Chapter 1

Introduction

As virtual worlds become more complex and visually believable, there is an expectation that the characters that inhabit these worlds likewise evolve to match their world. Personality is a key component of believable characters (Mateas, 1997). Often personality is approached from an animation point of view (e.g. Stuart, 2007) to the exclusion of considering how personality affects behaviour or reasoning. Visual appearance of a character can give the false impression of an underlying difference between characters. However, over time, character differences will only be noticed if they result in differences to observable behaviour. Personality should affect the way a character evaluates its personal success, makes decisions and, as a direct result, how the character behaves in the world. Within the field of games and virtual agents, there are many implementations of personality. Usually these implementations require handcrafting each individual character in a way that is not suitable for large numbers of characters, or rely on a set of fixed archetypes where every character within the type is exactly the same. As a result, environments with large numbers of characters often develop characters that do not withstand close inspection. We aim to develop a method of automatically generating individual characters with their own unique personalities in order to create more believable characters in virtual environments with large numbers of interactive characters.

Personality should be unique for each character, i.e. individual. Cognitive-social theories of personality address how personality is developed throughout our lifetimes. According to these theories, two essential components of personality are that it should be able to adapt or change; and that situation or context is important in determining

how we behave. That is, in real life, people react differently in different contexts based on their past experience and the underlying goals they are trying to achieve. The term *context* could relate to a physical location or other participants in the activity. We will be looking at the case where context relates to how close the character believes it is to achieving its goals. For example, one context could be that the character has achieved one of its goals, but not another. Over time people adapt their behaviour based on past experience of different contexts. People will adapt and change their own behaviour, and to some extent their personality, based on other people they interact with and changes in the environment. Characters within current state-of-the-art games and virtual agent environments do not demonstrate all of these features unless explicitly handcrafted to do so. We will be concentrating on observed personality as a sufficient requirement instead of "actual" personality. We do this because human users interacting with these characters can only distinguish characters based on their observed behaviour. From this basis of cognitive-social personality theories, we believe that personality should be: *individual*, *adaptive*, and based on goal *context*. We now present our motivation for this thesis followed by a description of the research questions that will be addressed within this thesis.

1.1 Motivation

This thesis addresses important issues about how to make game characters and virtual agents more realistic or believable by using cognitive-social personality theories. We begin motivation by discussing two examples. The first is from a television series and the second is a more concrete example that will be used throughout the thesis to explain concepts and the model. These examples represent the ideal goal of our work. After these examples have been presented, we discuss where the theories that motivate this work have come from, as well as why games and virtual agents are in need of the work presented here.

Our first example come from the new version of *Battlestar Galactica* (by the SciFi channel) which is about a future war between humans and Cylons. Cylons are similar to robots and have been built by humans according to the back story. Some Cylons have been built to look like humans externally. There are 12 different "models" of these human-looking Cylons. A model includes the physical appearance of the Cylon

(a) Boomer (b) Athena (c) Number Eight (d) Copies of Model 8

Figure 1.1: Model Eight from *Battlestar Galactica* (produced by the SciFi channel): There are many copies of model Eight (also known as Sharon Valerii). Each copy is generated from the same template, and yet takes on its own unique personality[1].

and some personality characteristics. There are many Cylon copies built of each of the models, so that externally each of these Cylons can appear to be the same "person" at first glance, e.g. Figure **??**. However, each individual Cylon, or copy, from the same model is able to build up their own memories based on their experience and this changes their future behaviour and eventual personality.

We look at the example of model number "Eight", Sharon Valerii, see Figure **??**. Within the series we are introduced to this character without realising she is a Cylon. In fact, she does not realise herself that she is a Cylon until later in the story. As the series progresses, we are introduced to other copies of number Eight. Each copy has her own experiences and after many episodes we are able to distinguish at least three different "Sharons". The first Sharon we meet is nicknamed "Boomer", Figure 1.1(a). Boomer has a romantic relationship with the chief human engineer but eventually sides with Cylons against the humans. Next is Athena (figure 1.1(b)), who has always known she was a Cylon, unlike Boomer. Athena fights on the side of the humans, is in a relationship with a human and has a child with him. There are also many copies who remain with the main Cylon attacking force. Mostly these copies are indistinguishable from each other. However, as the series progresses the "head" of the number Eights shows increasing interest in Athena, to the point that she downloads Athena's memories

[1]Figures a and c taken from `http://www.thescifiworld.net/interviews/grace_park_02.htm`; Figure b from `http://www.imdb.com/media/rm4068644352/tt0407362`; Figure d from `http://en.wikipedia.org/wiki/File:Battlestar_Galactica_1x12_Number_Eights.jpg`

and yet still remains different from Athena. This copy also heads a group of Cylons who want to help the humans rather than destroy them.

What is interesting about this example is that the unknown designers of the number Eight model set in place only one personality template. The different copies become distinguishable based on their own experience, with no extra input from the designers. The designers are able to generate a personality template for each model and then allow copies to become their own individuals over time, while retaining some behaviour similar to the other copies within that model number. This example is motivating because the designers handcraft only twelve Cylon models and yet each of the many Cylon copies can have its own individual personality. This ability to produce many different characters from a handful of templates represents the overall goal of this thesis, i.e. to be able to generate individual characters with personality without handcrafting all behaviour.

To motivate our work further, we describe an example of an ideal virtual environment that could be built using the techniques that are developed in this thesis. This example or vignette will be used throughout the thesis to explain concepts and the model.

1.1.1 Motivating Vignette

Imagine a virtual world populated with diverse characters who live in villages and who each have their own personality distinguishable from others. Usually in this kind of world, the personality templates that are used become visible to a game player after prolonged interaction with the world. For instance, if you have met the baker in one village and travel to another village, you will probably meet essentially the "same" baker. Imagine if this new baker still had many of the core traits of the other baker, but had a different personality, i.e was recognisably different. For example, the baker in the first village might sing while serving his customers, while the second might be sullen and grumpy towards his customers, but may make better quality bread. Finally, imagine that these differences are not handcrafted, but are automatically developed based on each character's personal experience with other characters, including the player, and the world. Further, if something changes in the world, such as a stranger entering it, the characters' personalities will change over time to reflect their new experiences. For

example, if the stranger is very hostile existing characters will change their behaviour around the stranger and perhaps with other characters as well.

This is the world we are attempting to achieve.

Let us look at this example in more detail, since we will be using it as a case study throughout the thesis, beginning by looking at how to build one village. In our goal model, a limited number of carefully crafted personality templates would be designed. A personality template includes information about how the character reacts to events, what their personal goals are and the actions or choices available to it within the world. These personality templates form the building blocks of all the characters. Characters that have the same initial personality template will not necessarily have the same resulting personality, similarly to identical twins and Cylons gaining their own experience and becoming different from their twin or, in the case of Cylons, their model number. Personality is generally considered to be a result of both nature and nurture. In our goal model, the personality template represents a character's nature, whereas experiences in the environment represent a character's nurture.

Using these personality templates, we generate a large number of characters for our initial village. Then we allow the characters to interact with each other and learn how to achieve their personal goals. This allows them to take on roles within the village, such as butcher, baker or candlestick maker. Characters learn which action they prefer to take and how they prefer to execute the actions. These preferences are different depending on the character's perceived context or state of its current goals. For example, when a character is not achieving its goal to make friends, it may prefer to talk to someone new, whereas when it has many friends it may prefer to do something different. Characters with the same personality template will not be exactly the same. This is because of the different relationships developed and the different methods the characters have found to achieve the common goals. Now we have a village full of interesting characters who all have their own history and relationships with the other characters.

Next we use the same starting personality templates, but with a different random seed, for a second village. Again we allow the characters to interact with each other and develop their preferences and personalities. We now have a new village full of interesting and different characters, and yet these characters will be observably different from the characters in the original village. A human player who walks between the two villages

will see some similarities, but will not be faced with "the baker" who is always the same in every village. For example, in one village the candlestick maker and the baker might be married, in another village they may never have met.

Villages will be able to generate a community personality of their own. For example, if all the characters have a goal to make sure everyone is not hungry, then one village might rely on giving away goods for free and then receiving other goods for free, whereas another village works only by the use of money or bartering. If a baker is trying to be generous and giving away bread, it will still need meat and candles. Unless the other characters cooperate and give it these items, it may become hungry or will not have light.

The villages will not be able to function well unless the roles of 'butcher', 'baker', and 'candlestick maker' are filled (there may be more roles in an actual game). However, who takes each of these roles is not defined by the personality template; it is discovered by the character as a suitable way of achieving their personal goals. For instance, the goal of making money could be achieved by any of these roles, but would probably not be successful unless the character specialises in only one of the roles. The model proposed in this thesis reflects the real world, where some personality "types" are more inclined towards certain professions or roles. However, personality is more than someone's job. The way that a character executes their role and their overall behaviour informs the observer of their entire unique personality.

One issue encountered in the games industry, is that characters capable of learning can be considered risky because they may develop new behaviour that may offend the players or change the gameplay significantly. The benefit of this model is that since the personality template restricts the actions available to the character (and this does not change over time), the character's actions cannot become entirely unpredictable. That is, characters cannot generate new actions, they can simply choose differently. What gives the diversity and appearance of a level of unpredictability is the choices the character makes and its preferences for those actions. An individual character, if watched, will eventually become predictable in their actions. However, if a player sees a character of the same personality type they will not be able to entirely predict the new character's behaviour without watching it for some time. By using the model in this way, characters will not be able to generate "dangerous" or unconstrained behaviour that becomes unsuitable in a shipped product.

1.1.2 Motivation from Psychology and Personality Theories

When considering the motivation to pursue this topic, we need to establish why personality should be an important aspect of any virtual character. Chuck Jones, a cartoonist for Warner Brothers, found that in creating believable characters, personality is the most important aspect: it is "the individual, the oddity, the peculiarity that counts" (Jones, 1994, p.14). Personality and emotions have been successfully used as "filters to constrain the decision process when selecting and instantiating the agent's behaviours" (André *et al.*, 1999). Personality can be considered as the engine that generates reactions and responses in a coherent, consistent and predictable manner (Ortony, 2002).

We believe that personality is visible in the observed behaviour of characters. Therefore, the action a character chooses out of many possibilities should be a reflection of the character's personality. Differences in behaviour for the same person and for different people are due to a number of factors including emotionality (differences in emotional reactions to events), current state and interpretations of the world situation (Ortony, 2002). Or, according to another theory (Lazarus, 1994), people respond differently to similar events depending on their individual goal hierarchies and perceived current context of themselves and the world. That is, behaviour or responses are context-dependent. We believe that context should be based on the current level of achievement of goals. For example, if a character currently has a lot of money but no friends, it may achieve his goals better if it gives away food, rather than selling it. We will take these concepts into consideration when building our model for personality that generates behaviour.

Trait-based theories of personality are popular for virtual characters. However, trait-based theories assume that personality is static and unchanging, and offer no explanation of how personality is developed Cervone & Pervin (2008). Although this may be appropriate for many environments, it assumes that the designer is able to develop a suitably complex personality for every character in the environment and predict the situations the character will encounter. Cognitive-social theories of personality (for example Bandura's social learning theory (Bandura, 1977)) address the issue of personality development. If we allow characters to develop their own personality according to cognitive-social theories, then we will be able to use a simple template to generate

many different complex personalities. This process will hopefully reduce the burden on the designer.

Cognitive-social theories believe in reciprocal determinism: that is, behaviour results from the complex interaction of persons and the environment, rather than from any single factor alone (Bandura, 1977). Hence, people are neither driven by inner forces nor buffeted by environmental stimulus. The traditional view of behaviour interaction is that a person's behaviour is a function of the person and the environment. However, people's actions contribute to the overall environment, which will in turn reflect behaviour in a reciprocal fashion (Bandura, 1977). In our village example, if one character gives away food it may cause those around it to also give away food, i.e. the environment has changed because of one character's behaviour. Experiences that a character generates through their own behaviour will affect what a person becomes. That is, if a person tries to bake bread (behaviour), and they are successful, they may continue to do this (changing the person). Both the person and the environment will in turn affect subsequent behaviour (Bandura, 1977). There are three main types of learning according to Bandura, by response consequences, through modeling, and self-reinforcement (Bandura, 1977). People learn and adapt via these mechanisms throughout their lives. Further, according to cognitive-social theories, behaviour is not determined by global traits, behaviour depends on the situation that the person is in (Cervone & Pervin, 2008).

An inspiration for this thesis comes from Damasio's somatic marker hypothesis (Damasio, 1994). According to this hypothesis, emotions in the form of "gut instinct", called somatic markers, guide the decisions we make. Without somatic markers it would be a struggle to interact rationally with other people (Damasio, 1994). Somatic markers are built up automatically throughout our lives, so that when faced with a possible decision in a particular context, they guide us towards or away from certain choices. After this (often unconscious) elimination process, we are able to make a more studied analysis of choices available in order to decide which to choose.

Cognitive-social theories and the somatic marker hypothesis represent ways that real humans interact with other humans. Somatic markers represent an aspect of personality, in that our gut instinct guides us towards certain choices in a way that is different from how another person may be guided. Somatic markers are based on past

experience, and we will use cognitive-social learning theories to inform how our characters acquire these somatic markers. This acquisition process and using somatic markers for decision-making is based on the character's context, to reflect that people choose different actions in different situations. In our thesis, we implement a combination of learning by self-reinforcement and by response consequences using a reinforcing function to build up actual somatic marker values and therefore influence behaviour and the environment.

From cognitive-social theories and the somatic marker hypothesis, we believe that personality is unique to each individual, it should adapt to the environment and should be context dependent. As we will show, these factors are not fully implemented in current computer games and virtual characters.

1.1.3 Motivation from Virtual Agents Domain

Virtual agents can be seen in a number of application areas, from military simulations to pedagogical environments to embodied conversational agents. Believability is one of the key goals for most research groups in this area. Believability will make the characters more engaging and users will have a more enjoyable experience (Johnson *et al.*, 2000). Characters that have unrealistic behaviour are more noticeable and distract users from the virtual world (Johnson *et al.*, 2000).

It has been thought for a long time that the use of emotions in virtual environments improves decision-making and believability of characters (Minsky, 1986). Over the past fifteen years, there has been increasing use of emotions to improve realism and believability of intelligent agents in the agent research field with much initial success. This research by others into how to implement emotions in virtual agents is important, but often does not fully address the issue of how to implement personality. Personality gives life to characters, not emotions (Lim *et al.*, 2005). Some research has concentrated on how to animate characters with their own personality, i.e. change the visual appearance of behaviour. Our work concentrates on how to imbue the character's decision-making and evaluation processes with their own personality without handcrafting every step. We believe this will contribute to character believability. We concentrate on implementing personality that is unique to each individual; adaptive over time; and reflects the context the character is in.

In the virtual agents domain, most characters developed can generally be considered to be individuals. However, this is often because there are no other characters to compare against, or because the individuality has been handcrafted, for example, the Oz project (Mateas, 1997). The Oz project gave equal attention to believable characters and story or interactive drama. However, the characters were handcrafted to obtain the desired effects.

Most implementations of personality are based on trait-based theories, such as the popular Five Factor Model (McCrae & John, 1992), and are static with respect to time and contexts. However, personality does develop over time, particularly during childhood experiences or when the environment itself changes in a substantial and long term manner. For instance, a character that is constantly ignored would be expected to change its behaviour over time to reflect this. Over extended periods of time, characters will be more compelling if they appear to learn from experience (Blumberg *et al.*, 2002). However, characters should adapt in a way that is consistent with their personality (Mateas, 1997). Characters such as Blumberg's dog (Isla *et al.*, 2001) can adapt, but not in conjunction with its personality and only when taught explicitly by the user, i.e. not by itself. Static implementations of personality do not provide support for context-aware behaviour. In a static implementation, the number of starting personalities is the same as the final number of personalities. Since the characters cannot learn, their personalities will be fixed over time, so two characters with the same personality will behave the exact same way without any variations unless explicitly included. In order to present a believable complex personality, characters should be seen to make decisions based on their past experience via adaptation, and based on their perceived context.

1.1.4 Motivation from Games Domain

The model we seek to build will ideally contribute to the body of work in computer games. However, much of the academic work in emotions, personality, agents and learning is difficult to implement. Learning techniques are often complex for a lay-person to understand, can result in unpredictable behaviour, require significant computational power and require handcrafting of individuals. This results in making many models unsuitable for use in games.

Many recent computer games include large numbers of computer-controlled characters, such as in *Grand Theft Auto*, *The Sims*, and *Oblivion*. In many cases, these

characters have been created using a limited number of handcrafted personality "types" or archetypes (Ellinger, 2008). Every character with a particular personality type is essentially the same and not distinguishable from others of the same type, that is the characters are not individuals and appear homogeneous (Russell, 2008). For example, the player can meet someone in a store in one part of town, move to another part of town and essentially meet the same character again, even though they are supposed to be different people. The longer a player spends in the world the more likely these similarities will be noticed and will decrease the player's enjoyment of the game due to excessive repetition and predictability of behaviour.

A different approach is used in some other games, such as *Half-Life 2* where the player is required to work with a single computer-controlled character for much of the duration of the game. In *Half-Life 2*, the character is explicitly scripted so that its behaviour will change throughout the game. Unless this scripting is implemented by the designer, the character will not be able to adapt its behaviour over time and can result in the appearance of one-dimensional characters. In *Black and White*, the creature character can adapt, however it is only via the explicit teachings or reinforcements of the player, and not independently or in conjunction with its personality. In-game learning is probably only suitable for characters that the player interacts with for prolonged periods of time. However, learning prior to shipping the game will allow unique characters to emerge and will improve the diversity of characters a player meets, even if the player only has a short interaction with the character.

To be believable, characters should behave differently depending on their context. Within games, the characters can only recognise different contexts if the game developer has explicitly included that capacity with hand-scripting. For example, many characters are developed using finite state machines. If the character is in a particular state (e.g. surrounded by enemies), it will perform one action; if it is in another state, it will perform a different action. However, these differences must be hard-coded by the designer, who must consider each and every state and all the action possibilities available to the character. In our ideal model, the character would be able to learn which actions are appropriate for it personally on its own without enforced player interaction or hard-coding preferences.

1.2 Research Questions

Personality can give life to virtual characters. However, current applications of personality require a large amount of handcrafting. We seek to reduce the amount of handcrafting required for a designer to build many individual characters. That is, this thesis aims to answer the question: How can characters with personality be created without handcrafting all behaviour? In addressing this question, we believe that cognitive-social personality theories offer a method of developing characters based on their individual experiences. Cognitive-social theories maintain that personality is both adaptive and context-aware (situated). By common definition, each personality is unique. So a requirement of personality itself is that characters should be individuals. These requirements lead to the following research questions:

1. How can a model of personality be created that uses *adaptation*? How does *adaptation* affect character behaviour?

2. How can a model of personality be created that uses *context*? How does *context* affect character behaviour?

3. How can personality be implemented so that the same template can be used to create a number of distinct, *individual* characters who behave differently?

These research questions can be broken down into a set of sub-questions relating to creating the model and testing the model: model-based and testing-based.

1. Adaptation:
   - Model-based sub-questions:
     - What aspects of personality can adapt?
     - How are decisions made?
     - How can characters calculate reward?
     - How can characters use reward to update behaviour?
   - Testing-based sub-questions:

   1.(a) Does behaviour change over time?
   1.(b) Can characters learn about specific, functional goals?
   1.(c) How does reward change with time?
   1.(d) What happens if adaptation is turned off?

2. Context:
   - Model-based sub-questions:

  - How can context be represented?
  - How can we fill in context information?
- Testing-based sub-questions:

2.(a) Is character behaviour different in different contexts?
2.(b) What happens if context is turned off?

3. Individual:

- Model-based sub-questions:
  - What *is* an individual within our model?
  - What is a personality template?
  - How does personality change over time (i.e. how can a character be different from another character with the same template)?
- Testing-based sub-questions:

3.(a) Are the behaviours of characters different from each other over time?
3.(b) Are any individuals obtained?

The model-based sub-questions can be answered by defining terms appropriately and implementing the model in an application. It is not possible to measure the validity of these subjective answers directly, the model itself is the "answer". To determine whether the model developed is successful, we implement a game to be run with many different starting conditions, we consider sub-questions relating to testing and establish a set of criteria for success. From the runs of the game, we can extract data relating to three quantifiable measures of effectiveness: behaviour, reward and individuality. *Behaviour* relates to what the player can see on the screen, the actions that the characters choose to execute. In particular, we count the number of times that characters choose different actions over regular output time periods. *Reward* relates to how well the character is achieving their own goals according to their own personal evaluation. High reward values mean that the character is achieving its personal goals in the current environment. *Individuality* measures the number of differences between all characters in the game based on their behaviour. The maximum number for individuality implies that each and every character is completely different from every other character in the game, in the respect that they choose different actions at the same time periods.

Using behaviour, reward and individuality as measures of effectiveness of the model, led to a set of quantitative criteria for success the model should satisfy to be determined successful. Each criterion for success addresses a specific testing-based research sub-question as shown in Table 1.1. We recognise that the choice of the majority as a cut-off

Research Questions and Testing-based Sub-questions	**Criteria for Success**
1. How does adaptation affect character behaviour?	
1.(a) Does behaviour change over time?	*Behaviour* changes over time.
1.(b) Can characters learn about specific, functional goals?	When given a functional goal to learn, the majority of characters choose the "correct" action the majority of the time, based on *behaviour*.
1.(c) How does reward change with time?	*Reward* values are on average higher using our model than when random choice is used.
1.(d) What happens if adaptation is turned off?	Compared to when adaptation is turned off, both *individuality* and *reward* are higher when adaptation is used.
2. How does context affect character behaviour?	
2.(a) Does character behaviour differ in different contexts?	For one character's *behaviour*, show that in different contexts the action chosen the majority of times is different.
2.(b) What happens if context is turned off?	Compared to when context is turned off; both *individuality* and *reward* are higher when context is used.
3. How can personality be implemented so that the same template can be used to create a number of distinct, **individual** characters, according to their behaviour?	
3.(a) Are the behaviours of characters different from each other over time?	Character *behaviour* passes the chi-squared test.
3.(b) Are any individuals obtained?	Based on their *individuality*, at least one character is different from the majority of the other characters when they are all based on the same template.

Table 1.1: Criteria for success to be used to evaluate testing-based research sub-questions. Words in *italics* are the measures of effectiveness.

for some of the criteria is somewhat arbitrary. Having at least the majority of characters demonstrate a desirable behaviour shows that they are behaving as intended. We simply use the majority as initial criteria so that we can quantitatively rate the model that is designed and implemented in this thesis. Note that a chi-squared test measures whether the characters' behaviour (actions over time) are actually independent from each other.

1.3 Thesis Overview

This thesis begins with a literature survey of related work and introduces pertinent theories. This allows us to place our work within the broader research field and explain the theories that are used in our personality model. After the literature survey, we introduce our model for agent personality development, Chapter 3. This chapter details the key components that characters need, as well as the adaptation loop that characters use to make decisions and evaluate their choices. It shows how we use cognitive-social theories and somatic markers to develop personality that affects decision-making and evaluation. The model is generic and can be applied to many domains. We detail how to build individual characters and how to attach domain-dependent specifics to the generic model.

To test our model, we developed an example game domain. The game developed is based on simplistic school children who can insult each other and move around their world. Limited implementations have been shown to generate large numbers of possible paths or actions for characters (Theune *et al.*, 2004). Within our simple game implementation, we were able to generate a large amount of complexity due to interactions between characters. In order to determine success of the model in terms of being able to generate different individuals from the same personality template, we use three measures of effectiveness: behaviour, reward and individuality. In Chapter 4 we introduce our implemented game, discuss our experimental setup, including how the measures of effectiveness are obtained, present our method for answering the research questions, the scenarios used and consider expected results.

In Chapter 5, we present our results to the testing-based research questions as well as some interesting side results that we found while answering the research questions. Results are based on the criteria for success and the measures of effectiveness. The side

results relate to domain-dependent aspects of our implementation. In this chapter we also discuss the implications of the results obtained.

In the final chapter results are summarised, particularly in relation to the research questions. We discuss future directions for research, implications for the game industry, and contributions made by our agent personality development model.