

Close Reading Oblivion: Character Believability and Intelligent Personalization in Games

Theresa Jean Tanenbaum
University of California, Irvine
Department of Informatics
ttanen@uci.edu

Jim Bizzocchi
Simon Fraser University
School of Interactive Arts & Technology
jimbiz@sfu.ca

Abstract

This paper investigates issues of character believability and intelligent personalization through a reading of the *Elder Scrolls: Oblivion*. *Oblivion*'s opening sequence simultaneously trains players in the function of the game, and allows them to customize their character class through the choices and actions they take.

Oblivion makes an ambitious attempt at intelligent personalization in the character creation process. Its strategy is to track early gameplay decisions and “stereotype” players into one of 21 possible classes. This approach has two advantages over a less adaptive system. First, it supports the illusion of the gameworld as a real world by embedding the process of character creation within a narrativised gameplay context. Second, the intelligent recommendation system responds to the player's desire to believe that the game “knows” something about her personality. This leads the players to conceptualize the system as an entity with autonomous, human-like knowledge.

This paper considers ways in which *Oblivion* both succeeds and fails at mapping player behaviour to appropriate class assignments. It does so through the analysis of multiple replays of the opening sequence, and the application of two theoretical lenses - character believability and intelligent personalization. The paper documents moments where the dialogue between player and game breaks down, and argues for alternative techniques to customize the play experience within the desires of the player.

Author Keywords

Game design; believable characters; adaptive systems; interactive narrative; close reading

Introduction: Reading Critically In Oblivion

Oblivion (Bethesda Softworks, 2006) is the fourth “anchor” game in a series of Computer Role-Playing Games (CRPGs) known as *The Elder Scrolls*. *Oblivion* was marketed as one of the first “next-gen” games, and was one of the flagship games for the release of Microsoft's X-Box 360 gaming console. Set in the fictional world of Tamriel, *Oblivion* builds on the narrative and gaming traditions established by its predecessors in the *Elder Scrolls Series* which includes *Arena* (Bethesda Softworks, 1994), *Daggerfall* (Bethesda Softworks, 1996) and *Morrowind* (Bethesda Softworks, 2002).

Oblivion is ambitious in almost every dimension of its design, from the opening training and character creation sequence to the vast world geography and sophisticated artificial intelligence system that governs the behaviours of the world's inhabitants. *Oblivion* spans about 16 square miles of virtual terrain according to a press release from Microsoft (Chihido, 2008) [Figure 1-1].

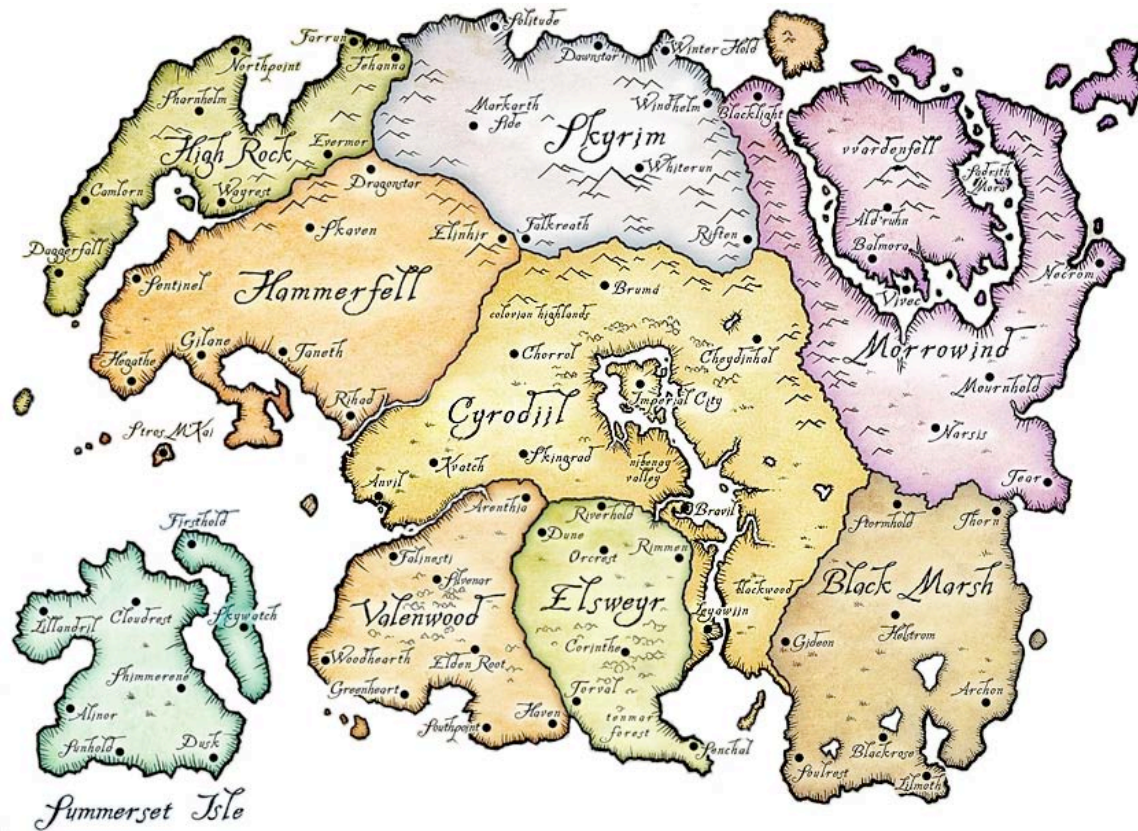


Figure 1-1 - Map of Tamriel (2006, Bethesda Softworks, Used with permission)

It is known for its rich and highly responsive world, a compelling central story, and a myriad of possible play experiences. Murray(1997) maintains that a combination of large scale and extensive detail can support immersion within digital environments. *Oblivion* meets this criterion, using scale and density to help combine aspects of a rails game with aspects of a sandbox game. However, the size and complexity of *Oblivion* are both a blessing and a curse for games theorists. On the one hand, it is a rich experience, rife with material that begs to be analysed. On the other hand, the scope of the game can be daunting; its open-ended nature makes it difficult to focus an analysis on any one facet of the experience.

Oblivion has already been the subject of several other studies. Most notably, Champion (2007) discusses the notion of presence in *Oblivion*, concluding that the game succeeds in conveying a sense of worldliness, but fails to communicate social and cultural presence. Lankoski & Björk (2007) perform a case study in *Oblivion*, considering how one of the NPC shopkeepers, Claudette Perrick, can be understood to partially embody a series of “design patterns” for believability within a game. Rozak (2006) discusses how *Oblivion* strikes a balance between

procedurally generated content and hand-authored content, while also considering the design implications of *Oblivion*'s dynamic difficulty system.



Figure 1-2 - The player's first view of Cyrodill in *Oblivion* (Screenshot by authors, used with permission)

Our investigation takes the form of a close reading, drawing on deconstructionist and hermeneutic traditions to strip away layers of the experience of playing *Oblivion*. Close reading is an epistemological stance that treats knowledge as arising from the act of reading.

“The bad reader...is the one who rushes with indecent, even journalistic haste, to decision, to decide on a reading, and thereby have done with reading, once and for all. Bearing this in mind, and seeking all the while to avoid becoming the bad reader, to have the last word or to close the book on reading, how do we read so as to avoid having read? How do we learn to read patiently, rigorously, in such a manner that we know all the while that we have not yet read, we have not yet done (with) reading...all we can do is practise acts of strong reading which will be, inevitably, misreading.”

(Wolfe, 2000, p. ix)

At the heart of the close reading methodology is the notion of reading to make a momentary meaning, and then of reading again to create a new meaning, and then of reading again, to make another meaning, and so on. This process is an iterative cycle that can not, and should not, be seen as completed or closed. This *hermeneutic circle* denies the possibility of reaching a final “true” reading, which would be counter to the act of reading. Reading is positioned as both present tense and continuous; to say something has been “read” is to suggest that it has been consumed and that the possibility of meaning creation has been exhausted. As close readers of

Oblivion we can thus only present a *strong reading* of the game, but one which is implicitly open for reinterpretation.

In this paper we report on the results of a close reading, which we have performed through iterative playings of *Oblivion*. The data we gathered over the course of this reading takes the form of a densely annotated and structured interaction record for specific sections of the game, and a set of unstructured general observations which were recorded through more casual play. [Figure 1-3]

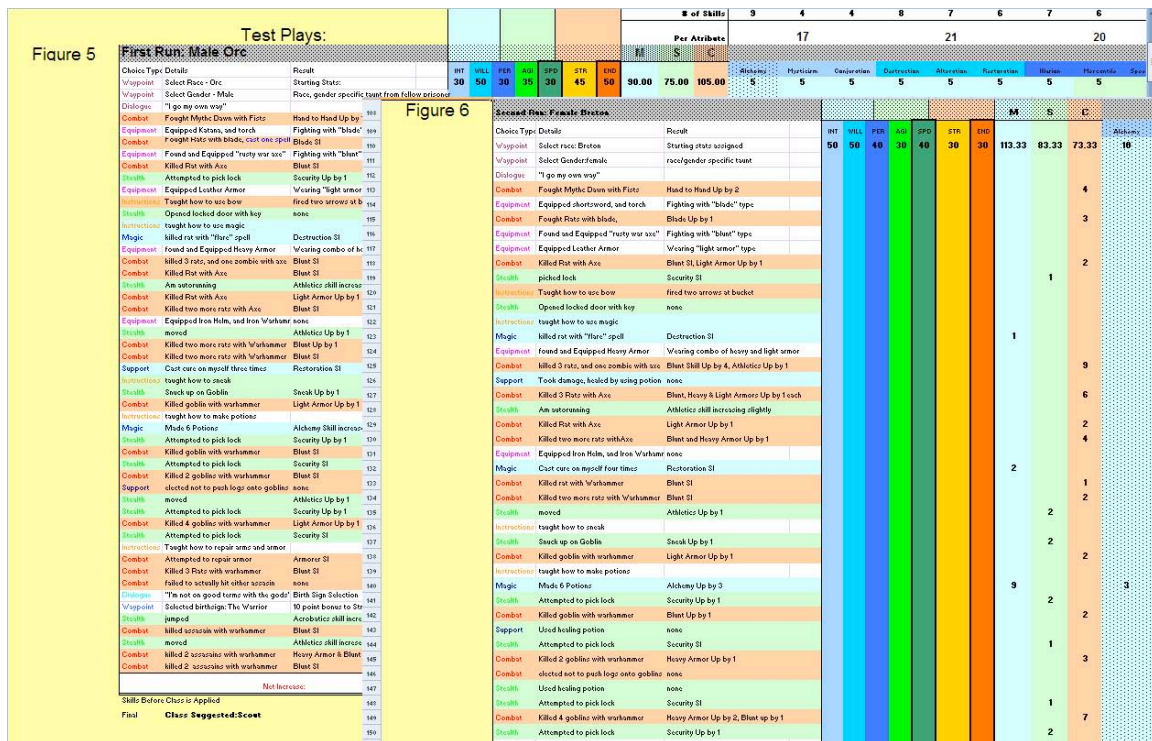


Figure 1-3 - A sample of the interaction record data gathered

Our intent with this paper is to present some of the high-level recommendations that emerged as a result of this directed play. In particular, we are interested in exploring the relationship between *adaptive personalization* and *narrative believability*.

Two Lenses

Close reading is often practiced in conjunction with a critical lens or perspective. One of the problems facing games research arises from the “slipperiness” of the artifacts around which the domain is centred. Games suffer from a certain degree of indeterminacy: one cannot guarantee that two theorists will encounter the same media assets while interacting with a game, or that they will experience them in the same order. Nor can one guarantee that they will observe and attend to the same details of the experience. As if this were not already problematic enough, the very size of many games often defies rigorous explication, with some CRPGs requiring upwards of 400 hours of gameplay time to traverse from beginning to end.

These twin issues of *indeterminacy* and *scope* require scholars performing close readings to carefully explicate their own experiential stance in relationship to the artifact. To accomplish this, we construct *analytical lenses*, which serve to filter our observations in productive ways. In this paper we describe two different lenses that have been used to guide our observations: *Adaptivity* and *Believability*. The use of multiple distinct analytical perspectives in our close reading allows us to isolate and compare specific phenomena within the larger experience of *Oblivion*. By focusing our readings on specific observable (and reproducible) phenomena in the game, we are able to at least partially mitigate issues of *indeterminacy* and *scope*.

We hope this reading teases out an understanding of the relationship between the design of the adaptive system and believability of the game's characters. For epistemological reasons we have attempted to isolate our lenses from each other. In practice, we shall demonstrate that there is much overlap between these two domains. In sections 4.1 and 4.2 we consider each of these lenses in turn, before concluding with a discussion of their interrelationship.

Adaptivity

Our first lens considers *Oblivion* from the perspective of adaptivity. This is a term that has seen much discussion in the fields of hypermedia and user modeling, but which has received little formal attention in games and interactive narrative. Here we consider some of the existing work on adaptivity and discuss how these ideas fit in with interactive narrative research and game studies. The vocabulary and the discourse surrounding adaptivity offer researchers a useful critical perspective for understanding and designing experiences. Our intent with this lens is to draw on the specific terminology and techniques from Adaptive Systems research in Artificial Intelligence in order to explicate *Oblivion's* use of user modeling and recommendation. Specifically, we are interested in how to understand the mechanism by which *Oblivion* executes its form of *adaptive personalization*.

Two of the primary uses for adaptive systems have been to provide *personalized learning assistance* for novice users and to *personalize the user interface* to the needs of advanced users. With the rise of more rich multimedia technologies, especially web-based, has come a third area of research in adaptive systems: adaptive hypermedia. Brusilovsky (2001), one of the central theorists of adaptive hypermedia, writes:

“Adaptive hypermedia is an alternative to the traditional ‘one-size-fits-all’ approach in the development of hypermedia systems. Adaptive hypermedia systems build a model of the goals, preferences and knowledge of each individual user, and use this model throughout the interaction with the user, in order to adapt to the needs of that user.”

(Brusilovsky, 2001, p. 87)

Adaptive systems operate on the notion that each user has a personal set of tastes and preferences that can be directly catered to. In order to best serve these preferences, adaptive systems need to be able to gauge the desires of their users.

Oppermann (1994) identifies the three components of adaptive systems using the metaphor of a biological organism: an *afferential component*, an *inferential component*, and an *effferential component*.

Afferential Component of Adaptivity. Adaptive systems observe and record user behaviour and system reactions, in addition to directly executing commands according to their input...

Inferential Component of Adaptivity. Adaptive systems analyze the gathered data in order to draw conclusions...The inferential component decides in which way the system should modify its behaviour in order to correspond to the actual usage profile. This implies that a basis must be specified (a theory, a set of rules) for drawing inferences. This also implies that the kind of data to be recorded (afferential component) and how the system should be adapted (effferential component) must be defined...

Effferential Component of Adaptivity. Adaptations lead to modifications of the system's behaviour. The change may concern presentation of objects, functions, or tools; default values for parameters; sequences of dialogues, or system messages."

(Oppermann, 1994, pp. 6-7)

These three components might be more simply described as *Input*, *Reasoning*, and *Output* respectively. This framework provides a good overview of what an adaptive system does: it gathers information about the user, it reasons about this information, and uses the results to modify the system in some way that is intended to improve the user's experience.

It is in the *inferential* component that most of the intelligence of an adaptive system is located. Many adaptive systems rely on a specific technique from AI known as *user modeling*. User models can be simply defined as "the models that systems have of users that reside inside a computational environment" (Fischer, 2001, p. 70). Alfred Kobsa (1995) describes two frequently used techniques for creating models of users: *stereotypes* and *interaction history*. The first technique, *stereotyping*, was initially devised in 1979 by Elaine Rich, and has remained relatively unchanged. Rich proposed that systems could use "clusters" of features to group users into categories. If the user triggered or activated enough features over a threshold, then she would be classified as belonging to the specific stereotype that those features were associated with. Rich argued that stereotyping was used by people in ordinary communication to quickly build models of each other, but observed that humans became emotionally attached to their stereotypes, often ignoring confounding information after the stereotype had been established. She claimed that computational systems would not suffer from this problem, due to the objective algorithmic nature of the evaluation (Rich, 1979).

Kobsa's second common technique for user modeling is *interaction history*, which has been more recently described by Brusilovsky & Millan as *feature-based modeling*. Feature-based modeling attempts to model specific features of individual users such as knowledge, interests, goals, etc. During the user's work with the system, these features may change, so the goal of

feature-based models is to track and represent an up-to-date state for modelled features (Brusilovsky & Millan, 2007).

Brusilovsky & Millan (2007) lay out a taxonomy of the features of a user which can be usefully modelled within adaptive systems including: Knowledge, Interests, Goals and Tasks, Background, Individual Traits (such as cognitive styles and learning styles), and the Context of the user's work. The advantage that feature-based modeling has over stereotyping is that it does not rely on a large pre-existing knowledge base of possible user types, and it does not risk miscategorising users into inappropriate containers should they exhibit behaviours that fall outside the set of expected user types. However, feature-based modeling is only as good as the inferential algorithm upon which the feature assignment is built. In one sense, the distinction between stereotype based modeling and feature-based modeling is one of granularity. Stereotype based models are comparatively low granularity in terms of what is measured and how the system can respond, whereas feature-based models allow for a much finer-grain measurement and response.

In games, the notion of adaptivity has great appeal. A number of commercially successful games already demonstrate some degree of adaptivity, often by changing the game environment to reflect the player's preferences, as in portions of *Black & White* (Lionhead Studios, 2001), or by making different abilities available to the player based on their in-game choices, as in *Star Wars, Knights of the Old Republic* (BioWare, 2003). Using the terminology of adaptive systems, we can evaluate games in terms of their *afferential*, *inferential*, and *afferential* components, as well as in terms of the user modeling technique employed. We can also begin to evaluate the extent to which these systems succeed at "inferring user preference" via designed interactions.

Believability

Believability is a notion that has been explored in several disciplines. In the humanities, believability is often associated with narrative quality: a good character is a believable one, and a good story is one that "rings true". Unbelievable characters and situations detract from the enjoyment of the narrative. However, there is little agreement on what formal traits make a character or a situation believable; often believability is a subjective and a highly personal metric, rather than a formal one. In both animation and artificial intelligence research, believability is mainly associated with "believable characters" or "intelligent agents". In these cases, believability is often framed as a issue of "seeming human" to a human observer or of creating "lifelike" behaviour.

Believability in games is a function of both of these perspectives. It is necessary for characters to maintain an illusion of life, while also adhering to the more subjective, aesthetic notion of character believability from the humanities. While this is still a fuzzy and difficult area to attempt to quantify, there are aspects of a character's design and behaviour which can be evaluated as signifiers of dramatic believability. Mateas provides the following definition of believability:

"For many people, the phrase believable agent conjures up some notion of an agent that tells the truth, or an agent you can trust. But this is not what is meant at all. Believable is a term coming from the character arts. A believable character is one

who seems lifelike, whose actions make sense, who allows you to suspend disbelief. This is not the same thing as realism. For example, Bugs Bunny is a believable character, but not a realistic character.”

(Mateas, 1997, pp. 5-6)

If believability is not always about making a character “realistic”, then solving the believability problem is not just about making characters act more like real people. Certainly, within novels and films there are many examples of characters that are believable within the context of the narrative, even when their actions and circumstances are a far cry from any actions or circumstances which we normally encounter in the “real world”. For example, dialog within much film and theatre follows a set of narrative conventions that make it “unrealistic” in comparison to an actual conversation between people, with its fits and starts and overlaps. So how do we perceive and experience these things as “believable” even though they do not conform to our perception and experience of the world?

In order to look at believability as an *experienced* phenomenon rather than an *encoded* phenomenon, we must look at work that draws upon human psychology and cognition. Magerko (2007) suggests that in order to understand and evaluate believability, it is necessary to first understand it as a phenomenon that is *experienced* or *perceived* by an observer, rather than a set of behaviours performed by a character.

“It has proven difficult to clearly define what this term ‘believability’ truly means when discussing (and more importantly comparing and contrasting) character behaviours. ‘Believability’ too often refers to the specific kinds of behaviours a particular agent has rather than a more general metric that is extrinsic of a particular approach and focuses on an observer’s perception of that behavior.”

(Magerko, 2007, p. 79)

Magerko goes on to describe a different approach to measuring dramatic believability and proposes that believability be deconstructed into two different components: “a) the user’s expectation of a performance and, b) the fulfilment of that expectation” (Magerko, 2007, p. 79). He separates the user’s expectations into two distinct types: *internal expectations*, which are the promises made by the fictional world, and *external expectations*, which arise out of the world knowledge that the reader brings to the experience (Magerko, 2007).

Magerko is not alone in describing this phenomenon; in 2006 Pasquinelli presented a poster at the *Enactive* conference entitled *The role of expectations in the believability of mediated interactions*. In this paper, Pasquinelli defines believability in the following way:

“The notion of believability in mediated conditions can be characterized as a judgment regarding the plausibility of a certain mediated experience, the judgment being positive when the experience respects the expectations of the subject which are activated by the contents and context of the experience itself.”

(Pasquinelli, 2006, p. 201)

This idea of the context and the contents of the experience “activating the expectations of the subject” is similar to Magerko’s concept of *internal* versus *external expectations*. Pasquinelli describes the mirror of this phenomenon in terms of “suspension of disbelief” when she discusses how mediated experiences *intentionally deactivate user expectations* in order to create more believable experiences (Pasquinelli, 2006).

Psychology has long acknowledged the salience of human expectation, not only with regard to our capacity to ‘believe’ in mediated experiences, but also to our overall perception of the world. In 1949 Bruner and Postman published an article entitled *On the Perception of Incongruity: A Paradigm*, in which they write:

“The organism in perception is in one way or another in a state of expectancy about the environment. It is a truism worth repeating that the perceptual effect of a stimulus is necessarily dependent upon the set or expectancy of the organism. And so, in many situations the student of perception must also specify the expectancies of the organism when exposed to stimulation.”

(Bruner & Postman, 1949, p. 206)

In other words, the expectant state of a person has the capacity to inform and alter the perception of phenomena in the environment. Our perceptions are at some level constituted by the attitudes and mindset. In psychological testing, one of the ways this is controlled for is by carefully monitoring and regulating the amount of knowledge that is given to study participants, in order to account for the impact of their expectations and assumptions on their experience.

This phenomenon provides at least a partial accounting for the relationship between mediated belief and expectations. To further explicate this relationship we turn to the work of Marsella and Gratch (2003), who have done extensive work on the creation of virtual humans for simulations and games. They propose two related psychological mechanisms that help “mediate between stimuli and response” in humans (Marsella & Gratch, 2003, p. 314). The first is the process of *appraisal*, which is the “process by which a person assesses their overall relationship with its environment, including not only their current condition but past events that led to this state as well as future prospects” (Marsella & Gratch, 2003, p. 314). The second is the process of *coping*, which “determines how one responds to the appraised significance of events” (Marsella & Gratch, 2003, p. 314). While Marsella and Gratch are primarily interested in modeling these cognitive processes in order to simulate them in virtual characters, they are rooted in actual models of human cognition that bear a striking resemblance to the treatment of expectations discussed above. We contend that it is possible to understand the process of *appraisal* and *coping* as one in which the individual initially *formulates a set of expectations* and then *responds (copes) based upon those expectations*. This coincides with the example given by Marsella and Gratch:

“People are motivated to respond to events differently depending on how they are appraised. For example, events appraised as undesirable but controllable motivate people to develop and execute plans to reverse these circumstances. On the other

hand, events appraised as uncontrollable lead people towards escapism or resignation.”

(Marsella & Gratch, 2003, p. 314)

In other words, the *beliefs* people have about a situation will influence their response to the situation. What all of these examples point toward is an understanding of the phenomenon of believability as one which is deeply grounded in the context of an encounter with a media object: in order to understand what a reader is going to find believable, it is necessary to take into account how she has *appraised* the situation, what her expectations are, and the extent to which the experience is able to *deactivate* unwanted expectations while *activating* desired expectations. This relationship between believability, expectations, appraisal, and response bears a close resemblance to what Masahiro Mori (1970) described as the *Uncanny Valley* [Figure 2-1].

Initially described by Mori in 1970, the Uncanny Valley was originally framed to describe a phenomenon in robotics, but has since proven useful to any situation where it is desirable to represent a human figure. Mori describes a “valley of familiarity” which occurs when a robot or toy grows close enough in appearance to a living healthy person that the elements that fall short of humanness result in an uncanny ‘negative familiarity’. One example he gives is of a prosthetic hand which visually appears very realistic, but which reveals its mechanical nature when shook, creating an unsettling inversion of familiarity (Mori, 1970).




Figure 2-1 - Mori's Uncanny Valley (Mori, 1970)

As a character comes closer and closer to a real living human in appearance and behaviour, it elevates our expectations. Rather than viewing it as an automaton, or an animated figure, we find ourselves wanting to view it as a human. The resulting dissonance between the almost-

human character and our expectations of a real human results in a *less-believable* experience of the character; we have fallen into the Uncanny Valley. One question that remains unanswered is whether this same idea can be applied to the believability of a gameworld or to specific interactions within a game. As the representation of reality becomes more and more realistic, the more it invites comparison to our actual experience of reality.

Descent Into Oblivion

These two perspectives, *adaptivity* and *believability*, provide us with the means to consider the design of *Oblivion* through two different lenses. As discussed above, *Oblivion* is a rich and sprawling game with hundreds of hours of content to explore. In order to constrain and focus our investigation, we look at a relatively small portion of the game: the opening sequence and character creation process. This section of play is of great interest because it trains the player in the basic game mechanics, it introduces her to the central narrative of the game, and it provides an extended *narrativised character creation* mechanism.

Narrativised Character Creation

One of the most interesting design elements that the *Elder Scrolls* games share is a narrativised character creation process.¹ In *Arena* and *Daggerfall*, this took the form of a questionnaire. Players were presented with ethical dilemmas, or problems to be solved, and a set of multiple choice responses intended to separate them into three broad categories: fighter, mage, or thief. In *Morrowind*, character creation also follows a questionnaire model, but it is integrated into the narrative context of the game: the player awakens on a prison transport ship, and is interrogated by a jailor about her past before being released back into society, thus establishing relevant details about the character. In *Oblivion*, the character creation process takes the form of an “introductory dungeon” where the player is slowly trained in the basic mechanics of the game. As she explores this dungeon she is given opportunities to use different play mechanics, to outfit her character, and to select aspects of her character’s background such as race, gender, and birthsign.

Intelligent Personalization

Associated with the character creation process is a unique and ambitious device: an adaptive system which observes the player as she explores the introductory dungeon and attempts to model her play style. At the end of this sequence, the game recommends a class to the player that it has derived from her behaviour in the game. In this way, the game attempts to reason about the player’s preferences based on her actions, and to place her into one of 21 possible stereotypes [Figure 3-1]. These stereotypes, or classes, are represented to the system as a series of “class skills” and attributes. For example, the Warrior class has a number of combat skills, such as Blade, Heavy Armour and Hand-to-Hand, while the Mage class specializes in Conjuration, Destruction, and Restoration magic. This opening sequence is a self contained, concentrated play experience, with one specific adaptive task, and a number of clearly observable user choices. As such, it lends itself readily to a study of adaptive system design in games, and of user modeling as applied to a player’s actions and choices.




Figure 3-1 - The 21 Classes in Oblivion

Oblivion: Oblivious to the Player's Preferences

In order to understand the operations of the adaptive system from a player's perspective, we engaged in a series of systematic playings of the *Oblivion* opening sequence. These sessions were carefully annotated in order to codify our understanding of the *observable* mechanisms that informed the system's creation of a player model. We simultaneously paid attention to the extent to which the game's NPCs behaved in narratively believable ways. Our annotations captured details about character-creation decisions such as race, gender, and astrological sign, and gameplay decisions including overall play-style preference, and fine grain measures of specific skill and ability increases. For each play-through we constructed a personal narrative for our character. This narrative guided the expectations that we brought to the experience, while also informing the type of play preference which we attempted to communicate to the system. While the volume of data gathered defies inclusion in this paper, a few high-level conclusions can be derived from the play experiences that it represents.

Reading for Adaptivity: Mis-mapping Player Decisions to System Responses

The first set of issues we identify in *Oblivion* is concerned with the functionality of the adaptive system that underlies the game's class recommendation mechanism. What we found through our observations is that the *inferential* aspect of the system does not effectively bridge between the *effential* and the *afferential* components. There are three specific aspects of the design that we find particularly problematic.

Problem 1: Choices that seem significant to the player are not significant to the system.

One of the most readily apparent opportunities for a player to express a preference to the system about her preferred style of play occurs in the first few moments of the character creation process. Here, the player is allowed to select her race and her gender from a diverse pool of humanoid species. Races in *Oblivion* are visually and narratively differentiated from each other, while also being strategically distinct. For example, the mystical “Bretons” are clearly intended to be magic users, receiving racial bonuses to the Intelligence and Willpower attributes that govern successful spell casting, as well as increases to several of the available schools of magic, such as Conjuration and Illusion. By contrast, the hulking “Orcs” benefit from racial bonuses that include physical attribute enhancements such as Endurance (the ability to withstand greater amounts of damage) or combat skills such as Heavy Armor and Hand-to-Hand. By selecting a race, a player is making two related decisions--the first of which is how he/she wishes to be narratively represented in the gameworld, the second is how he/she intends to play the game [Figure 4-1].



Figure 4-1 - The race selection screen (Screenshot by authors, used with permission)

One of the most surprising discoveries we made in our playings of the opening sequence was that the game engine disregards this choice entirely in its class recommendation system. For our first play-through we chose to create an Orc character, who we correctly predicted would be stereotyped into a combat class: the Scout. In order to determine the significance of the choice of race, we performed a second play-through, this time selecting a Breton character. In order to reduce confounding factors, we attempted to play the Breton in exactly the same way we had played our Orc. Our expectation was that rather than being assigned a combat class, we would be assigned a combat-heavy magic-using class, such as the Spellsworn. Instead, we were assigned the exact same class as before: our choice of race had no bearing on the recommendation system, even though it had significant impact on our experience of the opening

sequence. The fact that the class assignation did not vary across sessions indicated that it was non-random; however, it was difficult to determine which elements of our play were resulting in a particular recommendation. Most problematic was the fact that it seemed possible for the player to make choices which were personally meaningful and significant, but which had no bearing on the response of the system.

Problem 2: The available choices in the opening bias the system against certain playstyles.

For our next set of play-throughs, we focused on the range of choices provided by the system during the opening sequence. These choices constitute the palette by which a player may express her preferences to the system through taking different actions. Does she use a bow or a sword? Does she cast spells or sneak up on enemies? Does she run into a room swinging an axe or does she pick the pockets of her foes and ‘stealth’ away, leaving them alive? *Oblivion* divides gameplay into three categories: *Combat*, *Magic*, and *Stealth*. Within each of these three categories, there are seven skills, and thus seven *imaginable* ways for the player to represent preference through action along each of these valences.

However, within the opening sequence of *Oblivion* is an imbalance between the number of *imaginable* skill increases and the number of *possible* skill increases. In the game opening, the *possible* skill increases are limited by the spells and items provided to the player. While there is an opportunity to pick up and use shields, two different kinds of armor, and a number of different weapons including a bow, sword, and warhammer, there are only two spells provided to the player at the beginning of the game. As a result, the player is capable of ‘training-up’ seven out of seven combat skills, five out of seven stealth skills, and only three out of seven magic skills. It is consequentially much easier to represent a preference for combat or stealth than it is to express an interest in magic during the opening sequence. Should a player wish to play a pure magic user within the opening sequence, the limited opportunities to practice and utilise magical abilities prevents her from clearly communicating this desire to the system.

Problem 3: The way that the system tracks different types of skills biases the game in favour of certain play-styles.

The second bias within the opening sequence is connected to two specific skills: Athletics and Acrobatics. Unlike other skills in the game that can be intentionally and conditionally achieved by the player (i.e. by repeatedly performing or by *refraining from* performing associated tasks), these two skills increase *incidentally* as the player moves through the space by jumping or running. Increasing these two skills is unavoidable as the player navigates the game, unless the player is possessed of superhuman patience and chooses not to jump or run at any point. In one of our play sessions, we attempted to complete the entire opening sequence without increasing either of these statistics. In spite of a profoundly dreary three-hour long session, we were still unable to avoid gaining a few points in Acrobatics, and thus the *stealth* specialization. This introduces its own type of bias into the system; regardless of what area of interest a player may want to express to the system, she will also be expressing a significant interest in the stealth specialization, without *intentionally* choosing to. This bias towards the stealth specialization is active during all play, except when the character is *sneaking*; while in sneak mode, the chance of gaining skill increases to Acrobatics and Athletics is decreased, presumably because the character is moving slowly and subtly. This particular systemic contradiction is nothing short of ironic. This second form of bias within the system is an inversion of the first problem we

identified: in this situation, *Oblivion* treats certain choices as significant that may not be significant to the player.

These three problems with the design of the opening sequence of *Oblivion* result in an adaptive system that is fundamentally flawed. Not only does it disregard salient information provided by the player about her preferences, but it also has two significant sources of bias built into it, which make it difficult for a player to represent her desires to the game. To a certain extent, it seems that the available class *stereotypes* are insufficient to reflect the nuance of player performance within the game; however, even in cases where a clear stereotype preference is represented to the system by the player, it seems that the *inferential* component is unable to properly reason about it. Of course, these shortcomings of the intelligent recommendation system at the beginning of the game may not be initially apparent to the average player. The relative opacity of the system's workings initially supports the illusion that the game somehow knows something the player does not. The dissonance only becomes evident when a player attempts to actively represent a preference to the system, only to become frustrated by one or all of the abovementioned issues. Even in these circumstances, players who don't attempt to become pure magic users, or pure thieves, may not notice the biases of the system.

Reading for Believability: Expectations Confounded

The failure of the adaptive system in *Oblivion* has further consequences for the narrative believability of the game and the characters within it. We can consider the game as a whole in terms of our above mentioned metric for believability by evaluating how effectively it satisfies both the player's expectations (external expectations in Magerko's (2007) terminology) and the narrative conventions of the fictional gameworld (internal expectations to Magerko). The primary point of contact between the adaptive systems in the game and the player comes in the form of the Non-Player Characters (NPCs) who mediate the player's experience of the gameworld.ⁱⁱ When the system delivers its assessment of the player's class, for example, it does so via the NPC of Baurus, a bodyguard to the emperor of Tamriel [Figure 4-2].



Figure 4-2 Baurus makes a class recommendation. (Screenshot by authors, used with permission)

Because the game presents these characters as entities with human capabilities, the player's expectations are inflected by her knowledge of human interactions. As with any narrative artifact, these expectations may be augmented by the narrative conventions established within the fictional world. For example, the player may not expect humans to be able to fly, but within a game or story about superheroes, flight is a common ability, and so the internal rules of the fictional world overrule the rules of the "real world", permitting the player to suspend her disbelief. In order for this to operate, the difference between the narrative conventions of the fictive world and the player's expectations from the real world must be clearly delineated. Human-like characters in games automatically invoke the player's expectations of human beings until a new set of narrative conventions is established by the gameworld.

In *Oblivion* the characters suffer from this 'phenomenon of expectations', both at the visual/perceptual level and at the behavioural level. The appearance of the characters in the opening sequence is 'human enough' to invite comparison to real humans, however the motion and behaviours of the characters do not make that comparison a favourable one. The notion of the NPC "observing" the player and "guessing her character class" raises her expectations of the cognitive capabilities of the characters, but proves ultimately disappointing when the recommendation fails to conform to her intentions as she has represented them. The opening sequence raises the player's expectations of the responsiveness of the world, only to demonstrably fail to satisfy them. This is congruent with Crawford's concern about expectations, as presented in his book on Interactive Design:

"We set the expectations of our users with the cues we give them and with the language of interaction that we provide to them. Most programs inflate user

expectations and then confound those expectations. We do this by suggesting that the software universe inside our program is larger than it actually is.”

(Crawford, 2003, p. 85)

This opening sequence does exactly that: inflates expectations, and then confounds them. In *Oblivion* these expectations come from what Crawford describes as *conceivable states*. These are the number of possible outcomes that the user can imagine as he interacts with a system (Crawford, 2003). *Oblivion*'s minimum number of conceivable states for a character is 21; one for each of the premade classes. Crawford (2003) contrasts this with the notion of *accessible states*, which is the number of *actual* outcomes afforded by the system. The number of *accessible states* afforded by the adaptive system in *Oblivion* are limited to Combat, Stealth, and Magic, and even within this triad, there is significant bias against Magic.

This problem of non-believable NPC's unfortunately extends beyond the opening sequence. The NPC's continue to fall short in various measures of believability. They do not exhibit temporal realness, which is to say that they do not experience time in a believable way. *Oblivion* has an elaborate calendar system, but the passing of time has zero effect on the characters. They do not age, they do not meet, fall in love and raise a family, and they do not get frustrated with their endless days of standing around on street corners discussing the weather. In addition, NPC emotional reactions are not always consistent with the plot events that should affect their responses. For example, it is possible for the player to become a member of two guilds in the course of the gameplay. However, the NPC guild members can't satisfactorily reconcile this "dual citizenship" state. Instead, they alternate between giving the player warm response accorded to a fellow guild member, and rough or hostile response accorded to a member of an opposing guild – both in the course of a single conversation. In another example, an NPC will express joy at the player's rescue of her missing daughter - even if the rescuing player subsequently killed the returned daughter. These clumsy and contradictory response patterns break the envelope of believability.

Conclusions

By taking two distinct approaches in our reading of *Oblivion* we have uncovered a relationship between the technology of the game's design, and the experience of narrative believability. While each of our lenses is concerned with different dimensions of the game, our different readings have converged on a common point.

We suggest that *Oblivion* is the victim of its own ambition: it raises the player's expectations of the world and the characters but is unable to satisfy them. Like Mori's robots, *Oblivion* falls into an uncanny valley where the apparent offer of subtle and realistic gameplay is betrayed by shortcomings in actual performance. In the process, *Oblivion* undermines the narrative believability of its fictional gameworld. One of most disappointing manifestations of this is the system of "intelligent" personalisation that the game employs as part of its character creation mechanism. In this situation, we see a direct relationship between the limitations of the adaptive system and our experience of character believability. The failure of the adaptive system to

affectively infer a player preference undermines believability in both the NPCs that populate the gameworld and in the fictive reality of the player's own character.

How then can we learn from the design of *Oblivion* if we wish to avoid the pitfalls described in this paper? First, we must consider the ways in which the game raises player expectations, regardless of whether or not it succeeds in satisfying them. *Oblivion* seems to offer several desirable experiences: it promises a world populated with human-like entities to interact with, it promises a world in which the player's actions and choices can potentially result in meaningful responses from the inhabitants, and it promises a diverse range of possible approaches to any given situation without favouring any one approach over any other. The failure of *Oblivion* is not one of poor design or inappropriate promise, but rather of poor implementation and unfulfilled promises. A game in which human-like entities responded appropriately to any one of a multitude of player choices certainly sounds like a wonderful thing, but it is conceptually flawed in a fundamental way. These promises all operate in terms of the expectations of the player, who daily interacts within a world populated by humans who respond more-or-less meaningfully to her every action. Unfortunately, mimicking this level of intelligent interaction is currently beyond any artificial implementation. The dissonance in *Oblivion* comes from a failure to define the narrative conventions (or *internal expectations*) of the world and to adhere to them. As a result, the limitations of the game's functionality come directly into conflict with the expectations of the player, who has not been provided with any framework around which to adjust them in order to be in line with the experience the system is actually capable of providing.

One lesson that can be taken away from this is to design with both sets of expectations in mind. In order to deliver a satisfying and believable experience in a game with the level of narrative richness in *Oblivion*, it is desirable for designers to explicitly identify and shape the narrative conventions of the gameworld in such a way as to operate *harmoniously* with both the technological limitations of the game software and the narrative details of the storyworld. Creating a model of the player's preferences and personalizing the experience using that model is a useful technique for constraining the behaviours of the storyworld; however, as *Oblivion* demonstrates, failing to properly evaluate the player's preferences and interpret her actions can render this technique ineffective or even counterproductive.

References

- Bethesda Softworks. (1994). *The Elder Scrolls: Arena*. Bethesda Softworks,
- Bethesda Softworks. (1996). *The Elder Scrolls II: Daggerfall*. Bethesda Softworks,
- Bethesda Softworks. (2002). *The Elder Scrolls III: Morrowind*. Bethesda Softworks/ZeniMax, Ubisoft,
- Bethesda Softworks. (2006). *The Elder Scrolls IV: Oblivion*. ZeniMax Media Inc., March 24, 2008.
- BioWare. (2003). *Star Wars: Knights of the Old Republic*. [PC], LucasArts,
- Bruner, J. S., & Postman, L. (1949). On the Perception of Incongruity: A Paradigm. *Journal of Personality*, 18(2), 206 - 223.
- Brusilovsky, P. (2001). Adaptive Hypermedia. *User Modeling and User-Adapted Interaction*, 11(87-110).
- Brusilovsky, P., & Millan, E. (2007). User Models for Adaptive Hypermedia and Adaptive Educational Systems. In P. Brusilovsky, A. Kobsa & W. Nejdl (Eds.), *The Adaptive Web* (pp. 3-53). Berlin Heidelberg: Springer-Verlag.
- Champion, E. (2007). *Social Presence and Cultural Presence in Oblivion*. Paper presented at the PerthDAC.
- Chihido, D. (2008). Reinventing *Oblivion*. Retrieved March 23, 2008, 2008, from <http://www.xbox.com/en-US/games/t/theelderscrollsIVoblivion/20051110-fe.htm>
- Crawford, C. (2003). *The Art of Interactive Design*. San Francisco, CA: No Starch Press.
- Fischer, G. (2001). User Modeling in Human-Computer Interaction. *User Modeling and User-Adapted Interaction*, 11, 65-86.
- Kobsa, A. (1995). *Supporting User Interfaces for All Through User Modeling*. Paper presented at the HCI International '95.
- Lankoski, P., & Björk, S. (2007). *Gameplay Design Patterns for Believable Non-Player Characters*. Paper presented at the Digital Games Research Association (DiGRA) 2007.
- Lionhead Studios. (2001). *Black and White*. [PC], Electronic Arts,
- Magerko, B. (2007). *Measuring Dramatic Believability*. Paper presented at the AAAI Fall Symposium on Intelligent Narrative Technologies.
- Marsella, S., & Gratch, J. (2003). *Modeling coping behavior in virtual humans: don't worry, be happy*. Paper presented at the Proceedings of the second international joint conference on Autonomous agents and multiagent systems.
- Mateas, M. (1997). *An Oz-Centric Review of Iterative Drama and Believable Agents*. Pittsburgh, PA: Carnegie Mellon University.
- Mori, M. (1970). The Uncanny Valley. *Energy*, 7(4), 33-35.
- Murray, J. (1997). *Hamlet on the Holodeck: the future of narrative in cyberspace*. Cambridge, Massachusetts: The MIT Press.

- Oppermann, R. (Ed.). (1994). *Adaptive User Support*. Hillsdale, New Jearsy, Hove, UK: Lawrence Erlbaum Associates.
- Pasquinelli, E. (2006, November 20-21, 2006). *The role of expectations in the believability of mediated interactions*. Paper presented at the Third International Conference on Enactive Interfaces, Montpellier, France.
- Rich, E. (1979). User Modeling via Stereotypes. *Cognitive Science*, 3, 329-354.
- Rozak, M. (2006, April 4). Oblivion: Full spectrum content, from hand-generated to procedural. Retrieved March 29, 2008, 2008, from <http://www.mxac.com.au/drt/OblivionProcedural.htm>
- Wolfreys, J. (2000). *Readings: Acts of Close Reading in Literary Theory*. Edinburgh, Scotland: Edinburgh University Press.

ⁱ Technically any character creation process is “narrativised”, in the sense that creating a character is implicitly about narrative. We use the term in this case to highlight character creation mechanisms that are incorporated *diegetically* into the context of the game.

ⁱⁱ In Oppermann’s terminology, the NPCs represent the *afferential* component of the system.