Chapter 6

Conclusion

The main aim of our research was to develop a method of generating individual characters with a minimum of handcrafting so as to populate virtual environments with large numbers of interactive characters. We have developed a model that generates individual characters with the same initial personality template. The results showed that the personalities generated were adaptive, characters developed different activity preferences for different contexts, and the characters were distinctly individual (in a computational sense). In this chapter, we begin by discussing the research questions in relation to our final model. Next, we discuss further testing and future work that could improve the model. This is followed by the implications of this model for the games industry and some final words to conclude the thesis.

6.1 Addressing the Research Questions

The major research questions presented in the introduction (see Section 1.2, page 14) relate to building and testing a model of personality so that the characters generated are adaptive, context-aware and individual. We consider each of the research questions in turn and address the related model-based and testing-based sub-questions.

6.1.1 Research Question 1: Adaptation

How can a model of personality be created that uses *adaptation*? How does adaptation affect character behaviour? Personality can be seen in the actions taken by characters. Characters choose which actions to execute based on their

preferences (implemented as somatic markers). So, to include adaptation in personality, the model needs to adapt the preferences. This can be done by updating the preferences after execution of actions within a feedback loop. Self-reinforcement (personal reward) learning (from cognitive-social theories) can be used so that the preferences learnt are based on the personal soft goals of the character. Which soft goals a character is trying to achieve (and the importance of each soft goal) is part of the character's personality template. So the personality template affects how the character learns the adaptive part of its overall personality. We now consider the model and testing-based sub-questions relating to adaptation.

- Model-based sub-questions:
  - **What aspects of personality can adapt?** Somatic marker preferences, which are used to make decisions, can change based on experience and are part of an character's personality.
  - **How are decisions made?** Decisions are made by following the appraisal of coping choices process (inspired by the somatic marker hypothesis) that uses somatic markers to place preferences on available plans in the current context.
  - **How can characters calculate reward?** Reward can be calculated based on a weighted sum of an individual's current achievement levels of its soft goals (based on its soft goal personality).
  - **How can characters use reward to change behaviour?** Personal reward can update somatic markers according to the reinforcement comparison technique. Somatic markers affect behaviour by changing the character's preferences towards or away from choosing certain actions.

- Testing-based sub-questions:

1.(a) **Does behaviour change over time?** According to our results, the characters were able to change which actions they executed most frequently based on feedback from the environment and their personal goals. This means that characters can change to suit different players or design changes to the environment. We met this criterion because the character's somatic markers adapted over time, so that the preferences for action changed, based on personal reward.

1.(b) **Can characters learn about specific, functional goals?** Our results showed that characters could. This means that a designer can give a character a goal to achieve in the world or a role to play in the game and the character will be able to automatically learn how to do this (i.e. without further external intervention). Further, it means that the overall character personality that is generated can still fulfil functional requirements. We met this criterion because characters learn based on personal reward and personal reward is based on the soft goals the character is trying to achieve.

1.(c) **How does reward change with time?** Reward fluctuates greatly in response to fluctuating environmental feedback. The reward values that the character calculates are higher over time when the character can use adaptation. This means that adaptation improves achievement of soft goals.

1.(d) **What happens if adaptation is turned off?** The personal reward and the individuality of characters decreases. In terms of reward, this means that adaptation helps characters learn how to achieve their soft goals, which improves their reward. In terms of individuality, the results indicate that characters are more different from each other when they can adapt. Characters who are choosing actions randomly may appear different over one time period but, over an entire game, these characters will not be significantly different from each other. Characters who can learn their own preferred actions to achieve their soft goals will become different from the other characters over the entire game.

6.1.2 Research Question 2: Context

How can a model of personality be created that uses *context*? How does context affect character behaviour? Characters execute different actions depending on their context. Characters choose which actions to execute based on their preferences (somatic markers). So, the character's preferences need to be based on context as well. In our model, we choose context to relate to the soft goals the character is trying to achieve (part of its personality template). To know what to do in a specific context, the character learns, via adaptation, their preference for that particular action and context (rather than hard-coding the preference). We now consider the model and testing-based sub-questions relating to context.

- Model-based sub-questions:
  - **How can context be represented?** Context is a string representing a discrete version of the achievement level of each of a character's soft goals from its soft goal personality.
  - **How can context information be provided?** Characters update their somatic markers based on the context they were in when they made the decision they are evaluating.

- Testing-based sub-questions:

2.(a) **Does character behaviour differ in different contexts?** We found that the most preferred action was different for a character in their two most frequently used contexts. This means that characters can learn which action they personally consider most suitable according to whether specific soft goals are or are not currently being achieved. This results in more complex and realistic behaviour and therefore overall personality. We met this criterion because, when characters learn somatic markers, they link their reward to the action and the context in which the action is chosen.

2.(b) **What happens if context is turned off?** When context is turned off, reward and individuality are likely to improve, according to our results. This is the opposite of what was expected. In terms of individuality, this means that being able to learn about different contexts did not allow characters to consistently find a clear solution path that is different from the other characters. In terms of reward, this means that characters are less likely to learn how to achieve their soft goals when they have many soft goals to achieve simultaneously. This could be because there were too many contexts in some Cases or the way contexts operated within the specific implemented domain.

6.1.3 Research Question 3: Individuality

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? Personality is made up of a fixed template and an aspect based on individual personal experience. Different personal experiences will result in different

preferences for actions, and therefore different actions executed, leading to different observed personality, even when the template is the same. We now consider the model and testing-based sub-questions relating to individuals.

- Model-based sub-questions:
  - **What *is* an individual within our model?** An individual is comprised of a number of components and beliefs including personality template, somatic markers and domain-dependent knowledge. Individuality is distinguished based on observable behaviour, in particular the actions that a character executes over a time period.
  - **What is a personality template?** A personality template guides the character's learning and decision-making. In our model, it is made up of a goal/plan hierarchy, a soft goal personality and emotionality.
  - **How does personality change over time (i.e. how can a character be different from another character with the same template)?** Personality changes over time by adapting somatic marker preferences. Somatic markers are adapted based on personal reward experience, so no two characters will have exactly the same somatic markers, even if they began with the same initial personality template.

- Testing-based sub-questions:

3.(a) **Are the behaviours of characters different from each other over time?** Characters were different at the level of the actions they executed. Results showed that characters do not have the same most frequently chosen actions as other characters over an entire run. We found that the behaviour of one character cannot be used to predict the behaviour of another over a run. We obtained this because the learnt somatic markers allowed the character to find stable ways to achieve their goals in a different way compared to other characters.

3.(b) **Are any individuals obtained?** We found that in half the Cases, we were able to obtain at least one character that was significantly different from the majority of the other characters. This means that some characters are more different from others and are therefore more likely to have an observably different personality. We obtained this result because characters had their

own unique experiences, which changed the somatic markers they obtained, which in turn changed the actions that character preferred to execute.

6.1.4 Implications Arising from Research Questions

Our results indicate that our model can achieve our aim to build individual, adaptive and context-aware characters from the same template. This means that individual preferences for every character for every context do not need to be hard-coded, they can be learnt based on an automatic processes and their personality template. Traditionally in personality research for virtual characters, it is very difficult to quantify results. We have used a set of criteria for success for our model to satisfy prior to user feedback questionnaires and studies. These criteria allow us to determine the areas of our model that are in need of improvement, in this circumstance, context-aware behaviour. That is, the benefits of context were not as clear as was expected. Satisfying the criteria does not mean that the model can necessarily produce characters that are distinct to human observers. However, satisfying many of the criteria indicates that the model has promise. The model did not fully satisfy all criteria, so further testing is needed. In future testing, it might be possible to determine a value above which the individuality measure should be in order for characters to be observably different according to players. That is, players could be presented with several versions of the same game, but with different character individuality measures, and, based on feedback, a minimum individuality value may be discovered. The implemented domain was relatively simple domain, yet it could generate complex friendship networks and obtain distinct differences between the characters. The results showed that adaptation allowed characters to learn their own unique way of achieving their personal goals. However, there is scope for improvement in the level of individuality and the spread of differences across characters, particularly compared to when context is not used.

6.2 Potential for Future Research

The model presented in this thesis shows some promise as a method for generating many complex personalities from a single personality template. However, the results indicate that further work is needed to resolve some of the problems that emerged during testing, particularly relating to context and improving individuality per character. We

begin by discussing tests that could be done on the model as it is now relating to personality templates and input parameters. Next, we investigate how the model could be improved or extended, based on current results. Then, we consider visualisation of some of the complex components from our implementation. Finally, we discuss several different qualitative tests with humans that could be performed to test the model (or future models) further.

6.2.1 Personality Templates and Input Parameters

For our testing, we choose a small variety of Cases to demonstrate that the basic principles of our model were effective. When choosing the Cases, we fixed many parameters so that we could use simple Cases to test the model extensively for those Cases. However, there will always be more Cases and further tests that could be completed on the current implementation of the model. By investigating the current model and implementation with more Cases, we may be able to establish more precisely the effect of context on reward and individuality and whether high individuality always leads to low reward values. This process will allow us to pinpoint which areas of the model or implementation need more work, and which areas are already suitable.

Although we have validated the model against the criteria for success relating to our testing-based sub-questions, more stringent levels for the criteria could allow further benefits and problems to emerge. Characters had the same soft goal personality, meaning they had the same weights and ideal values for the soft goals to be pursued. It would be interesting to compare Cases where the soft goals were the same but the weights and/or ideal values were different for each character. Part of emotionality is the trade-off between exploration and exploitation, which was not fully investigated in the implemented domain. A more thorough investigation into the influence of the learning parameters (α, β and bucketing-related values) may give more stable learning and less abrupt changes in preferences. One possibility would be to change this parameter over time so that initial values encourage the character to explore its world and later encourage the character to stay with its preferred choices. The results from the current Cases indicate that there was a trade-off between individuality and personal reward. That is, some Cases that showed high individuality also had characters generating low personal reward values. This could have been a side-effect of the specific Cases chosen so more Cases would need to be constructed to resolve this. With different Cases it

could be necessary to investigate extending runtime and changing output time steps further.

6.2.1.1 Testing Adaptation Further

In our testing, we compared characters using our model in modes where adaptation was turned off. Turning off adaptation meant that the characters were effectively using random choice to make decisions. In order to further test the effect of adaptation, a method would be to run the simulation for a specified duration. After this, character data could be collected, that is, the character's somatic marker preferences that were developed. Then, the simulation should be run again twice: firstly where the characters can continue to adapt, and secondly where the characters cannot change their somatic marker preferences any further. We could then test each run against the criteria. This test would allow us to determine whether the benefit from adaptation is only during the development of character personality, or whether continued adaptation improves or stabilises behaviour, reward and individuality.

6.2.1.2 Designing Specific Characters

Although high individuality was one of our goals, if there is too much individuality within the same personality template the designer will lose control over the types of characters generated. There is obviously a trade-off between the amount of handcrafting of a character and the predictability of their behaviour. On the one hand, if a character is required to fulfil a specific, high-profile, role within the game, then its behaviour and personality probably need handcrafting. On the other hand, for the background characters who contribute to the overall feel of the world, high diversity or individuality with a minimum of time-consuming handcrafting would be more desirable. However, even with background characters a designer does not want them acting in entirely unpredictable ways. That is, a designer may want a number of characters who are "greedy" and some who are "friendly". In our model, these personality types can be set up using a different personality template for each major group. We have shown from our results that, within the same personality template, a number of different characters can be generated based on their preferences for different activities. We need to investigate further whether, after time, the range of characters generated still fit within the overall personality template they were given. That is, if a number of

characters are given a "greedy" personality template, then we need to check whether all the eventual personalities generated remain reflective of how a "greedy" character should act. Although testing would be needed for confirmation, given the way that the soft goal personality has been implemented to constrain characters to achieve specific goals, it is not envisioned that there would be any significant problems in relation to this question.

6.2.2 Model

When designing the model and testing it, a number of ideas emerged as worthy of further investigation. We considered how a somatic marker should be updated if the chosen plan or activity failed. For example, the character wants to talk to someone, but is unable to find someone to talk to. Initially, we decreased the preference value for such failed plans. This led to many plans having a negative preference and the characters did not want to choose any of the plans. In the final implementation, there was no penalty for failing to complete execution of a plan. However, it seems that the characters should be able to learn which plans are difficult to execute and that there should be a better method of adjusting the somatic marker value to reflect this. One possible method would be the redistribution of reward system proposed by Ponsen *et al.* (2006a) (see literature survey Section 2.2.1.3, page 46).

The results show that adaptation in our model improves personal reward and individuality but indicate that the use of context does not consistently improve either of these measures. Does this mean that context is only appropriate in certain domains? Or that the manner in which we implemented context was not appropriate? Perhaps context should relate to something other than goal achievement. Certainly, it was clear that even with only six soft goals there were a large number of contexts the character could be in. This means it is very difficult for characters to learn which action to choose in each context. One solution to this could be to determine a similarity measure across contexts, so that characters could use information from similar contexts to make a decision about their current context. One method is that adopted by Bakkes & Spronck (2006) (see literature survey Section 2.2.1.3, page 48). In this method, the characters can use preferences from similar contexts when encountering a new context. Another method is that every time the character makes a decision, it could use a weighted

function to use preferences not just from its current context but also from similar contexts. A final method is to allow characters to learn for some period of time without context, and then use this learning as a baseline for future learning including context information. We found that reward values fluctuated frequently and did not stabilise or converge. This may be desirable in some game applications, but if it is found to be a problem, then techniques to improve convergence of the learning could be considered, such as those presented by Matignon *et al.* (2006) (see literature survey Section 2.2.2.2, page 57).

6.2.3 Visualisation of Personality, Context and Friendship Networks

The implemented game could generate fairly complex personalities, contexts and friendship networks, despite the simplicity of the actual game structure. However, these three aspects were not directly observable to a player. We now consider some possible games or domains that could exploit further these aspects.

In our implementation, character personality was only visible in the number of times they choose each activity. The characters were only able to choose three obviously different activities and it was questionable whether a player watching a character would be able to distinguish the fact that one character chose a particular activity more frequently than any other activity and, as a result, link this to a different underlying personality. However, it is unclear how to make personality more explicitly obvious. A more in-depth game scenario would give the characters more observably different ways to achieve plans and to allow their personality to be expressed, which would hopefully result in more explicit personalities. The village example from the introduction could be a good starting point.

If contexts are confirmed to be a useful construct for learning personality, there needs to be a better way to enable the player to distinguish that a particular character is in a particular context. The characters were able to smile or frown to indicate their happiness, but this discrete indication cannot convey the complexity obtained when the context is based on many soft goals, such as "LLHLHM" or even "HM". If the domain leads to characters choosing very different activities in different contexts, then the player must be able to distinguish at least some of the different contexts so that the player can understand (and therefore believe) the behaviour of the character.

The domain-dependent friendship networks generated are interesting to examine from a real-world perspective to consider how these networks would be generated or function. For example, in some Cases there was a clear clique of popular characters. It would be interesting to design a game, that exploited these friendship networks in a more visible manner, so that players could notice how the friendships change and alter over time. Perhaps the friendship network could play an explicit role in the player's game, so that the changing network could be monitored and the player could attempt to put themselves at the centre of the network.

6.2.4 Qualitative Testing

In order to determine the success of the model in improving individuality and believability, it will eventually be necessary to use human participants to play the game and answer questionnaires. Although we determined a quantitative measure of individuality, the measure loses significance if the measured differences are unable to be detected by a human participant. The game scenario was chosen to be simplistic, due to time and resource constraints, and also in the hope that personalities would be easier to distinguish in a simple environment. However, to distinguish personalities, a human participant would need to play the game for a long period of time, something that is perhaps undesirable for such a simple game.

Ideally, we would like to measure the coherence of the characters. That is, determine whether their behaviour is stable, identifiable as uniquely theirs and not erratic across time and contexts. It is likely that this can also only be measured using questionnaires of participants. The problem with developing questionnaires is that some participants may be likely to notice differences where there are none and vice versa. To reduce this effect, it would be necessary to ensure that some of the participants are also given a version of the game without our model being utilised fully, i.e blind tests. In this way, participants could compare characters between the Cases.

Ruttkay, Dormann & Noot (2002) proposed a framework to compare embodied conversational agents (ECAs) to each other and to traditional input methods (see literature survey Section 2.1.1.4, page 28). However, this framework still relied on many subjective questions based on the opinions of individuals. To test for different personalities, we would need to attempt to answer such questions as: "how many different characters did you notice in the game?" What one participant uses to measure

"different" may be nothing like another participant's measure. Determining whether the model quantifiably "improves" believability and individuality may not be achievable depending on the robustness of the questionnaires used, the method and the number of participants in the study. We propose another method of determining whether distinct personalities are generated in any simulation: the personality guessing game.

6.2.5 The Personality Guessing Game

The original Turing test proposed by Alan Turing (Turing, 1950) in 1950 related to distinguishing a woman from a man and further whether a machine could be distinguished from a woman (see literature survey Section 2.1.1.4, page 30). Imagine a similar test where you are at a terminal typing to two of your friends. If they change their names, are you able to distinguish the two from each other? If they are good friends, you probably can.

Imagine a game world populated with some fixed number of characters. The interrogator (a player) watches and interacts with the characters for some amount of time to "get to know" who is who and what the characters like to do. That is, the interrogator attempts to observe each character's personality. After some time, we stop the simulation. We change the names and physical appearance of the characters in some way, but do not change their underlying personality and reasoning. That is, although the characters visibly look different, the control of their behaviour is unchanged. The interrogator now watches the characters again. If the interrogator is able to correctly guess which character matches to which original character name then the test would be said to be satisfied. This test gives the interrogator a quantifiable score for the test. This quantifiable value can then be used to measure the effectiveness of the model.

This personality guessing game represents an ideal test of personality in a simulation since it eliminates much of the subjectivity of standard questionnaires. If the underlying personalities are strong enough, they will be observable even when the virtual appearance of the character has changed. If we performed this test using characters that were all generated from the same personality template, we would be able to demonstrate more strongly that our model does automatically generate individual personalities for each character. That is, if the participants are able to distinguish the individual characters better than by random chance, then we can be confident that the personality of each character is individual.

6.3 Implications for Games

The implementation presented in this thesis was of a more academic nature (compared with commercial games), so as to test the model itself. In this section, we consider the model and results obtained and how these pertain more broadly to the games industry. The model developed is generic and can therefore be applied to a large number of different domains, unlike some models that require a complex domain-dependent appraisal of choices. To apply our model to a particular domain, the available personality templates and the equations to determine soft goal achievement levels must be created.

The characters developed using our model are adaptive, context-aware and individual. *Adaptation* improves the number of distinct individual characters, compared to randomising behaviour. Adaptation means that the characters can learn functional goals within the environment, can change automatically to match a new or changing environment, and will automatically learn about the player. The designer does not need to predict all the situations a character will be exposed to because the character can adapt automatically. Furthermore, the character cannot learn entirely new behaviour and become a liability to the game due to its personality template (in particular the goal/plan hierarchy) which limits the behaviour it can exhibit. In this way, characters are not themselves unpredictable, it is the actions they choose to execute that differ. The instability noticed in reward and character behaviour may be well suited to a dynamic game environments. In a dynamic environment, the character would need to be changing their strategies constantly, something the characters generated in our implementation do without effort by the designer. Some game applications take a personality template and use this combined with a new random factor for every decision. However, our results showed that random behaviour does not give distinct individuals over time.

Context allows characters to choose behaviour that is appropriate to the situation the character is in. By combining context with adaptation, the designer can allow the characters to learn what to do in every situation. Our results indicate that context may not be as useful as originally expected. Without further research, it may be simpler to remove contexts from characters in a game environment.

Individual characters are generated from a single personality template. This means that a designer can handcraft a small number of personality templates or archetypes to be used by a large number of characters. Although characters within an archetype will have some similar behaviour, each will be a distinct individual in their own right.

The player feels that they are continually meeting new characters, rather than the same character with the same archetype. This adds to the immersion and depth of the virtual world generated. In order to test our model, we developed a quantitative measure of individuality. This measure is useful in itself so that a designer can test character personality prior to obtaining potentially costly and lengthy player feedback via questionnaires and playtests.

Our model is intended for use in games needing many background characters. Since we sought to analyse characters in detail, we limited testing so that only eight different characters were considered for our illustrative Cases. However, in four of our Cases, all eight characters had the same starting template and yet they developed so that there were measurable differences between the characters in all Cases. Our implementation could easily be scaled to generate more characters from the same template since, once a personality template has been designed, only a single number in the input text file needs to be changed so that the implementation instantiates fifty characters of a particular template, compared to eight.

6.4 Final Words

The model we have designed and implemented shows some promise as a method of developing complex character preferences or personality without handcrafting behaviour for each character in every situation. One of the main technical contribution of this research is the inclusion of soft goals and somatic markers in a BDI architecture. As noted in this conclusion, further research is required to improve the implementation domain and contexts within the model, since the domain was found to be too simplistic and contexts did not clearly improve the characters. In the introduction, we presented an ideal game world that would benefit from our model, the village example. At the end of this thesis, we believe that an implementation of the village example would benefit from using our existing model. This thesis has shown that, through using soft goals and somatic markers to allow adaptation and to perceive contexts, we have developed a model that enables personality to influence decision-making and evaluation for characters in a virtual world. It demonstrates that it is possible to automatically generate, from a single personality template, multiple characters that are unique, adaptive and context-aware, without handcrafting all behaviour.