Stories, Agents and Videotapes: Agents that make up stories

(extended abstract of the MSc dissertation)

Joana Almeida
Departamento de Engenharia Informática
Instituto Superior Técnico

Advisor: Professor Ana Paiva

***Abstract*—With our work, we aim to study virtual cameras and how they provide a meaningful way for players to connect to many and diverse types of virtual worlds and games alike. We will explain how we researched virtual cameras in many fields and with many purposes, with the intent of building a system of our own and incorporate these ideas. All this to study the impact and influence the camera has on people's perceptions of these virtual worlds, from event awareness to understanding inter-character relationships.**

Through our results we were able to pinpoint the Third Person Camera as a strong candidate in leading the best possible experience on character-based interactive storytelling systems and games, by providing the best compromise between world awareness and player attachment to the character she controls. We believe this is a start to understanding people's engagement with these systems and a first step to providing them with a better experience that lets them get immersed in the story with greater ease.

I. Introduction

Many storytelling systems and games have graced both the academic and the mainstream worlds at large. These systems allow for players to express themselves while experiencing a story that weighs in their decisions to determine its outcomes through their actions, some more than others. Many times these stories involve other characters that either have Artificial Intelligence, called Agents, that are controlled solely by computer algorithms, or are controlled by other human players. Normally both the players and agents must interact with each other other and/or with items in various ways in order to move the story along.

The controlling of the main character by the player sets the mood for the whole game. A good control scheme will enforce the player's agency, while a weak one might put him off the game altogether. To further enhance this idea of agency, it is very important to keep the user visually focused on what is going on in the game and the events happening within its world. Camera work in games is a focal component to set properly in these cases, because it dictates how the game's story will be experienced.

Given this importance, we aim to study how camera placement can influence how a player perceives and understands a story from within a storytelling system or game while they play it.

II. Related Work

Studying camera in virtual systems is not a new field. It has been applied to many types of systems with different goals in mind. These systems have become increasingly specialized over the years, and require cameras that play well with their strengths. As such, we will present several camera work done over the years. Starting with **3D Virtual Environments**, we explore the core principles and concepts of cameras in general. We then introduce **Interactive Storytelling** and the ways one can develop one such system. This will be important in order both to understand where our work stands within the research community, but also to rationalize some of the camera decisions made in those systems. We then talk about cameras in more specialized virtual storytelling systems: **Games**. After analysing these systems, we will round up some of the more used camera patterns over all these systems and explain how they fit with **User Control**. Finally, we will perform a **Critical Evaluation** over the work we evaluated and see where it can benefit ours.

A. Camera in 3D Virtual Environments

As early as computer systems with graphic capabilities were created, the inherent question of how to best view the visual information these systems showed us followed. Although much of the initial research was initially turned towards how to make the representation more accurate, more and more researchers have started to realize that the way we see the information is just as important. If you don't see the information, or misinterpret it, its purpose is clearly downplayed in the eyes of the user. But this is not the only problem. These systems, as much as they need to show the information, need to also let the user explore the information at their own pace. This means that a new variable emerges that will complicate these camera studies even further: **user control**. Phillips et al. argued that providing good control of the viewpoint is essential to give users a sense of three dimensionality of the objects they are manipulating, in order to manipulate them better[18].

Ware and Osborne introduced three metaphors for camera control and mapped them to real world objects for easier understanding of their concept, as well as providing a tangible interface[31]. Of particular note among them is the **Flying Vehicle Control**, where the user operates a bat that, when moved, will propel the camera in that direction by means of a velocity change, and when rotated,the camera will rotate as well. When tested, it was rated the most useful when trying to navigate a vast 3D scene.

B. Interactive Storytelling

Stories vary wildly in terms of their focus. Sometimes the author will want to highlight the trials of a certain character (along with his feeling towards his predicament, etc.), while on the other hand some stories are better told when framed around other events that precede and follow it (even if the character within it are as important as the story itself), in very specific time frames. These types of storytelling focus respectively on characters and story.

Story-Based systems tend to favour an overall story structure. They clearly define the events that occur within a story, the connection between them and the events that both lead to and follow a specific plot point, providing a better cohesion between all events within the story. On the downside, the story becomes rigid, and user decisions alter the plot only at a very high level. Other researchers have tried to minimize the downsides of story-based systems, by incorporating constraints that help build the narrative as it develops, instead of setting it in stone from the very beginning[21].

Character-Based Interactive Storytelling systems, on the other hand, favour characters and their diversity over overall story structure. It also tends to give greater emphasis on inter-character relationships, and how they affect character's decisions and thoughts. This allows for greater variety in the cast, and makes characters that are much more identifiable and, as an extent, likeable and believable. These characters come across as more intelligent, because their decisions are made without a rigid story structure in mind. As such, they have more freedom over what they can accomplish and how. But this leads to some problems as well. In this type of systems, many events might occur, and while some of them might be good (in terms of pacing, tension, etc.) for the story, others might be completely irrelevant or even undesired. As a result, story structure suffers in the process. Several methods have been implemented in order for this problem to be alleviated, such as heuristic values for story states (preferring certain actions in detriment of others). "Emergence" is a high selling point, as well as a key feature in these systems. Cavazza et al. identified several key aspects for a system to be highly emergent in two of his works: **Initial Character Position, Action Duration and Speed**, **Random Outcomes for Actions**, **User intervention**, **Character-to-Character Interaction** and **Action Failure**[3][4].

Some systems tried to expand on this clear distinction between both previous types by merging them in a way that would be beneficial, combining their strengths. Often, these types of systems will be story-based, meaning their story will be planned in order to have stronger coherence. But these plans will be made with the characters' abilities and personality in mind (thus incorporating a character-based type of planning), reducing the sometimes thoughtless and schizophrenic behaviour that characters tend to have in story-based systems. This also limits the unneeded scenes or redundant interactions recurrent in character-based systems, giving the characters more focus.

Planning, much like the core of progression of a story, can be applied to both Story Structure (defining plot points and their order) as well as Characters (their plans to achieve their goals). Indeed, planning – be it story or character related – always needs to be present no matter the type, and it will be the characters that will make sure it is followed.

The **Hierarchical Task Network** was adapted as a formalism for Interactive Storytelling in many systems due to its appropriateness for use in domains that are very demanding knowledge-wise, a characteristic that Interactive Storytelling clearly has[6][7][26]. A HTN is composed of, as the name would imply, tasks, which can be of various types: primitive, compound , and goal tasks. Problems are solved on a top-to-bottom fashion, decomposing bigger problems into smaller ones until it is simple enough to be solved on its own. It is from this very decomposition that the solution plan is then extracted and executed.

The **Heuristic Search Planner** was initially adopted in order to overcome the main problem of the Hierarchical Task Network of not having a way of incorporating long-distance interactions between actions in a plan. These interactions state that actions in the narrative can influence (positively or negatively) other actions that occur after it (making them easier or harder, respectively)[4][19]. The operators (or actions) used in this kind of planning have, beyond the common pre-conditions and executability conditions, two types of lists: the add-list, one or more new facts and goals given by finishing the operator, and the delete-list, one or more facts and goals that are irrelevant or completed after the operator ends. These are then used to lock/unlock new operators, by providing/removing their pre and executability conditions. The add and delete lists are the main drive for long term dependencies between operators, updating information that may not necessarily be used immediately.

Façade introduced in its work the concept of **beats** and how they could be used for planning[14]. Beats were also adapted from *Façade* and used in a military training scenario application[25] when using Narrative Directive Behaviours.

Porteous and Cavazza introduced their work with planning of a story using constraints when trying to recreate a James Bond Novel, namely *Goldfinger*[21]. Another work that uses this kind of planning is the enactment of the Merchant of Venice, which uses Points of View to delve deeper into a character's motivations and actions[22].

In another kind of planning proposed by Riedl and Young, the core of the plan is given by Author Goals[23].

C. Camera in Interactive Storytelling

Whenever Interactive Storytelling is concerned, camera work needs to account for a very important variable: **User Control**. This is the key difference between Interactive Storytelling and other storytelling media. User control and virtual cinematography stand at odds with each other: by having greater control, the cinematography will become its "servant", and the reverse is true as well.

This also applies to narrative: greater control means that a narrative will have to become less focused in order to account for every possible interaction, what is called the dichotomy between story coherence and user control[20][28][32]. It is a very important balance to keep: the story should not be very derivative nor be very inflexible, thus risking removing the player from the experience by undermining her sense of presence and agency[24].

Also, giving more control to the user might mean she will never see certain events of the story unfold, see but not completely understanding them or even get distracted and miss events that would otherwise be visible[28].

Camera placement in Interactive Storytelling systems needs to account a lot of variables, and the characters are just a part of it. The ambient that they live in can also provide information that might be relevant to the plot, such as lighting, objects and how they are arranged in the set. Tomlinson et al., when developing an Autonomous Cinematographer, identified several elements that make up a camera shot when dealing with a world filled with autonomous characters (a Character-Based Interactive Storytelling system)[29]: the **Actor**, the **Angle**, the **Motion**, the **Transition**, **Occlusion Detection** and **Lighting Elements**.

D. Camera in Games

Games share many similarities with Interactive Storytelling systems, especially the genres that prominently rely on story to engage their audience. One such genre is the Role Playing Game (RPG). However, unlike Interactive Storytelling, games engage their audience with more than story, they also do it through *gameplay*. Gameplay describes the many types of different actions (or *mechanics*) players might do to move the story along (or moving from story point to story point): from combat, to exploration, even puzzle solving. Combined with mechanics that make up story progression, such as character conversations, most games nowadays feature more than one mechanic as they get progressively more complex. From our previous analysis, this means that the camera will need to change based on the current type of interaction taking place.

Much like Interactive Storytelling systems, games want the players to feel connected to their representing character. By putting them at the centre of events, and making them believe they are achieving their goals, they reinforce the player's feeling of agency, and thus motivate them to continue. This connection needs to be made at both the visual and control level: only by working together will the player feel immersed in the experience[13].

Developing skills is a very important factor in gaming experience, and the responsiveness of the input helps in this development. As we can see, this response lies in the visualization of what the input did in the world. The camera is pivotal in this visualization and also sets the mood for the interactions the payer will perform. Games have, until now, presented us with several types of camera control, all of which tend to favour distinct control schemes (more on that in the next section):

- **First Person Camera**: This type of camera places itself on the "eyes" of the character that the player is controlling. It serves to sell the player on the fact that she *is* the character, that the actions the character is performing is indeed the player's own.
- **Fixed Camera Angles**: In this type of camera, the user's visibility is at the whims of the camera, which is out of control of the user.
- **Third Person Camera**: This camera follows the character closely, as if the player was behind him. This type of camera is normally used to establish close proximity, though not as much as the First Person camera[12], but lets the player see the immediate surroundings of her character.
- **Top-Down/Free Roam/Godlike Camera**: In this type of camera, the player views the world and its characters from a high ground. This grants the user a bigger chunk of the world to see at one time, but at the cost of the player's connection to the character he controls[29] and details in the character's looks and expressions[30].

As can be seen, using the correct camera type in a game is very important, as it allows the players to better see the consequences of their actions. In order to keep players from not knowing what to do next or where to look at, level designers employ a multitude of ways to maintain their focus. Milam and El Nasr studied and categorized these patterns with the help of several level designers from several mainstream games[15]: **Collection**, **Path Target**, **Pursue AI**, **Path Movement and Resistance** and **Player is Vulnerable**.

Nnadi et al. also studied the effects of camera positioning on players' ability to read a game's virtual space. By employing different camera parameters and making them consistent with each type of scene present in the game, the players were able to understand the transitions between these scenes faster when compared to a traditional regular third person camera[16].

E. Camera and Types of User Control

There are many different ways for a user to interact with a computer system and the one used very much depends on the type of system that the control is destined for. After our introduction to camera work over various types of systems, it is no surprise that camera placement will directly influence the type of control that is expected from the system. The reverse is also true.

The **First Person Camera** and **Third Person Camera** tend to be very similar in terms of how they control. This is due to the proximity the player has to the character she's controlling, and focus on the character's actions. The most used scheme for this type of visualization is often Mouse + Keyboard or Joystick approach, using keys or an analog stick to direct the character's movement and use keys or other buttons to have the character interact with the world.

The **Fixed Camera** is commonly used in more cinematic experiences, such as Adventure Games. Since the camera does not have a discernible reference point at all times, incorporating a type of character control similar to a First Person or Third Person game tends to confuse the user. Most of these types of games tend to implement a "point and click" approach, where a curso directs the controlled character to where the player wants the character to go or what she wants her to interact with.

The **Top-Down or Godlike Camera** normally features control schemes more focused on gestures and sweeps. The world is being seen from a higher vantage point and, as such, precise control becomes harder to do for a specific character, depending on how far the camera is.

F. Critical Analysis

Visualization of a virtual world has been a constant point of thought and research ever since computer systems became capable of rendering data. From the very first research done on this subject, it was clear that this problem would not be solved by a single approach. In the end, the type of **camera** to use is highly dependent of the system that one tries to make. The problem lies when a system can use a multitude of camera types to view the data. Is one more viable than the other? Will it influence the users' actions in any way?

Most Interactive Storytelling systems and games research focuses on parts of work other than the camera, such as story building or character planning. While these are key components of interactive storytelling, very little work actually devotes itself to seeing how the camera can influence these system's ability to deliver a solid story that users can comprehend.

From the presented research, we have seen how a character can establish a connection to the world and its characters. In the **First Person Camera**, the player trades the ability to position himself with precision in order to increase his proximity to the character (often times *becoming* that character).

Fewer and fewer games use **Fixed Camera Angles** as a recurrent type of camera during gameplay. Its main strength is to provide a more cinematic experience to the players, but at the cost of visibility of the world and positioning, becoming the most restricting camera type.

In the **Third Person Camera**, the character is visible to the player, letting her see what the character looks like and does in greater detail. This lets the player connect better to the actions her character performs on the world, at the cost of distancing both from each other. However, it also lets the player see more of the world at once, and allows the player to analyse the character's surroundings faster when compared to the First Person Camera.

The **Top-Down or Godlike Camera** is more commonly used in games or simulations where the player either does not have a physical character to control in the world or he has to control/maintain more than one character. This camera is the best to use when trying to view the world as a whole, but ineffective to generate player attachment to it as well as the other types of cameras do. The characters will lack detail and expressivity, since the player will see them from far away.

Player control is a very important aspect when discussing the camera: not every type of control will work with a certain type of camera. All types of user controlled cameras we have discussed (save for the Fixed Camera) tend to favour a **Flying Vehicle** type of control. Depending on the type of camera however, certain parameters will be defaulted or severely limited in its value:

- The **First Person Camera** sticks its position with the character, and follows the character's movement.
- The **Third Person Camera** can be located either very closely to the character (by her shoulder) or far enough to see the character's whole body.
- The **Top-Down or Godlike Camera** is usually the one with most freedom. Depending on the game, it can zoom in and out of the world, move around in all directions and see from a wide range of angles.

Most of these cameras support a wide array of controllers. Usually, the keyboard and mouse combination is by far the most used, but some games can also go by with just one or the other at the cost of control of either the camera or the character, or make controls context sensitive. The use of motion controls is still very jarring to the user, but can achieve reasonable results with games that feature simpler controls. These normally encompass Top-Down or Godlike cameras, where the gestures are more broad and sweeping and thus are easier to recognise.

When developing a **Interactive Storytelling** system, it is useful to analyse what the focus of the work is. In this regard, Story-Centric and Character-Centric approaches are best for respectively representing a coherent plot structure or a strong characterization of the characters that populate the world. Mixed approaches are also more plot-focused than Character-based approaches, not having an equally high degree of generation in comparison. **Planning** depends greatly on the core of progression. When going for a more story-based approach, a more structured planning approach is normally used, like the HTN. Constraints are also a possibility, and can allow for a greater variety in story decomposition. When doing a Character-based approach, a more flexible planning tactic is normally a better option. For this purpose, HSP is widely used, since it allows a greater generation of rational actions by incorporating long-distance relations between them.

III. The Game and Camera Model

In order to establish the base of our work, we need to review what we set out to do in he first place. With our problem, we introduced our initial questions that led to the beginning of our research:

> *Given a camera mode, will the player be able to perceive the events of the story? Will he understand the overall arc of the story? The relationship between the characters? Can he understand their motivations and goals?*
>
> *Can the camera affect the way a player experiences a story?*

Just by this small set of questions, we have several needs for our system: a **virtual world**, a **story** that occurs in that world and **Characters** to populate the world. They will help with the story driving process by **Interact** between them based on their **relationships** and **Act** independently, based on their **motivations** and **goals**. Finally, and the key factor in our research, is the **camera**.

Considering that so much emphasis is put on the characters, including their relations and independence, it is inevitable that we should choose a character-based approach to telling our **story**, where our characters will drive the story forward by means of their planning abilities and actions. We decided that using HSP or a close approximation might be the best solution for our system in this case.

Each character will have an **ultimate goal**: what the character will strive to accomplish during the course of the story. This ultimate goal will be in turn composed of **abstract sub-goals**: the moment to moment decisions about what the character should do next, but without commitment to a specific course of action. Abstract sub-goals, once triggered, are refined into **concrete sub-goals**, where the character decides and commits to a specific idea that will satisfy the abstract goal. Concrete goals will in turn be completed by chaining **actions** in a logical way, based on their **pre-requisites** and **effects**.

When an action fails to complete for some reason (lack of resources, resource in use, inability to reach target, etc), the concrete sub-goal that action spawned from fails automatically as well. However, this does not mean that the abstract sub-goal nor ultimate goal will fail automatically: it is important to give the character several options to continue, so she can form a new either abstract goal that she can comply with. This will both make the character more believable and intelligent, as well as more flexible.

The **camera** will be our focal point, since it is on it that hinges much of our research. Our idea is to provide a way to evaluate how common camera patterns used all over games and Interactive Storytelling systems fare in giving users the context needed to understand an interactive story.

We identified these types of camera in section II-D as being the **First Person Camera**, the **Third Person Camera** and the **Top-Down or Godlike Camera**. But before delving into the specifics of each camera, it is a good idea to review how all cameras operate. We will continue by explaining the basics of how any camera operates.

The camera is a special object in any virtual system, because it declares parameters and functionality that define much of how a user sees a virtual system. The camera can be set up in any number of ways and provide different looks and feel depending on its configuration. In order to operate, the camera needs to have the following attributes correctly defined:

- **The Projection Type**, chosen from two possible types: **Perspective**, the most commonly used in 3D applications, since it provides a more realistic visualization for the viewer and **Orthographic**,which makes it hard to traverse a 3D world since the depth of the world is lost.
- **Clipping Planes**: The near and far planes represent respectively the minimum and maximum distance that the viewing volume will contain. Anything behind the near plane and in front of the far plane will not be drawn.

For the camera function, it also needs a variety of functions that help with its management, including one of the most important: the way to position and rotate itself so it focuses on a target, by applying a *LookAt* Matrix[27].

The **First Person Camera** will be intimately connected to the main character, by sharing its position with the position of the character's eyes, and looking to the character's front. Despite this type of camera being a proxy for the character's (and the player's) head, tilting (rotation along the z axis) is often disabled in order to reduce controls and confusion for the user.

The **Third Person Camera** is also connected to the main character, but with less emphasis than the First Person Camera. This camera hangs a certain distance from the character, far away enough to view her but close enough to still maintain a personal view and perspective. Also, this type of camera usually does not sit exactly behind the character, but instead looks over her shoulder in order to let the player view a bigger portion of what is in font of the character. To place this camera, there are more calculations required than in the First Person Camera. The character's front and right vectors are used to determine the point from where the camera should look over the character's shoulder by calculating three vectors:the **horizontal distance vector**, the **vertical distance vector** and the **side distance vector**. These are all summed up to determine the point where the camera should be placed. Also, we define an extra 3D point with the help of the character's forward vector and multiplying it by a distance to get our *LookAt* point, which will become the center of our camera perspective.

The **Godlike or Top-Down Camera**, unlike the previous two types, has little to no attachment to any characters in the game whatsoever. In this type of camera, its position and rotation provide much more freedom than any other type. The player can move it around (sometimes in all directions), zoom in and out to get a more detailed or broad view, even rotate the camera around to see the same scene from different angles. Some of these can still be constrained (such as preventing the camera from moving outside the world's borders or rotate the camera to view the sky/ceiling), but its lack of focal point and anchor in the world (which was the character in the previous two) allows the player much more freedom to see the scene how he wants and from where he wants to. There are two types of rotation used by these cameras **Rotation in place** or **Orbit**. Although this is not a very used type of camera when a single character is involved, we deemed important to compare its results against the other two most common used types. This, however, means that a different control scheme needed to be developed to play to this camera's strengths and weaknesses.

IV. Implementation

Our game took inspiration out of one of Charles et al.'s earlier work in their Interactive Storytelling research[6][3][8][5][2]. In it, the characters are all autonomous, and the player is a disembodied agent that can influence the world by manipulating its objects and thus mess or help the characters with their plans. In this case, the characters in the set are a group of friends, and one of the characters wants to take another on a date. In order to do that, he needs to know what she likes and woo her with gifts and other things she might like. By interacting with other characters and objects in the environment, he can collect clues that can lead him to success.

In order to make our Game a reality, we decided that using a game development engine was the best approach. To make our game, we used the Unity3D engine/game development tool, already used in several commercial titles[1] and that has proven to be an easy to use and fast prototyping tool.

[1] http://unity3d.com/gallery/made-with-unity/game-list

In Unity3D, a game is an hierarchy of **Scenes**, **Game Objects** and **Components**. The Architecture of Conceal features several types of Game Objects:

- The **Manager** is an invisible entity responsible for organizing and maintaining all other Game Objects.
- The **Camera** houses a *Camera Component* that Unity3D already features and that we extended to make our camera able to change between presets.
- The **Pathfinding Grid** was introduced in order to help our characters traverse through the scenario. It uses the AStar Path library[2], that creates a grid of the scene in real-time and determines paths point-to-point using the A* algorithm that characters can use to get to their objectives.
- The **Characters** are composed of three types of Game Objects: the **Body**, responsible for all of the Physical variables attached to the character, such as its position, orientation and inventory, the **Mind**, housing all the intelligence the character has, including the actions she can perform and variables used in planning, tying her to our Agent Framework FAtiMA, and the **Model**, which houses all resources and variables pertaining to her visual representation.
- **Interactable Items and Placeables** are the main type of objects seen in the world, which characters can interact with in order to complete their plans and advance the story.
- **Decorative items** fill the world to make the scene more recognisable, provide natural ways of restricting the playing space and offer greater visual variety to rooms.

In our project, we needed to ground the player's experience in a character from the world, in order to test how each camera would affect the player's understanding of the situations around that character. To do so, we took advantage of the already existing camera specific components present in Unity3D, but decided to add functionality to them to make it work like the **camera types** we mentioned in our Model.

Unity3D already has a pre-made **Camera Component**, that includes all the basic variables and functionality that a camera needs in order to function properly. After this, we expanded this component with one of our own: a *Script Component*. The script uses the manager to track which type of camera it should use and any other information it needs about the character.

The control scheme used in each camera is the scheme commonly used on most recent games and interactive storytelling systems that feature that type of camera, in order to provide some familiarity. In First and Third Person Camera, by moving the mouse away or towards her, the player can have the character look up or down. By moving the mouse sideways, the character rotates in that direction. Also, the left mouse button is used to interact with close objects that are targeted by the mouse pointer. The character's movement is controlled by the WASD keys or the directional keys.

In the Godlike or Top-Down Camera, since it can be placed and orientated freely, we decided that using a point-and-click style of interaction was better suited in order to keep the player from becoming confused and still be able to perform all actions with a high degree of accuracy. The player can move the camera in all directions using the WASD or arrow keys, zoom in and out using the mouse wheel and rotate it by right-clicking and dragging the mouse.

As we have introduced in our model, our **characters** will be the main focus of our system when it comes to story progression. In this section, we present how we developed these characters in terms of characterization and action: the basis for their intelligence.

In order to focus on building the world, we decided that using an existing architecture for our agents' minds was the best solution. FAtiMA[9][11][10] is a fully mature model for character artificial intelligence with a good track record. It is used in several systems where characters are at the center of the experience, such as FearNot![17] and ORIENT[1], and uses a very similar method for planning and character action that we described in our model. But before going into detail on how our characters reason, we'll introduce how FAtiMA's architecture functions in a general sense.

FAtiMA employs a empty core template that defines how the whole architecture works. This template can then be extended by layers (or components) that add up to form progressively more complex lines of reasoning for the characters. This is done by running each layer in turn, executing its main functions. If new events occurred, they are processed into the character's memory and fed once again to all layers to update them on the new status. Finally, it chooses an action to perform based on its Action Selection components. This process is repeated each update.

The **Deliberative Layer**, like the name implies, is the component that is responsible for the character's reasoning and planning. This was the most important component we used, that helped us integrate planning capabilities for our agents.

FAtiMA does not have a direct definition of an **ultimate goal** like we described in our model, but instead each character (represented by a role in FAtiMA) is imbued with a list of her goals, ordered by the importance that the character gives to their success and failure. **Goals** are defined by a number of properties and events that must occur for it to succeed. If the goal is actually pursuable, then FAtiMA's planner creates a logical sequence of **actions** to accomplish it.

The **ION simulation framework**, as the names implies, simulates virtual environments, allowing the bi-directional contact between the agent's minds and the world properties they need to manipulate. The core aspect of this framework is to be able to keep an eye on concurrency when performing changes to the world.

Integrating of ION in Unity3D is done trough scripting Components. The main component that needs to be present somewhere in the scene (and that we placed in the Manager Game Object) is the IONSimulation. This component is used to issue direct control over the updates carried over in the ION Layer. For characters and other interactable objects, the components are stored in the "Mind" or "Entity" sub-Object of each and represent the various components of the ION Framework: **IONEntity**, **IONProperty** and **IONAction**. Characters need one extra component: the **FAtiMAMind**. This script launches a FAtiMA process and creates a representation of the Character's mind in ION, a RemoteMind, linking them both afterwards and allowing communication between them.

[2]http://www.arongranberg.com/unity/a-pathfinding/

V. Conceal

On top of the systems presented in the last chapter, we devised a world that would let us use all these components to create a story. That world is *Conceal*, where you play one of the characters involved in a wedding proposition: the Bride or the Groom. We have three characters in our story, each serving a different function.

A. Goals and Relationships

Relationships are very important in order to establish what each character will accomplish and how, and these are reinforced by the characters' goals.

The **Bride's** ultimate goal will be to not marry the Groom. In order to do so, she will employ several tactics, which may compliment each-other or be used on their own, in order to discredit her husband-to-be. These tactics are related to Gifts and Prizes and, as can be surmised, are directly tied to the Groom's plans (further defining the Antagonism relation that was introduced earlier). The Bride can steal and hide or destroy them, or exchange gifts with less desirable ones.

The **Groom's** ultimate goal, as we have stated, is to gain favour with the Bride's Father in order for him to accept him as a worthy husband for his daughter. He will employ two main tactics to be successful: giving gifts and boast about his accomplishments through prizes.

The **Father's** ultimate goal will be to and maintain his family's good name, with a secondary goal to increase his prestige. The marriage is the primary source for the increase in prestige he seeks. Marrying her daughter with the Groom, provided he is a worthy husband, will satisfy both his goals.

Because of this, the Bride and Groom share a relation of **Antagonism**. The Bride will try everything in her power to stop or hinder the Groom's plans by intercepting what he is doing in order to please her father and increase his opinion of him. This is a implicit relation, and will not manifest itself directly. Since both the Bride's Father and Groom want the marriage proposal to succeed, they share a relation of indirect **Complicity**. This in turn makes the Bride and her Father at **Tension** with each-other.

B. The Objects

When approaching our game, we wanted the interaction mechanic to remain similar to our inspiring system in the way it engages the player, mainly with the repercussions the taking of objects from the scene might cause to the story's structure. To keep the interactions between our characters a bit more interesting, we decided to include several types of objects that can work in combination with each other to produce different results.

In our story, the Groom brought several objects with him to impress the Bride's father. These items fall into two categories: **Gifts** that the Groom may give the Bride's father or **Prizes** that the Groom collected during his life and can show the Bride's father to boast about his accomplishments. Both these types of items increase the Father's opinion of the Groom when gifted or shown.

However, we devised other types of items that can help the Bride turn the tide to her favour: **Hiders** that can be used to hide **Gifts**, **Destroyers** that can be used to destroy **Prizes** and **Tamperers** that can be used with **Prizes** to alter them. If the Bride tampers with a Prize and drops it at the same spot it was before for the Groom to pick up and show her father, the father will be negatively impressed with him.

By placing these items around the world, we can increase the complexity of both our story and characters. The interactions the characters have through these items serve as both a way to make them interact with the world in general, but also to further establish their relationships, even without having direct interactions between them. For example, by stealing, destroying and tampering with the Groom's belongings, it should become clear that the Bride's attitude towards the Groom is nothing short of negative.

VI. Evaluation

In this research, we set out to investigate the relation between the player and the character she portrays in various games and storytelling systems. We proposed that this connection can be made stronger or weaker depending on the type of camera that the game or system used to ground the player in the world. Going further, we hypothesized that a Third Person Perspective might bring the best balance between world awareness and player immersion in order for player to better understand the story that unfolds in these types of systems and its events.

In order to evaluate our system in full, we needed to have a small set of users to test it before committing to a larger audience. These **preliminary tests** were done to discover and iron out the main problems that might arise when people play or watch the game that we were not aware at first.

Our preliminary tests consisted of letting people play the full game and asking them to comment as they played, while we collected their comments. Also, the game also collected logs that allowed us to later analyse what actions the player did, where they clicked, where they were and how long it took them to perform each action. The test was conducted on six people, ages ranging from 21 to 26 years old, all with some form of experience playing games similar to the ones we based our interactions out of.

From the players' **feedback**, as well as well as checking the logs we realized that our initial house was to little and had too few rooms, which clamped up the objects too much and did not allow for player exploration. Also, the items littered the floor and were hard to spot with cameras at eye level. To correct this, we ended up increasing the size of the house by four times the size of the original, created several new rooms and spread the items throughout them. Also, we created new decorative items to place interactable objects on top of and to keep the new rooms from being too empty.

We also realized that the main character's controls were inadequate. Based on feedback, we ended up fine-tuning much of the control attributes, such as speed of the character, ending up having two distinct speeds whether she's carrying an item or not. The Third Person Camera was the one most extensibility changed, pulling it a bit further from the character, allowing the player to see more around her, and also by increasing the side distance and allow the player to see more over the character's shoulder.

We also received some feedback based on the Camera angles and the player's understanding of what was going on in the game world. One of our Godlike Camera testers ended up confused since so many things were happening at once. This was less of a problem related to our system, and more of a drawback of the camera type itself, which bode well for our research.

Making the corrections we deemed necessary after our first feedback loop, we decided to finally entail with our main batch of feedback to make our **final evaluation** and validate our main research hypothesis:

> *In a storytelling environment, having a camera too close or too far away to the main character is detrimental to understanding what is happening around her. By having the camera focus on the main character but also let the player see her immediate surroundings, the player will be able to place himself and connect better to what her character is experiencing.*

This means that we are placing our focus on the **Third Person Camera** which has the best chance of achieving the balance between world awareness and player attachment to the character.

In this second evaluation, we decided to open up and broaden our audience. Since our focus was not so much the integrity of our controls or camera but rather peoples' perception of the world based on it, we decided to employ a combination of video and a questionnaire about it.

In the video, people can see a normal playthrough of the game from the perspective of the Bride, which goes about and interacts with several objects. We had three versions of the questionnaire, with three different videos based on the camera types we have developed. We made sure that the actions performed by the characters were fairly equivalent and that their results were visible during the video, so people could potentially understand their impact.

The questionnaire itself contained four main batches of questions that touched on the subjects of **what the Bride did**, **what the Groom did**, the **perceived relationship between the characters** and **general information**.

The first part's main goal was to test how easy it was to understand what the main character was doing. The second part was designed to test the ability of the person to infer or understand the consequences of the Bride's actions. Also, we wanted to test how easy it was to perceive what the other characters in the set were doing. In the third part, we wanted to test how the characters' relations transpired depending on the player's vantage point, since we discussed that cameras closer to the character tended to bring the character and the world closer to the player. By general information, we wanted to capture a glimpse of the person's age, gender and gaming ability. With it we were hoping to understand their background and analyse their responses based on their previous experience if the results proved too disparate.

The **First and Second parts**, we measured the people's understanding of what had transpired by calculating the percentage of correct answers.

After determining the percentage of correct responses per person, we then conducted tests to see if this percentage significantly changed based on the type of test done. To do so, we used the Test of Homogeneity of Variances and the one-way ANOVA test, since we had three main cases to test between. Note that all ANOVA tests performed in the present research had a Significance level of $.05$.

In both cases, Levene's test indicates that the assumption of homogeneity of variance has been violated, ($F(2,87) = 34.394, p < .001$ for the Bride and $F(2,87) = 22.775, p < .001$ for the Groom). Transforming the data did not rectify the problem and so F-tests are reported nevertheless.

There was a statistically significant difference between groups when it comes to the percentage of correct responses in the first part of the questionnaire, as determined by one-way ANOVA ($F(2,87) = 6.232, p = .003$). A Games-Howell post-hoc test revealed that the percentage of correct responses was statistically significantly higher in the Third Person Camera opposed to the First Person Camera ($p = .005$) and the Godlike Camera ($p = .012$). There were no statistically significant differences between the First Person and Godlike Camera ($p = .606$).

Much like the first part of the questionnaire, there was a statistically significant difference between groups when it comes to the percentage of correct responses in the second part of the questionnaire, as determined by one-way ANOVA ($F(2,87) = 9.840, p < .001$). Also like the first part of the questionnaire, a Games-Howell post-hoc test revealed that the percentage of correct responses was statistically significantly higher in the Third Person Camera opposed to the First Person Camera ($p = .001$) and the Godlike Camera($p < .001$) There were no statistically significant differences between the First Person and Godlike Camera ($p = .991$).

The **Third part** of the questionnaire dealt with how the characters related and felt towards one another. This was a more subjective measure, but one that we deemed important to analyse nonetheless. We evaluated how the interviewees perceived the relationship between characters by providing a scale from 1 to 5 (*Disagree Completely* to *Completely Agree*) on several questions about their behaviour and internal thoughts.

The average values we obtained pretty much corresponded to the ideas we were trying to convey. The Bride did not like (2.13) nor help the Groom (1.76), due to not wanting this marriage (2.16). The Groom seems interested in her though (4.19), and wants the marriage to happen (4.48). He got in the good graces of her father (3.86) and managed to impress him (4.02). As a result, he's likely to be chosen as his daughters fiancé (3.60).

After gathering these results, we needed to study them to determine if there was any significant change between camera types. Once again, we used both the test of Homogeneity of Variances and the one-way ANOVA test.

Like the first two parts, in half of these questions' results, Levene's test indicates that the assumption of homogeneity of variance has been violated. Transforming the data did not rectify the problem and so F-tests are reported nevertheless.

As can be seen in the ANOVA test table, seven out of eight questions have significant differences to how people perceived the interactions between their characters and their feelings for one another based on the camera type.

For the questions that do not violate the assumption of homogeneity of variance, we applied an additional post hoc test that is more appropriate, the Tukey Test.

The results show that there was:

- No significantly statistical difference between any camera types and the opinion they had on the fact that the Bride Helped the Groom, $F(2, 87) = 1.271, p = .286$.
- A significantly statistical difference between camera types and the opinion they had on the fact that the Bride Likes the Groom, $F(2, 87) = 7.883, p = .001$. Tukey post hoc tests revealed no significant differences except between First Person - Third Person Camera ($p = .031$).
- A significantly statistical difference between camera types and the opinion they had on the fact that the Father Likes the Groom, $F(2, 87) = 6.703, p = .002$. Tukey post hoc tests revealed significant differences between First Person - Third Person Camera ($p = .003$) and Third Person - Godlike Camera ($p = .012$).
- A statistical difference between camera types and the opinion they had on the fact that the Bride wants the Marriage, $F(2, 87) = 3.176, p = .047$. However, Tukey post hoc tests revealed no significant differences between Camera Types.

To analyse the others, we once again recurred to a Games-Howell post-hoc test.

The results show that there was:

- A significantly statistical difference between any camera types and the opinion they had on the fact that the Groom Impressed the Father, $F(2, 87) = 7.883, p = .001$. Games-Howell post hoc tests revealed no significant differences except between First Person - Third Person Camera ($p = .001$) and Third Person - Godlike Camera ($p = .038$).
- A significantly statistical difference between camera types and the opinion they had on the fact that the Groom Likes the Bride, $F(2, 87) = 8.024, p = .001$. Games-Howell post hoc tests revealed no significant differences except between First Person - Third Person Camera ($p = .001$) and Third Person - Godlike Camera ($p = .036$).
- A significantly statistical difference between camera types and the opinion they had on the fact that the Groom wants the Marriage, $F(2, 87) = 5.623, p = .005$. Games-Howell post hoc tests revealed no significant differences except between First Person - Third Person Camera ($p = .008$) and Third Person - Godlike Camera ($p = .013$).
- A statistical difference between camera types and the opinion they had on the fact that the Father will choose the Groom, $F(2, 87) = 7.618, p = .001$. Games-Howell post hoc tests revealed no significant differences except between Third Person - Godlike Camera ($p = .001$).

These values clearly show us that the Third Person Camera stands out in most tests over the First Person Camera, even trumping the Godlike Camera in a few tests as well.

The **Fourth part** of the questionnaire was comprised of several general questions to help us analyse the background of our interviewees. It contained questions about gender, age and playing habits by game genre. These game genres helped us filter the types of cameras people are accustomed to playing with, based on the most common type of camera implemented in them. Here is what we learned about our population:

- They are closely tied together on gender, with males (58.9%) occupying a slightly larger portion than the females.
- They are young overall, with most people occupying the 18 to 25 years tier (78.9%).
- Most are avid gamers, playing at least once a week. The majority either play everyday (34.4%) or once every two days (22.2%).
- The majority played or plays games that usually feature the type of camera that they tested (78.9%).

The last point is especially important to note: using a type of camera regularly might be a good start to figure out how it works and helps in the creation of a more familiar environment, but ultimately it is not a instant defining factor to understand what is going on in the virtual world as events occur.

Concluding from these results, most of our tests show that the Third Person Camera allows for a greater retention of transpired events and heightened insight on character relationships when compared to the First Person Camera. The line blurs a little more when compared to the Godlike Camera: the perception of events or character relationships does not have a discernible difference in most cases. Still, the Third Person Camera faired better against the Godlike Camera where character relations were concerned, proving that the Godlike Camera can distance the characters from the viewer more.

Regardless of the results' statistical significance, the Third Person Camera still gained the upper hand when considering the overall mean of the answers. All in all, the Third Person camera clearly shows the potential to overthrow the other two types when character-controlling character-based interactive storytelling systems are concerned.

This does not mean that people got very disparate takes on the characters, however: in the end everyone was able to find a common ground and share the same opinions overall about character relations. The Third Person Camera stood out merely for garnering more definitive answers.

Taking all of this into account, these tests are not meant to be taken as a definitive result. Measuring a person's story understanding can be quite a daunting task, mainly because there's no clear or assured way to evaluate such an abstract concept. Also, there are shortcomings with the type of test we did that could potentially skew the results. The first and most prominent is that we had no control over the number of times people could watch the video, even thought we discouraged multiple viewings. This could potentially mean that some of the people might have second guesses when it comes to what happened and the relationships between the characters.

This could potentially be solved by having interviewees actually test the prototype and having people supervise and collect information from the player while they play. Although this could be potentially a better type of test to conduct, it presents other challenges: people would invariably experience it at a different pace and commit to different plans, compromising the homogeneity of the tests. This factor was what led us to conduct the tests we have showed, since they allowed us to reach conclusions based on tests under the exact same conditions.

Despite all this, we still believe that the tests and the difference they showed are meaningful, and could be backed up with further and more diverse testing.

VII. Conclusions

With this work, we sought to study how the camera works with a game to provide the player with the information he needs to understand what is going on in the virtual world during play. We started by asking ourselves whether or not the camera actually impacts this exchange of information, and proposed that the camera should be far away from character enough that player sees her immediate surroundings, but also close enough to maintain a connection between them, as our research question showed:

> *In a storytelling environment, having a camera too close or too far away to the main character is detrimental to understanding what is happening around her. By having the camera focus on the main character but also let the player see her immediate surroundings, the player will be able to place himself and connect better to what her character is experiencing.*

This, we postulated, could be the key to making stories in virtual systems more approachable to the common person.

In order to pursue this research, we started by gathering related work from various sources and dealing with many topics. These would lay the foundations for our own system which we would develop.

We started by going over how cameras work in virtual settings in a general sense. This allowed us to put many of the current camera systems into perspective, by learning about the fundamentals that make or break a camera. It provided us with the basic knowledge to build these cameras in our own virtual environment later on. Also, there was a need to study how Interactive storytelling had progressed over the years, and how it was able to bring a world and its characters to life in order to tell a story. We focused especially on the relationship between story and characters in order to decide which path to take when developing our system. With this information, we were able to better analyse how today's cameras work in many and varied virtual systems. We studied many interactive stories and games to provide us with key information about choice of camera and design decisions that either influenced it or were made because of it. Finally, we delved into the relationship between camera and character control by analysing how different types of control could be mapped to certain camera types and vice versa, verifying the pros and cons of each approach.

After all this research, we started building the model for our interactive storytelling system based on the information we collected. We established it would be a character-based storytelling system, which would support the three main cameras used today in games and interactive storytelling systems: the First Person, Third Person and Godlike Cameras. We developed a 3Dimensional virtual world using the Unity3D Game Engine and populated it with characters and items that they could interact with. These interactions made the story flourish.

Done with its development, we finally tested our system by employing a combination of playtesting, used as a preliminary measure, and questionnaires, used as the main means of testing. Evaluating the responses to these questionnaires, we were able to verify our initial claim: indeed a camera like the Third Person can be beneficial to people's understanding of the world, its events and the relationships between the characters that live there, when compared to the other types of cameras. Although this difference was not statistically significant for all the parts of our test, we can still say with confidence that it was the camera which fared better in every tested aspect overall, from understanding the main character's actions to more accurately discerning characters' relationships and feelings towards others.

Even though our results seem promising at first, many variables are still left to be tested. As examples, the following points can provide good starts for **future research** that could potentially improve other areas related to gameplay cameras on interactive storytelling systems and games:

- When building our system, we had to scale down our initial idea of having a lot of characters on screen (somewhere around 5 to 7 people to keep the system relatively stable and simple) to having just the main character plus the two supporting characters that were utterly essential to keep the story going forward. It would be interesting to test the values we proposed when more characters (and more items for them to interact with) join in, especially if those characters can also disrupt the other character's plans as well.
- During development, and following with ideas on our related work, we wanted to have a system where the player could choose the type of camera he wanted on the fly during gameplay. Our initial idea was to study how and when the players would change the type of the camera based on world events. Even though we had this functionality present in our system, we deemed it too complicated to test under such a short time frame. However, it would be interesting to figure out how players use the camera to their advantage and the reasons for changing to a certain type.

This kind of research would prove invaluable to developers of storytelling systems and games to provide better worlds and improve the interactions players need to perform in them. Also, having a set camera, or knowing the places where the probability players will switch to a certain type is big enough could lead to further improvements in level design, making a world's environments more fun to play and explore through and interactions more interesting.

References

[1] Ruth Aylett, Natalie Vannini, Elisabeth Andre, Ana Paiva, Sibylle Enz, and Lynne Hall. But that was in another country: agents and intercultural empathy. In *Proceedings of The 8th International Conference on Autonomous Agents and Multiagent Systems - Volume 1*, AAMAS '09, pages 329–336, Richland, SC, 2009. International Foundation for Autonomous Agents and Multiagent Systems. ISBN 978-0-9817381-6-1. URL `http://dl.acm.org/citation.cfm?id=1558013.1558058`.

[2] Marc Cavazza, Fred Charles, and Steven J Mead. Planning characters' behaviour in interactive storytelling. *The Journal of Visualization and Computer Animation*, 13:121–131, 2002.

[3] Marc Cavazza, Fred Charles, and Steven J. Mead. Emergent situations in interactive storytelling. In *Proceedings of the 2002 ACM symposium on Applied computing*, 2002. ISBN 1-58113-445-2.

[4] Marc Cavazza, Fred Charles, and Steven Mead. Intelligent virtual actors that plan ... to fail. In *Smart Graphics*. Springer Berlin / Heidelberg, 2003.

[5] F. Charles, M. Cavazza, and S.J. Mead. User intervention in virtual interactive storytelling. In *VRIC, Virtual Reality International Conference*, 2001.

[6] F. Charles, M. Cavazza, and S.J. Mead. Generating dynamic storylines through characters' interactions. *International Journal on Intelligent Games and Simulation*, 2002.

[7] Fred Charles and Marc Cavazza. Exploring the scalability of character-based storytelling. In *Proceedings of the Third International Joint Conference on Autonomous Agents and Multiagent Systems - Volume 2*, 2004. ISBN 1-58113-864-4.

[8] Fred Charles, Miguel Lozano, Steven J. Mead, Alicia Fornes Bisquerra, and Marc Cavazza. Planning formalisms and authoring in interactive storytelling. In *Proceedings of the 1 st International Conference on Technologies for Interactive Digital Storytelling and Entertainment*, 2003.

[9] J. Dias and Ana Paiva. Feeling and reasoning: a computational model for emotional characters. In *Proceedings of the 12th Portuguese conference on Progress in Artificial Intelligence*, 2005. ISBN 3-540-30737-0, 978-3-540-30737-2.

[10] J. Dias, I. Leite, M. Kriegel, M. Y.Lim, C. Martinho, S. Mascarenhas, A. Paiva, A. Pereira, R. Prada, and A. Wichert. Architecture development with action-selection mechanism. Unpublished Deliverable 5.3, 2010.

[11] S. Mascarenhas J. Dias and A. Paiva. Fatima modular: Towards an agent architecture with a generic appraisal framework. *Workshop in Standards in Emotion Modeling, Leiden, Netherlands*, 2011.

[12] Tanya Krzywinska. *Hands-On Horror*, pages 12–23. Harmony Wu, 2002.

[13] Bride Mallon. Towards a taxonomy of perceived agency in narrative game-play. *Comput. Entertain.*, 5(4):4:1–4:15, March 2008. ISSN 1544-3574. doi: 10.1145/1324198.1324202. URL http://doi.acm.org/10.1145/1324198.1324202.

[14] Michael Matheas and Andrew Stern. *Writing Façade: A Case Study in Procedural Authorship*, chapter II. Computational Fictions. MIT Press, 2007.

[15] David Milam and Magy Seif El Nasr. Design patterns to guide player movement in 3d games. In *Proceedings of the 5th ACM SIGGRAPH Symposium on Video Games*, Sandbox '10, pages 37–42, New York, NY, USA, 2010. ACM. ISBN 978-1-4503-0097-1. doi: 10.1145/1836135.1836141. URL http://doi.acm.org/10.1145/1836135.1836141.

[16] Ogechi Nnadi, Ute Fischer, Michael Boyce, and Michael Nitsche. Effect of dynamic camera control on spatial reasoning in 3d spaces. In *Proceedings of the 2008 ACM SIGGRAPH symposium on Video games*, Sandbox '08, pages 157–162, New York, NY, USA, 2008. ACM. ISBN 978-1-60558-173-6. doi: 10.1145/1401843.1401873. URL http://doi.acm.org/10.1145/1401843.1401873.

[17] Ana Paiva, Joao Dias, Daniel Sobral, Ruth Aylett, Sarah Woods, Lynne Hall, and Carsten Zoll. Learning by feeling: Evoking empathy with synthetic characters. *Applied Artificial Intelligence*, 19 (3-4):235–266, 2005. doi: doi:10.1080/08839510590910165. URL http://www.ingentaconnect.com/content/tandf/uaai/2005/00000019/F0020003/art00003.

[18] Cary B. Phillips, Norman I. Badler, and John Granieri. Automatic viewing control for 3d direct manipulation. In *Proceedings of the 1992 symposium on Interactive 3D graphics*, I3D '92, pages 71–74, New York, NY, USA, 1992. ACM. ISBN 0-89791-467-8. doi: 10.1145/147156.147167. URL http://doi.acm.org/10.1145/147156.147167.

[19] D. Pizzi and M. Cavazza. Affective storytelling based on characters' feelings. In *AAAI Fall Symposium on Intelligent Narrative Technologies*, 2007.

[20] Andrew Polaine. The flow principle in interactivity. In *Proceedings of the second Australasian conference on Interactive entertainment*, IE 2005, pages 151–158, Sydney, Australia, Australia, 2005. Creativity & Cognition Studios Press. ISBN 0-9751533-2-3. URL http://dl.acm.org/citation.cfm?id=1109180.1109204.

[21] Julie Porteous and Marc Cavazza. Controlling narrative generation with planning trajectories: The role of constraints. In *Interactive Storytelling*. Springer Berlin / Heidelberg, 2009.

[22] Julie Porteous, Marc Cavazza, and Fred Charles. Narrative generation through characters' point of view. In *Proceedings of the 9th International Conference on Autonomous Agents and Multiagent Systems: volume 1 - Volume 1*, 2010. ISBN 978-0-9826571-1-9.

[23] Mark Riedl and R. Michael Young. Character-focused narrative planning for execution in virtual worlds. In *Proceedings of the Second International Conference on Virtual Storytelling*, 2003.

[24] Mark Riedl, C. J. Saretto, and R. Michael Young. Managing interaction between users and agents in a multi-agent storytelling environment. In *Proceedings of the second international joint conference on Autonomous agents and multiagent systems*, AAMAS '03, pages 741–748, New York, NY, USA, 2003. ACM. ISBN 1-58113-683-8. doi: 10.1145/860575.860694. URL http://doi.acm.org/10.1145/860575.860694.

[25] Mark O. Riedl and Andrew Stern. Believable agents and intelligent story adaptation for interactive storytelling. In *3rd International Conference on Technologies for Interactive Digital Storytelling and Entertainment*, 2006.

[26] Mark O. Riedl and Andrew Stern. Failing believably: Toward drama management with autonomous actors in interactive narratives. In *3rd International Conference on Technologies for Interactive Digital Storytelling and Entertainment*, 2006.

[27] Dave Shreiner. *OpenGL Programming Guide: The Official Guide to Learning OpenGL*. Addison Wesley, 7th edition, 2010.

[28] Karl E. Steiner and Jay Tomkins. Narrative event adaptation in virtual environments. In *Proceedings of the 9th international conference on Intelligent user interfaces*, IUI '04, pages 46–53, New York, NY, USA, 2004. ACM. ISBN 1-58113-815-6. doi: 10.1145/964442.964453. URL http://doi.acm.org/10.1145/964442.964453.

[29] Bill Tomlinson, Bruce Blumberg, and Delphine Nain. Expressive autonomous cinematography for interactive virtual environments. In *Proceedings of the fourth international conference on Autonomous agents*, AGENTS '00, pages 317–324, New York, NY, USA, 2000. ACM. ISBN 1-58113-230-1. doi: 10.1145/336595.337513. URL http://doi.acm.org/10.1145/336595.337513.

[30] Jaroslav Švelch. What you can't see is what you don't get: paradigms of game world visualization. In *Proceedings of the 2008 Conference on Future Play: Research, Play, Share*, Future Play '08, pages 212–215, New York, NY, USA, 2008. ACM. ISBN 978-1-60558-218-4. doi: 10.1145/1496984.1497026. URL http://doi.acm.org/10.1145/1496984.1497026.

[31] Colin Ware and Steven Osborne. Exploration and virtual camera control in virtual three dimensional environments. *SIGGRAPH Comput. Graph.*, 24(2): 175–183, February 1990. ISSN 0097-8930. doi: 10.1145/91394.91442. URL http://doi.acm.org/10.1145/91394.91442.

[32] R. Michael Young and Mark Riedl. Towards an architecture for intelligent control of narrative in interactive virtual worlds. In *Proceedings of the 8th international conference on Intelligent user interfaces*, IUI '03, pages 310–312, New York, NY, USA, 2003. ACM. ISBN 1-58113-586-6. doi: 10.1145/604045.604108. URL http://doi.acm.org/10.1145/604045.604108.