemented. Often story control is achieved in a centrally coordinated fashion through the use of a decision-making agent, which attempts to coerce the experience of the player to conform to a pre-existing story structure by directing the story world characters [13] (e.g., games with strong narrative focus such as "BioShock").

In this paper, we present an approach to modeling *interactive strong stories* as solutions of automated planning problems [4] in Artificial Intelligence (AI). Planning systems are problem-solving algorithms that operate on explicit representations of states and actions that are expressed in compact form through the standardized Planning Domain Definition Language (PDDL) [5]. Automated planning addresses the problem of generating actions sequences (i.e., *plans*) fulfilling a desired goal state from an initial state of the world. Two ingredients are required to generate a plan: *(i)* a *planning domain* that describes how *actions* work in a domain/world of interest. Any action is characterized by *preconditions* that state under which world conditions the action can be executed, and effects on the state of the world; *(ii)* a *planning problem* that indicates what the initial state of the world is and the goal state to be achieved. Similarly, however simplistic it might appear, the common denominator of interactive (strong) story formalization is to consider a story as a sequence of actions related through some form of causality, whose aim is to generate a credible narration with a well-defined target.

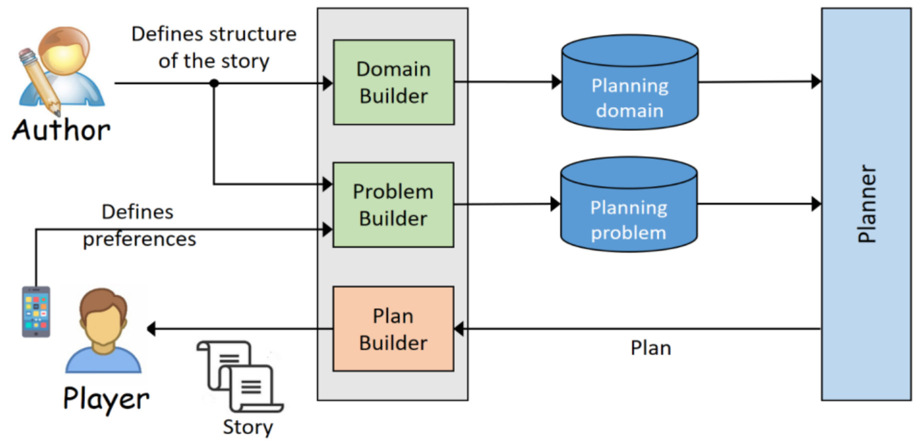

Figure 1: Overview of our approach to interactive storytelling.

Thus, by representing the underlying dynamics of the world in which a story unfolds as a planning domain and the characteristics of the specific target story instance that we intend to generate as a planning problem, then we can use an off-the-shelf planner to generate a story according to the given description. One motivation for this approach is that a wide range of story alternatives can be handled by a concise representation in PDDL, and these stories can be adapted and tweaked both at design-time and run-time in order to satisfy requirements that arise during the execution, e.g., in a video game context.

It is worth to mention that the use of automated planning for interactive storytelling is not new in the HCI literature [3, 8, 12]. However, while the existing approaches are targeted to generate entire plots of a story, in this paper we describe our own approach to interactive storytelling, in which automated planning is used to interpret player attitudes and preferences to produce a coherent narrative progression consisting of predefined story elements.

The rest of the paper is organized as follows. In Section 2, we discuss our approach to represent interactive stories as solutions of automated planning problems. In Section 3, we present a mobile application for cultural heritage based on mini games that we realized to show the feasibility of our approach. Finally, in Section 4, we provide some details on how a planning problem for story generation can be modeled, while in Section 5 we conclude the paper.

2 APPROACH

Our approach to interactive storytelling is particularly suited for stories that consist of interleaved periods of interactive play and cut scenes, i.e., short non-interactive scenes that transition from one play to the next, providing the player with goals and motivation for the next segment of play. In this mode of alternating between story and cut scenes, story elements and narrative progression are kept strictly separate.

As shown in Figure 1, our approach consists of the following methodological steps:

1. the *author* of the story is responsible of creating, with the support of the *Domain Builder* and the *Problem Builder* components, the "structural" aspects of the story, such as the locations where the story takes place, the interactive objects, the plot constraints, the characters, etc. In addition, a list of actions to progress from a story element (i.e., an interactive play) to another must be formalized as well. The Domain/Problem Builder components translate the above inputs in a PDDL planning domain and problem, respectively;
2. the *player* can request the compilation of a personalized narration through an application installed on her/his mobile device, providing her/his attitudes/preferences for the story, e.g., the length and the complexity of the interactive story. In a nutshell, the player is implicitly defining some aspects of the story s/he wants to experience, thus modifying the planning problem;
3. the *planner* interprets the above information and generates a *plan*, i.e., a coherent narrative progression between story elements that satisfies the player's preferences and plot constraints. To this end, the *Plan Builder* component translates a generated plan in a format that is interpretable by the player's mobile device.

Note that we assume interactive story elements to be already existing and predefined, i.e., they are seen by the planner as black boxes to be properly organized to create a valid narration. At play-time, if the completion of a story element produces an outcome that deviates from the player's preferences, *replanning* activities can be performed to support *story variation*, i.e., to adapt on-the-fly the flow of the narration.

3 MOBILE APPLICATION

To test the feasibility of our approach, we implemented a mobile application called "Cúltura", which supports a player to the visit

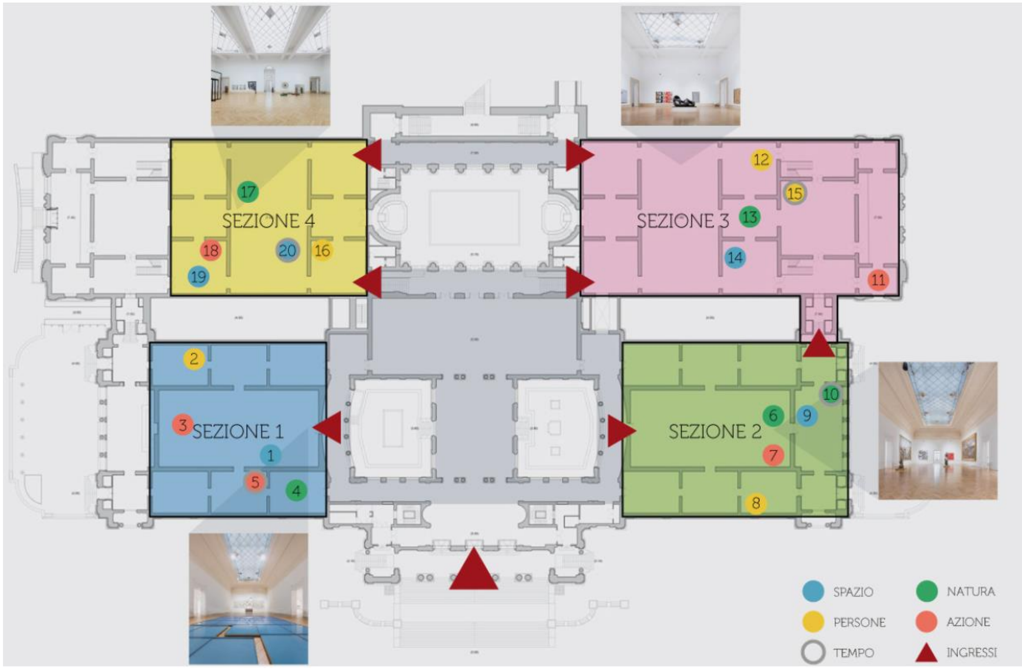

Figure 2: Map of the museum with the locations and thematics of the 20 paintings of interest.

of the most relevant paintings of National Gallery of Modern and Contemporary Art[1] (situated in Rome, Italy) through playing 5 out of 20 available mini games. The plot behind the playing experience is inspired to the well known fantasy novel: "*The Neverending Story*". The idea is that by playing 5 out of the 20 mini games in some specific (and not predefined) orders allows the player to experience different narrations towards the objectives underlying the story. Specifically, each mini game is focused to improve the understanding of a specific painting of the museum, e.g., by asking the player to complete a puzzle made up of mixed pieces of the painting (e.g., see the mini game in Figure 3), to identify anomalies or hidden details in the painting, etc. A score is assigned to the player depending on the time required and mistakes done to complete the mini game.

As shown in in Figure 2, the 20 paintings are grouped in 5 different *thematics* (4 paintings per thematic) located in 4 different *sections* of the museum. Any section hosts exactly 5 paintings. Before starting the visit, the player can provide preferences to personalize her/his gaming experience, such as which thematic s/he is interested in, which sections of the museum s/he wants to visit, the complexity of the gaming experience (easy, medium or hard), etc. On the basis of these inputs, the planner generates a personalized narration for the player, based on a flow of 5 games to be executed in sequence that respect the player's preferences and the story constraints.

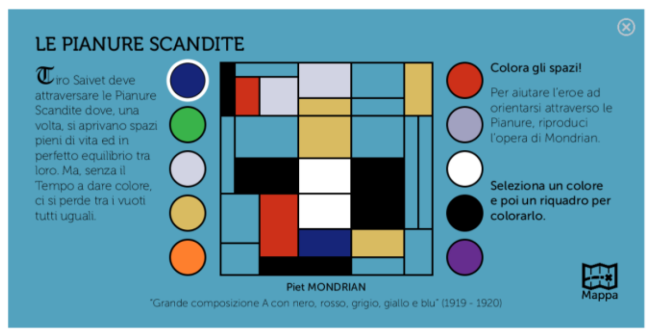

Figure 3: A puzzle-based mini game.

After the completion of a mini game (i.e., of an interactive play), a cut scene providing more details on the painting just visited is shown to the player, together with the information to move to the next location of the story. In this non-interactive time span, the mobile application interprets and reasons over the outcomes of the previously completed mini game (in terms of time needed/errors made to complete it), in order to quantify if its *real complexity* has been in line with the *expected complexity* as declared by the player at the beginning of the visit. In case that a game is perceived as too complex for a player, rather than executing the next expected game, the planner automatically generates a new (partial) narration (with a replanning activity) made up of "easier" mini games to complete the gaming experience.

[1] http://lagallerianazionale.com/en/

From a technical perspective, we developed the mobile application for Android smartphones and tablets using the Processing[2] programming language. The Fast-Downward planner[3] was used for generating the personalized stories.

4 MODELING STORIES AS PLANNING PROBLEMS

In this section, on the basis of the case study presented in Section 3, we provide some details on how modeling stories as planning problems in PDDL.

Technically, PDDL allows one to formulate a *planning problem* $P = \langle I, G, P_D \rangle$, where I is an initial state of the world, G is a goal state, and P_D is the planning domain. In turn, a planning domain P_D is built from a set of predicates describing the state of the world (a state is characterized by the set of predicates that are true) and a set of actions that can be executed in the domain. Each action is of the form $a = \langle Par_a, Pre_a, Eff_a, Cost_a \rangle$, where Par_a is the list of input parameters for a, Pre_a and Eff_a specify the preconditions and effects of a, in terms of the set of predicates, and $Cost_a$ expresses the action's cost. A solution for a planning problem is a sequence of actions—*a plan*—whose execution transforms the initial state I into a state satisfying the goal G.

One immediate observation is that PDDL focuses on an action-centered representation that models single aspects of the world, namely which properties are *true* at a given moment. In other words, PDDL is used to model facts like:

"Painting 1 is included in the thematic Space."
"Painting 2 is located in section 1 of the museum, etc."

The statements above can of course be directly represented by means of predicates in the planning domain, as for instance, through a pair of predicates *At(x,y)* and *In(x, z)*, being *x* a painting, *y* a thematic and *z* a section of the museum, i.e., the objects of the planning domain. Similar predicates can be used to keep track of the paintings visited (e.g., *visited(x)*), and of the pairs of paintings *x1* and *x2* that can not be visited together in the context of a same narration (e.g., *incompatible(x1,x2)*).

In the planning problem, the author must first clarify which are the concrete instances of the objects declared in the planning domain, e.g., all the twenty paintings, the four sections of the museum, and the available thematics, that are of kind *Space*, *Time*, *People*, *Action*, and *Nature*. Then, s/he defines which domain predicates are true in the initial state. For example, *At*(2, 1) and *At*(20, 4) indicate that paintings 2 and 20 are located in sections 1 and 4 of the museum, respectively. Similar instantiations, whose value depends on the structure of the museum (cf. Figure 2), can be provided for the other predicates defined in the planning domain. While the values of the above predicates remain unchanged in any gaming experience, the starting value of predicate *incompatible(x1,x2)* changes depending on the preferences expressed by the player. For example, if the player wants to visit a single section (or is interested in a single thematic) of the museum, all the paintings belonging to the other sections (or thematics) must be declared as incompatible with the paintings of the preferred section (or thematic). Finally, the goal condition states that five paintings must be visited at the end of the narration, guaranteeing that the complexity of the interactive experience matches the player's preferences.

To formalize the visit of a painting in the planning domain, we need to model a single planning action *visit(x)*, which can be executed if the painting *x* has not been already visited (i.e., *not(visited(x))* is true), and if it does not exists any other painting *x2* that has been already visited and is incompatible with *x*. The effect of the action is that *x* turns out to have been visited, and the total cost of the plan under construction is increased of a value that corresponds to the action's cost, i.e., to the cost of visiting *x*, which is related to the complexity of playing the mini game associated to *x*.

A plan will have the form of a sequence of *visit* actions, which is interpreted by the Plan Builder (cf. Figure 1) as an ordered sequence of mini games to be played to enact the personalized visit experience through the mobile application. The author is required to define the planning domain and most of the planning problem (in particular, the part capturing the static structure of the world where the story takes place) just once. On the other hand, the fragment of the planning problem that reflects the player's preferences (in our case, the values of predicate *incompatible* and the maximum allowed complexity of the playing experience) will be dynamically generated any time a new story begins.

5 CONCLUDING REMARKS

In this paper, we have discussed how the use of automated planning in AI can be employed to generate personalized narrative experiences in interactive storytelling. We observe that our planning-based approach is independent of the paradigm selected to represent the story, whether plot-based or character-based. The key idea is that if the structure and the dynamics of a story can be converted into a planning problem in PDDL, one can seamlessly leverage the recent version of the best performing planner to automatically generate a narration customized to the player's preferences. Since our approach is intended to be a practical and intuitive way to support the story developers as a kind of programming methodology, as future work we aim at devising an explicit list of guidelines to systematically create interactive stories as planning domains and problems in PDDL.

Acknowledgments. This work has been partly supported by the H2020 ERC project NOTAE (grant #786572), the Sapienza grants IT-SHIRT and METRICS, the Lazio regional initiative "Centro di eccellenza DTC Lazio", the project ARCA, and the Italian project SM&ST.

[2] https://processing.org/
[3] http://www.fast-downward.org/

REFERENCES

[1] Ruth Aylett, Sandy Louchart, Joao Dias, Ana Paiva, Marco Vala, Sarah Woods, and Lynne Hall. 2006. Unscripted narrative for affectively driven characters. IEEE Comp. Graph. and App. 26, 3 (2006).

[2] Marc Cavazza, R. Michael Young. 2017. Introduction to interactive storytelling. Handbook of Digital Games and Entertainment Technologies, 377-392 (2017).

[3] Fred Charles, Miguel Lozano, Steven J Mead, Alicia Fornes Bisquerra, and Marc Cavazza. 2003. Planning formalisms and authoring in interactive storytelling. Fraunhofer IRB Verlag.

[4] Hector Geffner and Blai Bonet. 2013. A concise introduction to models and methods for automated planning. Synthesis Lectures on Artificial Intelligence and Machine Learning 8, 1 (2013), 1–141.

[5] Malik et al. Ghallab. 1998. PDDL - The Planning Domain Definition Language. AIPS-98 (1998).

[6] Andrew Glassner. 2009. Interactive storytelling: Techniques for 21st century fiction. AK Peters/CRC Press.

[7] Akrivi Katifori, Manos Karvounis, Vassilis Kourtis, Sara Perry, Maria Roussou, Yannis Ioanidis. 2018. Applying Interactive Storytelling in Cultural Heritage: Opportunities, Challenges and Lessons Learned. In International Conference on Interactive Digital Storytelling (pp. 603-612). Springer, Cham

[8] Andrea Marrella and Stavros Vassos. 2014. Story Generation in PDDL Using Character Moods: A Case Study on Iliad's First Book. In Artificial Intelligence: Methods and Applications. Springer.

[9] Michael Mateas and Andrew Stern. 2002. Towards integrating plot and character for interactive drama. In Socially Intelligent Agents. Springer.

[10] Massimo Mecella, Francesco Leotta, Andrea Marrella, Federico Palucci, Chiara Seri, and Tiziana Catarci. 2018. Encouraging persons to visit cultural sites through mini-games. EAI Endorsed Transactions on Serious Games 18, 14 (1 2018). DOI:http://dx.doi.org/10.4108/eai.4-1-2018. 153531

[11] Mark O Riedl and AndrewStern. 2006. Believable agents and intelligent story adaptation for interactive storytelling. In Int. Conf. on Technologies for Interactive Digital Storytelling and Entertainment. Springer.

[12] R Michael Young. 2000. Creating interactive narrative structures: The potential for AI approaches. Psychology 13 (2000), 1–26.

[13] R Michael Young, Mark O Riedl, Mark Branly, Arnav Jhala, RJ Martin, and CJ Saretto. 2004. An architecture for integrating plan-based behavior generation with interactive game environments. Journal of Game Development 1, 1 (2004), 51–70.