Game Cinematography: from Camera Control to Player Emotions

Paolo Burelli

Abstract Building on the definition of cinematography (Soanes and Stevenson, 2005), game cinematography can be defined as the art of visualizing the content of a computer game. The relationship between game cinematography and its traditional counterpart is extremely tight as, in both cases, the aim of cinematography is to control the viewer's perspective and affect his or her perception of the events represented. However, game events are not necessarily pre-scripted and player interaction has a major role on the quality of a game experience; therefore, the role of the camera and the challenges connected to it are different in game cinematography as the virtual camera has to both dynamically react to unexpected events to correctly convey the game story and take into consideration player actions and desires to support her interaction with the virtual world. This chapter provides an overview of the evolution of the research in virtual and game cinematography, ranging from its early focus on how to control and animate the virtual camera to support interaction to its relationship with player experience and emotions. Furthermore, we will show and discuss a number of emerging research directions.

1 Introduction

Cinematography, over the last two centuries, has undergone a constant evolution: from the first experiments with machines such as the zoetrope (Enticknap, 2005) to the last advancements in three-dimensional cinematography and computer graphics. Throughout its history, it developed into a complex field dealing with techniques and methods to present the visual discourse. With the advent of three-dimensional computer graphics, a new branch of cinematography developed called virtual cinematography. At first limited to special effects and short films, after the first release of a fully computer generated film (Lasseter, 1995), virtual cinematography has grad-

Paolo Burelli
Aalborg University, A.C. Meyers Vænge 15, 2450 København (DK). e-mail: pabu@create.aau.dk

ually become a field on its own, even if largely intertwined with its traditional counterpart.

One of the most clear distinctions between the two fields lies in the differences between a virtual and a real-world camera. A virtual camera is an abstract construct that defines the way the virtual world is presented to the user; it is designed to simulate the behaviour of a real-world camera and, at the current state-of-the-art in computer graphics, images produced through a virtual camera can be indistinguishable real footage. However, contrary to a real-world camera, a virtual camera has no physical presence in the virtual world and its properties can change over time and adapt to the events filmed (Haigh-Hutchinson, 2009).

These advantages allow virtual cameras to film scenes with much higher freedom and expressiveness; however, this can potentially contrast with the filming conventions developed in traditional cinematography. Such conventions describe aspects like the way the camera should be placed or the way it should be moved to make a transition between two different scenes, and adherence to these conventions is often important to generate a *cinematographic* experience and not to disorient the viewer – e.g. by crossing the line of action (Arijon, 1991).

Furthermore, new application areas such as interactive narratives and computer games have further expanded the difference between real and virtual cinematography as new conventions and methods needed to be studied to address the specific characteristics of these new media. Interactive narratives, for instance, reduce the control of the designer over the *mise en scene* – i.e. the arragement of the elements in the scene – thus requiring more flexible and intelligent methods to control the camera.

Researchers have been studying effective and efficient solutions to address these differences and assist virtual designers and programmers to translate cinematographic conventions to virtual cinematography. Early research works focused on the problem of manually handling a virtual camera through input devices such as a keyboard and a mouse (Ware and Osborne, 1990); thereafter, the focus gradually shifted towards the problem of automating and assisting camera movements so that the camera could be animated in complex and constantly changing environments (Christie et al., 2008). This is the case, for instance, of interactive narratives, in which the events of a story depend on the choices and actions of the viewer. This medium, which can be in many ways seen as natural evolution of films, is one of the dominant application areas in the current state-of-the-art, and different researchers have investigated the translations of classic cinematography conventions to this medium: Jhala and Young (2005), for example, have investigated the automatic generation of shot plans for emergent stories, while Lino et al. (2010) have investigated automatic camera placement and animation in such contexts.

Beyond the differences in terms of unpredictability of the environment and events, the purpose of virtual cinematography in interactive narratives is, in large part, similar to non-interactive narratives: the camera has to support the storytelling. However, if interactive narrative is analysed as a component of a medium such as computer games, the focus of the cinematography and the purpose of the camera shifts away from solely supporting narration.

Computer games are a highly interactive medium and the virtual camera is responsible for supporting both the interaction and the visualisation of the game events, and, while traditionally the virtual camera supports narration and interaction in different phases of a game (Haigh-Hutchinson, 2009), there are a number of commercial examples of games in which these two aspects overlap (Raynal, 1992; Mikami, 1996; Toyama, 1999; Cage, 2010). Furthermore, different studies on the relationship between player emotions and cognition reveal that, in computer games, virtual cinematography is deeply intertwined with player experience. Evidence suggests that this relationship goes beyond the conventions of classical cinematography (Martinez et al., 2009; Burelli, 2013; Burelli and Yannakakis, 2015).

Both the role of the player as well as its relationship with the virtual camera and the visual experience define a clear distinction between classical cinematography and *game cinematography*. In game cinematography, the player has an active role in changing the game events thus it directly influences the movements of the camera. Furthermore, to affect the player experience, in game cinematography, the camera needs to be aware of the current state of the player both inside and outside of the game, thus, creating an indirect relationship between the player and the camera.

In this chapter, we present the concept of game cinematography from its foundations in virtual cinematography to the latest studies on player experience. We start by giving an overview of the game industry's perspective on game cinematography in Section 2. In Section 3, we present the state of the art in camera animation and placement. In Section 4, we discuss different methods to create plans of camera movements and shots to present an interactive story. In Sections 5 and 6, the focus shifts on the relationship between the player and different aspects of game cinematography. Finally, in Section 7, we highlight a number of possible future directions of research in game cinematography.

2 Camera control in computer games

Game cinematography has shown a low degree of experimentation in the game industry, especially in comparison to other aspects such as rendering or physics. With few exceptions, there is a strict dichotomy between interactive cameras and cinematic cameras (Haigh-Hutchinson, 2009) in which the first one is used during the gameplay, while the second one is used for storytelling during cut-scenes.

Christie et al. (2008) divide the camera control styles in games in the following three categories:

First person: The camera position and orientation corresponds to the player's location in the virtual environment; therefore, the camera control scheme follows the character control scheme. Examples of games adopting such a camera control scheme include *Doom* (Id Software, 1993) and *Halo: Combat Evolved* (Bungie Studios, 2001).

Third person: The camera shows the events in the game from an external perspective. This perspective can be freely controllable by the player or bound to

(a) A screen-shot from Heavy Rain (Cage, 2010), demonstrating usage of cinematographic techniques.

(b) A screen-shot from Gears Of War (Bleszinski, 2007) during a running action; in such context the camera moves downward and shakes to enhance the haste sensation.

Fig. 1: Examples of advanced camera control in modern computer games.

specific locations, orientations or characters. Examples of games using such type of camera control paradigm are action games such as *Tomb Raider* (Core Design, 1996), in which the camera follows the character from a fixed distance with different angles to avoid obstacles in the environment, or strategy and sport games – e.g. *Starcraft* (Blizzard Entertainment, 1998) – in which the camera is freely movable by the player who can select different targets. In another form of third-person camera control scheme, which Haigh-Hutchinson (2009) calls *predetermined*, multiple cameras are pre placed around the environment and, during the game, the perspective switches between them – e.g. Devil May Cry (Kamiya, 2001).

Cut-scenes and replays: In these non-interactive phases of the games, the camera focuses on representing the important elements of the story without the need to support interaction. It is often used in sport games (replays) and in story-heavy games (cut-scenes). Games featuring such a camera control scheme include *Metal Gear Solid* (Kojima, 1998) or most sport video games.

In recent years, the separation between interactive and cinematic cameras is becoming less distinct as more games are employing cinematographic techniques to portrait narrative and interaction in games. Examples such as *Heavy Rain* (Cage, 2010) or *Silent Hill* (Toyama, 1999) show extensive usage of cinematic techniques to frame the in-game actions (see Figure 1a). In such games, however, the cameras are set manually in place during the development of the game; reducing heavily the movement and the actions the player can take. Furthermore, achieving the same level of quality in a game in which the content is not known in advance (e.g. it is procedurally generated (Yannakakis and Togelius, 2011; Shaker et al., 2010) or crowd sourced, such as in *World Of Warcraft* (Blizzard Entertainment, 2004)) is still an open challenge.

Some custom dynamic techniques have been implemented in different games to achieve a more cinematographic experience in more action oriented games. For instance, in *Gears Of War* (Bleszinski, 2007) (see Figure 1b), the camera changes relative position and look-at direction automatically to enhance some actions or to

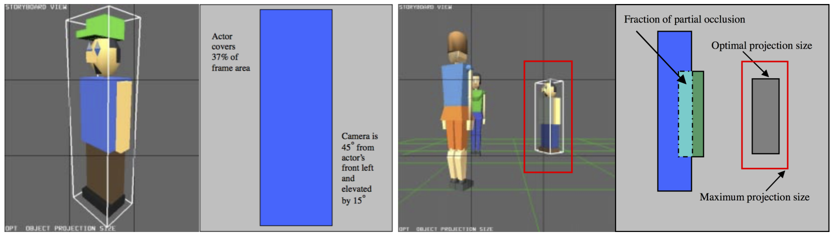

(a) (b)

Fig. 2: Examples of frame constraints and their relative geometrical definition for a medium three-quarter shot and three character shot. Image adapted from (Bares et al., 2000).

allow for a better view of the environment. Another example is the slow motion feature implemented in the Max Payne series (Remedy Entertainment, 2001). One of the few examples of a general camera control system capable of handling aspects such as composition and camera movements in dynamic environments has been developed by 10Tacle Studios (Hamaide, 2008) as an extension of a system proposed by Bourne et al. (2008).

3 Automatic camera control

Under different labels, the research in virtual cinematography has distributed along three main directions: planning and definition of the shots, automatic composition and real-time camera animation. The first can be described as the task of defining a sequence of shots to visualise one or more events in a virtual environment. The second is the process of translating these shots in actual camera configurations, while the last one is the process of animating the camera during the shots and ensuring smoothness between them.

One of the first examples of a system addressing automatic camera placement and animation was presented by Blinn (1988) who designed a system to automatically generate views of planets in a space simulator of NASA. Although limited in its expressiveness and flexibility, Blinn's work inspired many other researchers trying to investigate efficient methods and more flexible mathematical models to handle more complex aspects such as camera motion and frame composition (Arijon, 1991).

More generic approaches model camera control as a constraint satisfaction problem. These approaches require the designer to define a set of desired frame properties, which are then modelled either as an objective function to be maximised by the solver or as a set of constraints that the camera configuration must satisfy. These constraints describe how the frame should look like in terms of object size, visibility and positioning. Bares et al. (2000) presented a detailed definition of these constraints, which became the standard input of most automatic camera control methods. Examples of a few of these constraints can be seen in Figure 2: fore example, in Figure 2a, the projection size of the character is set to 37% of the frame area, while in Figure 2b, the occluded area of the rightmost character is set to a maximum of 10% of its overall projected area.

The problem of finding one or more camera configurations that satisfy a given set of frame constraints has been initially tackled by Bares et al. (1998) using a constraint satisfaction method with a constraint relaxation technique to identify and unselect incompatible constraints. This approach used a bi-dimensional spherical map, which proved inaccurate when multiple far subjects were evaluated at the same time. Bares et al. (2000) extended this initial approach by defining a sub-space of valid camera configurations for each constraint; the intersection of the valid spaces is then sampled to find the best camera configuration. The same principle of combining constraint satisfaction and search to find the best camera configuration has been further extended by improving the volume selection and integrating more sophisticated search algorithms (Lino et al., 2010; Burelli et al., 2008).

Pure optimization approaches, such as CAMPLAN (Olivier et al., 1999) or the Smart Viewpoint Computation Library (Ranon and Urli, 2014), implement a more flexible search strategy that models all frame constraints as an objective function (a weighted sum of each constraint) allowing for partial satisfaction of any constraint. These approaches do not prune any part of the search space and the satisfaction of the different frame constraints is prioritized by associating a weight to each constraint. The flexibility of such approaches comes with the price of a high computational cost. This aspect becomes a particularly critical factor when the algorithm is intended to deal with real-time dynamic virtual environments. In this context, the controller has to be able to calculate a reasonable camera configuration at short intervals (a few milliseconds) to be able to ensure synchronization with the scene changes and to have minimal impact on the overall application execution.

A more efficient approach to optimization for camera composition consists of employing local-search to find the best solution. Beckhaus et al. (2000) investigated first the application of local search algorithms to camera control. Their system used Artificial Potential Fields (APFs) to guide the camera through a museum and generate smooth virtual tours. Bourne et al. (2008) proposed a system that employed sliding octrees to guide the camera to the optimal camera configuration. Burelli and Jhala (2009) extended these two approaches to include frame composition and support multiple-object visibility. Local search approaches offer reasonable real-time performance as they perform a small sampling of all the possible camera configurations; however, they are often unable to calculate correct camera configurations when visibility for a specific subject is required on the frame (Burelli and Yannakakis, 2010).

Ensuring visibility of one or more subjects is one of the most critical objectives of virtual camera composition as object visibility plays a key role in frame composition (Christie et al., 2008). Evaluation of a subjects' occlusion can be either integrated in the search process as one part of the objective function to be optimized or as an extra heuristic to prune the search space or guide the search process. Bourne et al. (2008) propose to override the current search process in case of missing visi-

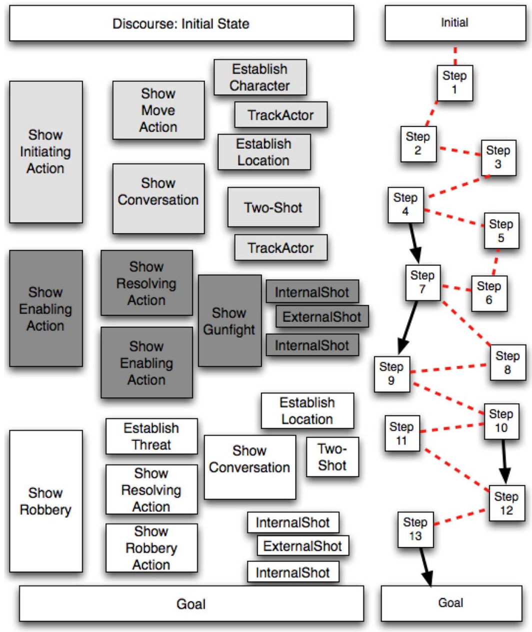

Fig. 3: An example of a cinematic discourse plan generated (left side) from a story plan (right side). Image adapted from (Jhala and Young, 2010).

bility of the tracked subject by introducing a cut; their approach, however, considers just one object of interest. Pickering (2002) suggested a shadow-volume based approach to identify volumes of space without occlusion: the search process would take place only in these volumes. Christie et al. (2012), in a similar fashion, prune the space of the possible camera configurations by calculating a visibility volume – i.e. a space of camera configurations in which visibility is guaranteed – for one or more targets; moreover, this approach takes into consideration also the temporal aspect of the occlusion to make the camera more robust to temporary and sudden occlusions.

4 Story-driven interactive cinematography

The research works described in the previous section focus primarily on the translation of high-level cinematographic requirements (e.g. an object's visibility or its position on the screen) into low level camera parameters (e.g. position and rotation). A number of researchers have instead focused on the problem of identifying the best high-level requirements for a specific event in a virtual environment with the objective of creating a coherent a cinematographically correct visualisation of a given story. Figure 3 shows and example of a list of camera actions described as abstract shots (third column on the left side), which should be used to drive the camera and visualised the story described in the plan on the right side of the picture.

The first research work focusing on shot planning was published by Christianson et al. (1996): they proposed a language (DCCL) to define shot sequences and to automatically relate such sequences to events in the virtual world. Each shot is encoded in an idiom that describes also the conditions in which the shot should be selected. He et al. (1996) extended the concept of idioms within DCCL by modelling them as hierarchical finite state machines and allowing for richer expressiveness.

McDermott et al. (2002) developed further this idea by allowing conditional transitions between idioms and a visual definition of the shot for each idiom. Likewise, Charles et al. (2002) expanded the logic employed to select a shot by proposing a number of semantic rules to prioritise the shooting of specific objects or actions during the story. El-Nasr (2002) followed a slightly different direction in her work, employing reactive planning and focusing on the integration of the shot selection process and scene lighting.

Based on the same principle – i.e. visualising a story from a given set of events – different researchers have refined further the aforementioned planning approaches (Jhala and Young, 2010) and have shown different applications such as comics generation (Thawonmas et al., 2010) and game replays generation (Dominguez et al., 2011).

One common aspect among these approaches is their focus on story-visualisation: on one hand only a handful of these studies explicitly target non-interactive productions (Thawonmas et al., 2010; Dominguez et al., 2011), on the other hand, the studies that target interactive narratives, focus primarily on the emergency aspect of the narrative rather than the user's interaction. In other words, user interaction is seen just as the cause of the changes in the narrative that in turn drive the changes in the cinematography, which aims at correctly supporting the communication of the new narrative.

5 Camera and player interaction

Although, story-driven cinematography can be easily identified as the dominant approach to virtual and game cinematography, a number of alternative approaches taking into consideration player preferences and player interaction have been proposed.

One of the first studies following this direction has been presented by Bares et al. (1998). In their work, they investigated how the user can influence the cinematographic experience and proposed a system that selects the most appropriate camera settings depending on the user's tasks in a virtual learning environment. Halper et al. (2001) followed a similar direction by adjusting the camera to accommodate the player's actions; furthermore, they devised a mechanism to predict short-term player movements to improve the camera animation smoothness.

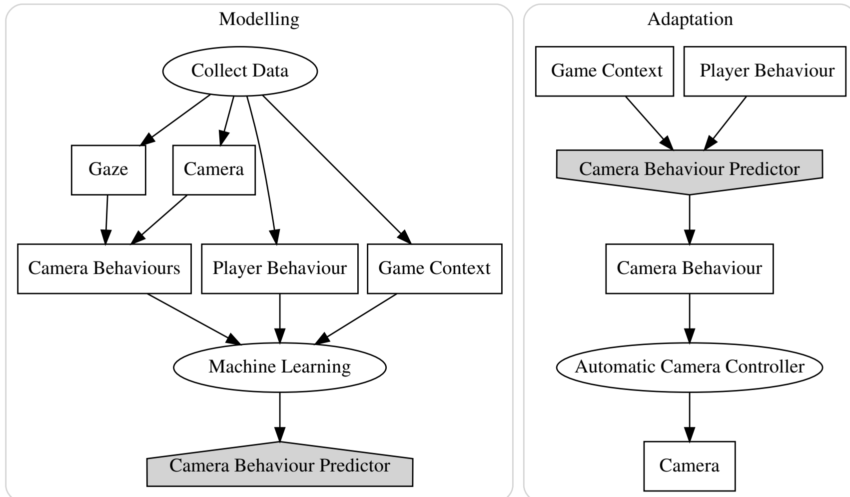

Fig. 4: Camera behaviour modelling and adaptation phases of the adaptive camera control methodology proposed by Burelli and Yannakakis (2015)
.

Both the work by Halper et al. and the one by Bares et al. adapt the camera requirements based on a number of predefined directives designed to support user interaction; however, while the behaviour of the system would adapt depending on the task, the user preferences and the effect on the user experience is not taken into consideration. Bares and Lester (1997) investigated the idea of modelling the camera behaviour according to the user preferences to generate a personalised cinematographic experience. In the system they proposed and evaluated, the user model construction required the user to specifically express some preferences about the style for the virtual camera movements.

The results of the evaluation of the work by Bares and Lester highlighted, for the first time, the importance of the relationship between the viewpoint and user interaction and suggested a direction to follow to leverage this relationship. However, the profile building procedure suggested was explicit and required the users to be conscious about the desired camera behaviour and also to be sufficiently competent to instruct the system to achieve such a behaviour.

One method to extrapolate information about the above aspects has been employed by Sundstedt et al. (2008) who conducted an experimental study to analyse players' gaze behaviour during a maze puzzle solving game. The results of their experiment show that gaze movements, such as fixations, are heavily influenced by the game task. They conclude that the direct use of eye tracking during the design phase of a game can be extremely valuable to understand where players focus their attention, in relation to the goal of the game.

Picardi et al. (2011) and Burelli and Yannakakis (2011) investigated the possibility to employ players' gaze to build user models of camera behaviour; in these works, the virtual camera behaviour is modelled on the amount of time the player spends framing and observing different objects in a virtual environment while playing a game. Combining camera movements with eye movements (i.e. fixations and pursuits) in a visually rich virtual environment such as a computer game, allows to identify exactly which objects drive the player attention (Irwin, 2004; Thorson, 1994) and, therefore, can be used to build a user model of visual attention.

This model has been later employed by Burelli and Yannakakis (2015) who extended the idea of user modelling of camera behaviour (Bares and Lester, 1997) by implicitly building the models from in-game player behaviour and its relationship with the player's eye movements. As show in Figure 4, through these models, the camera controller detects in real-time what objects will the player desire to see and it can generate appropriate camera requirements to keep these objects on screen. The results of the study by Burelli and Yannakakis (2015) show that adapting the camera behaviour based on user models of visual attention has the potential to improve the quality of the user experience. In particular, while not effective for all users, the user models proved effective in supporting the player interaction mostly improving the results achieved by the players.

6 Affective cameras

Another fundamental aspect of player experience that has been studied in relationship to virtual cinematography is the player's affective state, both in terms of the ability of the viewer to understand character emotions and in terms of the effect of the cinematographic choices on the viewer's affective state.

Studies on emotions and cinematography in non-interactive media have investigated connections between low level cinematographic features and the viewer's experience; for instance, Simons et al. (1999) studied the relationship between camera and object motion and the emotional responses of humans, concluding that an increase of motion intensity on the screen causes an increase in the viewer's arousal. Hanjalic and Xu (2005) studied further the relationship between viewer's emotions and cinematography and developed a deterministic model of arousal and valence based on a combination of on screen motion, shots rhythm and sound energy. Sun and Yu (2007) followed the same approach employing a non-deterministic method (i.e. Hidden Markov Models) to model the relationship between the aforementioned cinematographic features and four affective states: joy, anger, sadness and fear.

The first study exploring the role of emotions in interactive cinematography has been presented by Tomlinson et al. (2000), who prosed a bottom-up approach to shot definition based on a number of camera's affective states. In *CameraCreature*, there is no plan driving the movements of the camera, which is instead modelled as an agent moving in the virtual environment and reacting to the actions of the other agents. The agent controlling the viewpoint has an ethologically inspired structure based on sensors, emotions, motivations and action-selection mechanisms. The camera agent shares this structure with all non-player characters (*Actors*) in the virtual environment and is able, through virtual sensors, to detect their emotions and the type of action they are performing. Tomlinson et al. envision their work as *"a first step toward a future of interactive, emotional cinematography"*, which can be seen as an early definition of the concept of game cinematography. However, while the architecture proposed is potentially flexible enough to consider player emotions and actions, no directions are given on how this could be done.

Yannakakis et al. (2010) studied the impact of camera viewpoints on player experience and developed a computational model to predict this impact, demonstrating the existence of a relationship between player emotions, physiological signals and camera parameters. However, the features employed in the model described cinematography using low-level camera parameters such as height or distance, which are unable to express the content of the visualised images. Burelli (2013) performed a study that extended the aforementioned work by analysing the camera behaviour in terms of composition and by extending the analysis across different genres with richer game mechanics. The results confirm some of the findings revealed by Yannakakis et al. (2010), but there is evidence that the relationship between camera behaviour and player experience can be better explained by describing the cinematographic experience through more high-level features, such as shot spacing or symmetry, as these features allow us to understand what is the visual content that is reproduced on screen. Furthermore, the results reveal that the task the player performs affects the relationship between experience and visualisation.

7 Future directions in game cinematography

Cinematographic games are a rising genre in the computer games industry and an increasing number of titles published include some aspects of cinematography in the gameplay or the storytelling. At the present state, camera handling in computer games is managed primarily through custom scripts and animations, and there is an inverse relationship between player freedom and cinematographic quality. However, the studies described in the previous sections show that there is a strong potential for improvement on the current state-of-art in game cinematography, especially towards building a better understanding of the impact of cinematography on player experience and how this could be leveraged to make better and new types of games.

Consequently, we see a number of future research directions that could be pursued. For instance, to foster a stronger awareness in the application of cinematography to games it would be extremely important to develop a taxonomy of game cinematography similar to the one that has been developed in classic cinematography throughout its history. Such taxonomy should be built following the approach delineated by Yannakakis et al. (2010) and Burelli (2013) as this would allow game designers and developers to make choices on game cinematography with awareness of their impact on player experience and player interaction.

Furthermore, we envision a stronger interconnection of game cinematography with procedural content generation: this would allow to further develop the presentation aspect of game generation helping to move the field towards complete automatic game generation (Togelius et al., 2013). In order to do this, it would be necessary to extend the focus of game cinematography research towards a more general analysis of game presentation, including aspects such as visual aesthetics (Liapis et al., 2012) and to study its perception and its effects on the user's understanding of the virtual world events. For instance, investigating how semiotics and narrative cognition can be integrated within the frameworks of automatic content generation and virtual cinematography, would allow the generation and the visualisation to become one coherent process which takes into account the signification of the content and the events (Eco, 1984).

Finally, following the trends in embodied artificial intelligence (Chrisley, 2003) and artificial intelligence for physical games (Frazier and Riedl, 2014), it could be possible to investigate the application of the results achieved in virtual game cinematography back to the physical world. For instance, thanks to the recent advancements in fields such as computer vision and robotics (He et al., 2006; Krizhevsky et al., 2012; Meier et al., 2012) and the introduction of ever more miniaturised filming equipment, micro unmanned aerial vehicle could be used as intelligent autonomous agents that could film a physical game for remote gaming or game broadcasting and recording.

8 Conclusion

In this chapter, we have presented an overview of the field of virtual cinematography and its application to computer games, and we defined what specifically characterised game cinematography. Moreover, we have analysed how the focus of the research in the field is shifting from pure algorithmic studies, focused on developing more robust and efficient algorithms to automatically animate the virtual camera, to studies which analyse different aspects of the relationship between visualisation, story and player interaction. In particular, in both game and traditional cinematography, there is a growing interest in understanding the impact on the viewer's cognitive processes and affective state of camera movements, editing and other cinematographic aspects. Finally, we have highlighted a number of possible future directions for game cinematography both towards a deeper understanding of the player's cinematographic experience and towards new applications of game cinematography beyond traditional computer games.

References

Arijon, D. (1991). *Grammar of the Film Language*. Silman-James Press LA.

Bares, W. H. and Lester, J. C. (1997). Cinematographic User Models for Automated Realtime Camera Control in Dynamic 3D Environments. In *International Conference on User Modeling*, pages 215–226, Chia Laguna, Italy. Springer-Verlag.

Bares, W. H., McDermott, S., Boudreaux, C., and Thainimit, S. (2000). Virtual 3D camera composition from frame constraints. In *ACM Multimedia*, pages 177–186, Marina del Rey, California, USA. ACM Press.

Bares, W. H., Zettlemoyer, L. S., Rodriguez, D. W., and Lester, J. C. (1998). Task-Sensitive Cinematography Interfaces for Interactive 3D Learning Environments. In *International Conference on Intelligent User Interfaces*, pages 81–88, San Francisco, California, USA. ACM Press.

Beckhaus, S., Ritter, F., and Strothotte, T. (2000). CubicalPath - Dynamic Potential Fields for Guided Exploration in Virtual Environments. In *Pacific Conference on Computer Graphics and Applications*, pages 387–459.

Bleszinski, C. (2007). Gears Of War.

Blinn, J. (1988). Where Am I? What Am I Looking At? *IEEE Computer Graphics and Applications*, 8(4):76–81.

Blizzard Entertainment (1998). Starcraft.

Blizzard Entertainment (2004). World Of Warcraft.

Bourne, O., Sattar, A., and Goodwin, S. (2008). A Constraint-Based Autonomous 3D Camera System. *Journal of Constraints*, 13(1-2):180–205.

Bungie Studios (2001). Halo: Combat Evolved.

Burelli, P. (2013). Virtual Cinematography in Games : Investigating the Impact on Player Experience. In *International Conference On The Foundations Of Digital Games*, pages 134–141, Chania, Greece. Society for the Advancement of the Science of Digital Games.

Burelli, P., Di Gaspero, L., Ermetici, A., and Ranon, R. (2008). Virtual Camera Composition with Particle Swarm Optimization. In Butz, A., Fisher, B., Krüger, A., Olivier, P., and Christie, M., editors, *International symposium on Smart Graphics*, volume 5166 of *Lecture Notes in Computer Science*, pages 130–141, Berlin, Heidelberg. Springer Berlin Heidelberg.

Burelli, P. and Jhala, A. (2009). Dynamic Artificial Potential Fields for Autonomous Camera Control. In *AAAI Conference On Artificial Intelligence In Interactive Digitale Entertainment Conference*, Palo Alto. AAAI.

Burelli, P. and Yannakakis, G. N. (2010). Global search for occlusion minimisation in virtual camera control. In *IEEE Congress on Evolutionary Computation*, pages 1–8, Barcelona. IEEE.

Burelli, P. and Yannakakis, G. N. (2011). Towards Adaptive Virtual Camera Control in Computer Games. In Dickmann, L., Volkmann, G., Malaka, R., Boll, S., Krüger, A., and Olivier, P., editors, *International symposium on Smart Graphics*, volume 6815 of *Lecture Notes in Computer Science*, pages 25–36, Bremen, Germany. Springer Berlin Heidelberg.

Burelli, P. and Yannakakis, G. N. (2015). Adaptive Virtual Camera Control Trough Player Modelling. *User Modeling and User-Adapted Interaction*.

Cage, D. (2010). Heavy Rain.

Charles, F., Lugrin, J.-l., Cavazza, M., and Mead, S. J. (2002). Real-time camera control for interactive storytelling. In *International Conference for Intelligent Games and Simulations*, pages 1–4, London.

Chrisley, R. (2003). Embodied artificial intelligence. *Artificial Intelligence*, 149(1):131–150.

Christianson, D., Anderson, S., He, L.-w., Salesin, D. H., Weld, D., and Cohen, M. F. (1996). Declarative Camera Control for Automatic Cinematography. In *AAAI*, pages 148–155, Portland, Oregon, USA. AAAI Press.

Christie, M., Normand, J.-M., and Olivier, P. (2012). Occlusion-free camera control for multiple targets. In *ACM SIGGRAPH/Eurographics Symposium on Computer Animation*, pages 59–64. Eurographics Association.

Christie, M., Olivier, P., and Normand, J. M. (2008). Camera control in computer graphics. In *Computer Graphics Forum*, volume 27, pages 2197–2218.

Core Design (1996). Tomb Raider.

Dominguez, M., Young, R. M., and Roller, S. (2011). Design and Evaluation of Afterthought , A System that Automatically Creates Highlight Cinematics for 3D Games. In *AAAI Conference On Artificial Intelligence In Interactive Digitale Entertainment*.

Eco, U. (1984). *Semiotics and the Philosophy of Language*.

El-Nasr, M. S. (2002). Story Visualization Techniques for Interactive Drama. In *AAAI Spring Symposium*, pages 23–28.

Enticknap, L. (2005). *Moving Image Technology: From Zoetrope To Digital*. Wallflower Press.

Frazier, S. J. and Riedl, M. O. (2014). Toward Using Games and Artificial Intelligence to Proactively Sense the Real World. In *AI & GAME Symposium*, London. AISB.

Haigh-Hutchinson, M. (2009). *Real-Time Cameras*. Elsevier.

Halper, N., Helbing, R., and Strothotte, T. (2001). A Camera Engine for Computer Games: Managing the Trade-Off Between Constraint Satisfaction and Frame Coherence. *Computer Graphics Forum*, 20(3):174–183.

Hamaide, J. (2008). A Versatile Constraint-Based Camera System. In *AI Game Programming Wisdom 4*, pages 467–477.

Hanjalic, A. and Xu, L. Q. (2005). Affective video content representation and modeling. *IEEE Transactions on Multimedia*, 7(1):143–154.

He, L.-w., Cohen, M. F., and Salesin, D. H. (1996). The virtual cinematographer: a paradigm for automatic real-time camera control and directing. In *ACM SIGGRAPH*, pages 217–224, New Orleans, Louisiana, USA. ACM Press.

He, Z. H. Z., Iyer, R., and Chandler, P. (2006). Vision-based UAV flight control and obstacle avoidance. In *American Control Conference*, pages 2166–2170, Minneapolis. IEEE.

Id Software (1993). Doom.

Irwin, D. E. (2004). Fixation Location and Fixation Duration as Indices of Cognitive Processing. In *The interface of language, vision, and action: Eye movements and the visual world*, chapter 3, pages 105–133. Psychology Press, New York.

Jhala, A. and Young, R. M. (2005). A discourse planning approach to cinematic camera control for narratives in virtual environments. In *AAAI*, number July, pages 307–312, Pittsburgh, Pennsylvania, USA. AAAI Press.

Jhala, A. and Young, R. M. (2010). Cinematic Visual Discourse: Representation, Generation, and Evaluation. *IEEE Transactions on Computational Intelligence and AI in Games*, 2(2):69–81.

Kamiya, H. (2001). Devil May Cry.

Kojima, H. (1998). Metal Gear Solid.

Krizhevsky, A., Sutskever, I., and Hinton, G. E. (2012). Imagenet classification with deep convolutional neural networks. In *Advances in neural information processing systems*, pages 1097–1105, Lake Tahoe, NV, USA. Neural Information Processing Systems Foundation.

Lasseter, J. (1995). Toy Story.

Liapis, A., Yannakakis, G. N., and Togelius, J. (2012). Adapting Models of Visual Aesthetics for Personalized Content Creation. *IEEE Transactions on Computational Intelligence and AI in Games*, 4(3):213–228.

Lino, C., Christie, M., Lamarche, F., Schofield, G., and Olivier, P. (2010). A Real-time Cinematography System for Interactive 3D Environments. In *ACM SIGGRAPH/Eurographics Symposium on Computer Animation*, pages 139–148.

Martinez, H. P., Jhala, A., and Yannakakis, G. N. (2009). Analyzing the impact of camera viewpoint on player psychophysiology. In *International Conference on Affective Computing and Intelligent Interaction and Workshops*, pages 1–6. IEEE.

McDermott, S., Li, J., and Bares, W. H. (2002). Storyboard frame editing for cinematic composition. In *International Conference on Intelligent User Interfaces*, pages 206–207, San Francisco, California, USA. ACM Press.

Meier, L., Tanskanen, P., Heng, L., Lee, G. H., Fraundorfer, F., and Pollefeys, M. (2012). PIXHAWK: A micro aerial vehicle design for autonomous flight using onboard computer vision. *Autonomous Robots*, 33:21–39.

Mikami, S. (1996). Resident Evil.

Olivier, P., Halper, N., Pickering, J., and Luna, P. (1999). Visual Composition as Optimisation. In *Artificial Intelligence and Simulation of Behaviour*.

Picardi, A., Burelli, P., and Yannakakis, G. N. (2011). Modelling virtual camera behaviour through player gaze. In *International Conference on Foundations of Digital Games*, pages 107–114, Bordeaux, France. ACM Press.

Pickering, J. (2002). *Intelligent Camera Planning for Computer Graphics*. PhD thesis, University of York.

Ranon, R. and Urli, T. (2014). Improving the Efficiency of Viewpoint Composition. *IEEE Transactions on Visualization and Computer Graphics*, 2626(c):1–1.

Raynal, F. (1992). Alone in the Dark.

Remedy Entertainment (2001). Max Payne.

Shaker, N., Yannakakis, G. N., and Togelius, J. (2010). Towards automatic personalized content generation for platform games. In *AAAI Conference On Artificial Intelligence In Interactive Digitale Entertainment*.

Simons, R. F., Detenber, B. H., Roedema, T. M., and Reiss, J. E. (1999). Emotion processing in three systems: the medium and the message. *Psychophysiology*, 36(5):619–627.

Soanes, C. and Stevenson, A. (2005). *Oxford Dictionary of English*. Oxford University Press.

Sun, K. and Yu, J. (2007). Video affective content representation and recognition using video affective tree and hidden Markov models. In Paiva, A., Prada, R., and Picard, R., editors, *Affective Computing and Intelligent Interaction*, pages 594–605, Lisbon, Portugal. Springer.

Sundstedt, V., Stavrakis, E., Wimmer, M., and Reinhard, E. (2008). A psychophysical study of fixation behavior in a computer game. In *Symposium on Applied perception in graphics and visualization*, pages 43–50, Los Angeles, California, USA. ACM.

Thawonmas, R., Oda, K., and Shuda, T. (2010). Rule-Based Camerawork Controller for Automatic Comic Generation from Game Log. In *IFIP International Conference on Entertainment Computing*, pages 326–333, Seoul.

Thorson, E. (1994). Using eyes on screen as a measure of attention to television. In Lang, A., editor, *Measuring Psychological Responses To Media Messages*. Routledge.

Togelius, J., Champandard, A. J., Lanzi, P. L., Mateas, M., Paiva, A., Preuss, M., and Stanley, K. O. (2013). Procedural Content Generation : Goals, Challenges and Actionable Steps. In *Artificial and Computational Intelligence in Games*, volume 6, pages 61–75.

Tomlinson, B., Blumberg, B., and Nain, D. (2000). Expressive autonomous cinematography for interactive virtual environments. In *International Conference on Autonomous Agents*, page 317, Barcelona, Spain.

Toyama, K. (1999). Silent Hill.

Ware, C. and Osborne, S. (1990). Exploration and virtual camera control in virtual three dimensional environments. *ACM SIGGRAPH*, 24(2):175–183.

Yannakakis, G. N., Martínez, H. P., and Jhala, A. (2010). Towards affective camera control in games. *User Modelling and User-Adapted Interaction*, 20:313–340.

Yannakakis, G. N. and Togelius, J. (2011). Experience-driven procedural content generation. *IEEE Transactions on Affective Computing*, pages 1–16.