

Using Adaptation and Goal Context to Automatically Generate Individual Personalities for Virtual Characters

A thesis submitted in fulfilment of the requirements for the degree of
Master of Computer Science by Research

Jennifer Sandercock
(BEng(Hons)/BSc)

School of Computer Science and Information Technology
College of Science, Engineering and Health
RMIT University

August 2009

Declaration

I certify that except where due acknowledgement has been made, the work is that of the author alone; the work has not been submitted previously, in whole or in part, to qualify for any other academic award; the context of the thesis is the result of work which has been carried out since the official commencement date of the approved research program; any editorial work, paid or unpaid, carried out by a third party is acknowledged; and, ethics procedures and guidelines have been followed.

Jennifer Sandercock

14th August 2009

I dedicate this thesis to:

Jonathan Gratch and Stacy Marsella whose work inspired me to pursue
this research area.

The “Duval de L’Epinoy” by Maurice-Quentin de La Tour for keeping me
smiling.

My agents who kept me on my toes and guessing.

**“For there is nothing either good or bad
but thinking makes it so.”**

Shakespeare, Hamlet, Act 2, Scene 2.

Acknowledgements

Firstly, I would like to thank Alex Holkner for implementing the GUI-side of my game implementation and for assisting in the brainstorming process when deciding possible game designs.

I would like to thank Kaye Marion and Andrew Buelke for their generous help with clustering and statistics.

For assistance with JACK programming, I would like to thank Ralph Rönquist and the JACK support group.

Thanks go to Ana Paiva and João Dias of GAIPS at IST in Lisbon, Portugal for help in making the model more understandable.

Massive thanks goes to my supervisors, Vic Ciesielski and James Harland, who stepped up to save me from certain devastation.

Thanks go to my mother, Ellen Stoddart, for helping edit my thesis and for believing in me.

On a personal note I would like thank my friends who helped me keep going, particularly: Rebecca Zurrer, Neil Kirby, David Graham, Tina Smart and Wen-Ching Un. Thanks go, as well, to all my other friends who listened when I needed it. To my companions at university, for helping me stay relatively sane, thank you: Simon Duff, Antony Iorio and David Poutakidis.

Contents

List of Figures	xv
List of Tables	xvii
Glossary	xix
1 Introduction	3
1.1 Motivation	4
1.1.1 Motivating Vignette	6
1.1.2 Motivation from Psychology and Personality Theories	9
1.1.3 Motivation from Virtual Agents Domain	11
1.1.4 Motivation from Games Domain	12
1.2 Research Questions	14
1.3 Thesis Overview	17
2 Literature Survey	19
2.1 Theories	20
2.1.1 Agent Theories	21
2.1.1.1 Beliefs, Desires, Intentions (BDI) Agents	21
2.1.1.2 Goal Types and Motivations	24
2.1.1.3 Cognitive Appraisal Model	27
2.1.1.4 Measurement Techniques for Believability	28
2.1.2 Personality Theories	31
2.1.2.1 Trait-based Personality Theories	32
2.1.2.2 Cognitive-Social Theories	32
2.1.2.3 Individual Differences	35

CONTENTS

2.1.3	Somatic Marker Hypothesis	36
2.1.4	Adaptation Theories	37
2.1.4.1	Aspects that can be Learned	38
2.1.4.2	Reinforcement Learning	38
2.2	Applications	40
2.2.1	Game Applications	41
2.2.1.1	Games with a Personality Emphasis	42
2.2.1.2	Games with an Adaptation Emphasis	45
2.2.1.3	Focus on Research by Spronck <i>et al.</i>	46
2.2.2	Intelligent Virtual Agent Applications	50
2.2.2.1	IVAs with a Personality Emphasis	51
2.2.2.2	IVAs with an Adaptation Emphasis	56
2.2.2.3	IVAs with a Somatic Marker Emphasis	59
2.2.2.4	Focus on Research by Blumberg <i>et al.</i>	61
2.3	Building Blocks: Theories to Be Used in this Thesis	65
2.4	Summary of Literature Survey	67
3	Agent Personality Development Model	69
3.1	Model Components	73
3.1.1	Soft Goal Related Components	73
3.1.1.1	Achievement Level of Soft Goals and Soft Goal Equations	74
3.1.1.2	Context	75
3.1.2	Somatic Markers	76
3.1.3	Personality Template	78
3.1.3.1	Goal/Plan Hierarchy	79
3.1.3.2	Soft Goal Personality	80
3.1.3.3	Emotionality	81
3.1.4	Domain-dependent Knowledge	82
3.2	Adaptation Process	82
3.2.1	Appraisal of Coping Choices	87
3.2.2	Evaluation	90
3.2.2.1	Step 1: Individual Soft Goal Reward Calculation	91
3.2.2.2	Step 2: Calculate Personal Reward	93

3.2.2.3	Step 3: Update Somatic Marker Preferences	94
3.2.2.4	Step 4: Update Emotions	96
3.2.2.5	Step 5: Update the Reference Reward	97
3.3	Summary of the Model	97
3.3.1	Generic Infrastructure Implementation	98
3.3.2	Building an Individual	99
4	Implementation and Experimental Setup	101
4.1	Implementation	102
4.1.1	Pilot Implementation	103
4.1.2	Game Description	104
4.1.2.1	Context	109
4.1.2.2	Domain-dependent Knowledge	110
4.1.2.3	Soft Goals	115
4.2	Experimental Setup and Considerations	118
4.2.1	Updating Somatic Marker Preferences	119
4.2.2	Emotionality	122
4.2.3	Measures of Effectiveness	123
4.2.3.1	Behaviour as a Measure of Effectiveness	124
4.2.3.2	Personal Reward as a Measure of Effectiveness	124
4.2.3.3	Individuality as a Measure of Effectiveness	126
4.2.4	Cases and Modes Constructed	133
4.2.4.1	Modes	134
4.2.4.2	Cases	135
4.2.5	Expected results	137
4.2.5.1	Expected Behaviour	138
4.2.5.2	Expected Personal Reward	140
4.2.5.3	Expected Individuality	141
4.3	Summary of Implementation and Experimental Setup	143
5	Results and Discussion	145
5.1	Behaviour	146
5.1.1	Behaviour Over Time (Research Sub-question 1a)	146
5.1.1.1	Comparison to Expectations	151

CONTENTS

5.1.2	Learning A Functional Soft Goal (Research Sub-question 1b)	152
5.1.2.1	Comparison to Expectations	153
5.1.3	Behaviour in Different Contexts (Research Sub-question 2a)	154
5.1.3.1	Comparison to Expectations	159
5.1.4	Chi-squared Test (Research Sub-question 3a)	160
5.1.4.1	Comparison to Expectations	160
5.1.5	Summary of Behaviour Results	160
5.2	Personal Reward	160
5.2.1	Reward Compared to Random Choice (Research Sub-question 1c)	162
5.2.1.1	Comparison to Expectations	162
5.2.2	Effect of Adaptation and Context on Reward (Research Sub-questions 1d and 2b)	162
5.2.2.1	Reward Across Modes	162
5.2.2.2	Reward Across Cases	163
5.2.2.3	Summary of Personal Reward Results	164
5.2.2.4	Comparison to Expectations	164
5.3	Individuality	165
5.3.1	Effect of Adaptation and Context on Individuality (Research Sub-questions 1d and 2b)	167
5.3.1.1	Individuality Across Modes	167
5.3.1.2	Individuality Across Cases	168
5.3.1.3	Summary	169
5.3.1.4	Comparison to Expectations	169
5.3.2	Individuality per Character (Research Sub-question 3b)	169
5.3.2.1	Comparison to Expectations	170
5.4	Discussion	171
5.5	Extra Results: Domain-dependent Measures	173
5.5.1	Happiness	174
5.5.2	Friendships	175
5.6	Summary	176

6 Conclusion	179
6.1 Addressing the Research Questions	179
6.1.1 Research Question 1: Adaptation	179
6.1.2 Research Question 2: Context	181
6.1.3 Research Question 3: Individuality	182
6.1.4 Implications Arising from Research Questions	184
6.2 Potential for Future Research	184
6.2.1 Personality Templates and Input Parameters	185
6.2.1.1 Testing Adaptation Further	186
6.2.1.2 Designing Specific Characters	186
6.2.2 Model	187
6.2.3 Visualisation of Personality, Context and Friendship Networks	188
6.2.4 Qualitative Testing	189
6.2.5 The Personality Guessing Game	190
6.3 Implications for Games	191
6.4 Final Words	192
 Appendices	
A Behaviour Over Time	193
B Learning A Functional Soft Goal	199
C Domain-dependent Results	205
C.1 Happiness	205
C.2 Friendships	206
C.2.1 Case 1: Clear Preference Against One Activity	209
C.2.2 Case 2: Multiple Ways to Achieve Goals	211
C.2.3 Case 3: Conflicting Goals	213
C.2.4 Case 4: Complex Soft Goal Personality	215
C.2.5 Case 5: Different Soft Goal Personalities	217
 References	 219

CONTENTS

List of Figures

1.1	Battlestar Galactica: Model Eight	5
2.1	Generic goal/plan hierarchy	22
2.2	Standard BDI execution cycle	24
2.3	Reciprocal determinism in cognitive-social theories	33
3.1	Agent components and beliefs	74
3.2	Possible partial goal/plan hierarchy from the village domain. Part A. . .	78
3.3	Possible partial goal/plan hierarchy from the village domain. Part B. . .	79
3.4	Expanded version of the BDI execution loop	83
3.5	Adaptation process	83
3.6	The adaptation process with a generic goal/plan hierarchy	85
3.7	Effect of adaptation processes on components and beliefs	86
3.8	What makes an individual an individual in our model	100
4.1	Three characters from the game	105
4.2	Implemented domain-dependent goal/plan hierarchy used for all characters	106
4.3	Screenshot from the implemented game	107
4.4	Screenshot from the implemented game	107
4.5	Screenshot from the implemented game without a player included	108
4.6	Flowchart to update attraction values the character has insulted someone	113
4.7	Flowchart to update attraction values when the character has been told an insult	114
4.8	Extract of sample reward output	119
4.9	Extract of sample somatic marker output	120
4.10	Two examples of how characters can be the <i>same</i>	129

LIST OF FIGURES

4.11	Four examples of how characters can be <i>different</i>	130
5.1	Sample character graphs for Case 1 of behaviour based on the individual	148
5.2	Sample character graphs for Case 2 of behaviour based on the individual	148
5.3	Sample character graphs for Case 3 of behaviour based on the individual	149
5.4	Sample character graphs for Case 4 of behaviour based on the individual	150
5.5	Sample character graphs for Case 5 of behaviour based on the individual	150
5.6	Anna's Behaviour when moving for Case 4	154
5.7	Bec's sample character behaviour graphs for two contexts in Case 1 . . .	156
5.8	Chloe's sample character behaviour graphs for two contexts in Case 2 .	157
5.9	Deb's sample character behaviour graphs for two contexts in Case 3 . .	158
5.10	Gina's sample character behaviour graphs for two contexts in Case 4 . .	159
A.1	Case 1 behaviour based on the individual	194
A.2	Case 2 behaviour based on the individual	195
A.3	Case 3 behaviour based on the individual	196
A.4	Case 4 behaviour based on the individual	197
A.5	Case 5 behaviour based on the individual	198
B.1	Anna's Behaviour when moving for Case 4	200
B.2	Bec's Behaviour when moving for Case 4	200
B.3	Chloe's Behaviour when moving for Case 4	201
B.4	Deb's Behaviour when moving for Case 4	201
B.5	Elle's Behaviour when moving for Case 4	202
B.6	Fran's Behaviour when moving for Case 4	202
B.7	Gina's Behaviour when moving for Case 4	203
B.8	Heidi's Behaviour when moving for Case 4	203
C.1	Happiness and reference reward for two example runs for two example characters	206
C.2	Friendship network and attraction levels for Case 1 'normal'.	208
C.3	Friendship network and attraction levels for Case 2 'normal'.	210
C.4	Friendship network and attraction levels for Case 3 'normal'.	212
C.5	Friendship network and attraction levels for Case 4 'normal'.	214
C.6	Friendship network and attraction levels for Case 5 'normal'.	216

List of Tables

1.1	Criteria for success to be used to evaluate testing-based research sub-questions	16
2.1	Aspects Influenced by Personality in IVA Applications	53
4.1	Calculation of friendship points for two characters.	116
4.2	Cases used with soft goal personality templates	136
4.3	Expected results for criteria for success related to behaviour	138
5.1	Results for criteria for success based on behaviour	160
5.2	Average reward values for test Cases.	161
5.3	Significant differences for reward	161
5.4	Individuality for test Cases	166
5.5	Average number of characters who are different from the majority . . .	171
5.6	Summary of results of testing-based research sub-questions based on criteria for success	177

GLOSSARY

Glossary

Action An *activity* or a *plan*. A set of steps that the *character* can execute that are visible to a *player*.

Activity (See also *top-level activity*) A collection of plans and sub-goals that are used to achieve a high level plan that constitutes doing “something” within the domain. Once finished, success can be determined using *evaluation*. For instance, one activity could be “make something”. A plan within that activity could be “make bread”.

Adaptation The process by which behaviour changes over time. It is a simple form of learning based on experience. In our model this is done using a form of self-reinforcement via reinforcement learning, in particular the reinforcement comparison technique, see Section 3.2 (page 82).

Adaptive Can change behaviour over time. Uses *adaptation*.

Agent The reasoning part of a *character* (as compared to the visual aspects of a character). That is, the part that decides what to do and how to evaluate itself.

Appraisal of (coping) choices Used synonymously with decision-making, see Section 3.2.1 (page 87). The way that appraisal of choices is used within our model matches to secondary (not primary) appraisal in the cognitive appraisal model (Lazarus, 1991) (see literature survey Section 2.1.1.3, page 27). Reappraisal is implemented as *evaluation*.

BDI Beliefs, Desires, Intentions. Used in reference to the BDI paradigm that agents are embodied in their virtual world (see literature survey Section 2.1.1.1, page 21). Agents can hold beliefs or knowledge about their world and have desires of what they would like to achieve. An agent’s intentions are a list of the current *plans* it is using to achieve its desires. Intentions should be non-conflicting.

Behaviour Manner of acting; the observable actions and reactions of a person. Always considered over a specific time period. Behaviour is used as a *measure of effectiveness* of the model to test *criteria for success* (see Section 4.2.3.1, page 124). Measured in our implementation by counting the number of times characters choose different *actions* (*activities* and *plans*) over a fixed output time period.

Beliefs From the BDI paradigm (see literature survey Section 2.1.1.1, page 21), knowledge (*facts* and subjective *opinions*) that the *agent* has or stores about others and the environment.

Case A scenario used for testing. In each Case, the characters are given different *soft goal personality templates*. See Section 4.2.4.2 (page 135) for a listing of Cases.

Character The visual appearance combined with an *agent*, i.e. what a game player sees.

GLOSSARY

Conflicting goals *Goals* that cannot both be achieved at the same time. For example, you cannot have your cake and eat it too.

Context The perceived current situation of a character. Current level of achievement of soft goals that the agent is pursuing, regardless of importance, see Section 3.1.1.2 (page 75). This value is not based on the agent's history, it is simply what the agent is achieving now. For each soft goal the agent is pursuing, the achievement level is converted to a single letter representing: high (close to achieving this goal), medium and low (this goal is not being achieved well currently). The letters are combined based on the alphabetisation of the soft goal name to create the context. For example, if the soft goals are "have friends" and "have money", the context "LH" represents the state where the agent has hardly any friends and a lot of money.

Context-aware Characters who's behaviour depends on their perceived situation. In our model this is, behaviour based on knowledge of current *context*, in terms of the *soft goals* the agent is trying to achieve.

Coping According to the cognitive appraisal model of emotions (Lazarus, 1991) (see literature survey Section 2.1.1.3, page 27) coping is a mechanism that we engage in to improve our overall emotional wellbeing. Coping can be physical actions in the real world, such as running away when scared, or an emotional re-evaluation, such as realising there is no need to be scared in the first place. In our model, coping refers to the domain-dependent *plans* that the agents can use to act within the world in order to improve their overall wellbeing based on achieving their *soft goals*.

Criteria for Success The minimal set of tests (shown in Table 1.1, page 16) that must be satisfied for the implemented model to be considered to have addressed the testing-based *research sub-questions*. The criteria test whether the characters generated by the model are *adaptive*, *context-aware* and *individual*. The criteria are measured based on the quantitative values of: *behaviour*, *reward* and *individuality*.

Domain-dependent Knowledge Beliefs (facts and opinions) specific to the implemented domain. In the theoretical model, they are used to calculate achievement levels of soft goals, and therefore individual soft goal rewards, which leads to personal reward. In our implemented domain, the beliefs are opinions - happiness, attraction towards others and facts - attraction from others, location, insults said and told.

Emotion Related to feelings. There are many types of emotion, such as happiness, anger, fear etc... In our model we use the term to relate to any or all of these types. In our implementation we use the term to refer to a happy/sad scale that represents how close the agent is to achieving all their *soft goals*.

Emotionality A set of values that represent how the agent reacts to events, the thresholds above or below which they define "good" and "bad", and other learning related parameters, see Section 3.1.3.3 (page 81). Used in a similar sense to Ortony (2002) (see literature survey Section 2.1.2.3, page 35).

Evaluation The process by which *personal reward* is determined, see Section 3.2.2 (page 90). During this process the following are updated: *achievement levels*, *context*, *somatic markers* and *emotion*. This process is an implementation of reappraisal according to the cognitive appraisal

model of emotions (Lazarus, 1991) (see literature survey Section 2.1.1.3, page 27) and occurs after every *activity* has been completed.

Execution of a plan *Plans* are like simple functions, they are executed consecutively line by line. The execution of a plan is simply following the steps in the plan and ensuring that none of the individual steps fail.

Facts *Beliefs* that cannot be changed by an individual agent (i.e. different from *opinions*). It is a belief that is based on information given to the character by the environment or from other characters.

Goal/Plan Hierarchy From the BDI paradigm (see literature survey Section 2.1.1.1, page 21). In particular, see the generic figure for the goal/plan hierarchy, Figure 2.1 (page 22). The goal/plan hierarchy is a representation of *hard goals* and the *plans* that can be used to directly achieve these goals. The hierarchy begins with a hard goal placed at the top. Underneath this are a number of plans that can achieve this goal. Each of these plans can post a number of sub-goals that each must be achieved for the plan to succeed. This leads to a hierarchy, for example domain-dependent hierarchies see Figure 3.2 (page 78) and Figure 4.2 (page 106).

Goals Something that an agent wants to achieve or maintain. There are a number of different types of goals in the literature (see Section 2.1.1.2, page 24). In our model, we use *hard goals* and *soft goals*. Hard goals are implemented in the *goal/plan hierarchy*. Soft goals are part of an agent's *personality template* in the form of an agent's *soft goal personality*. Soft goals can be conflicting, whereas hard goals cannot. Agents are given no knowledge of how to achieve soft goals, but achievement of hard goals is explicit within the goal/plan hierarchy.

Hard goals Concrete goals within an explicit *goal/plan hierarchy*. The designer must explicitly state how an agent can achieve these goals. For example, to achieve the goal “make something”, an agent can choose a plan such as “make bread”.

Importance Used for *soft goals*, also known as *weight*. This is a number on a scale of $[0, 1]$, where 1 represents a soft goal that the agent really wants to achieve, and 0 is one they do not care whether they achieve or not. It is part of *soft goal personality*, see Section 3.1.3.2 (page 80).

Individual Different from others based on observable *behaviour*. A property a character can possess. In our model, an individual character is comprised of a number components and beliefs. Primarily the components and beliefs include *personality template*, *somatic markers* and *domain-dependent knowledge*. The extent to which a character is individual is measured using *individuality*.

Individual Soft Goal Reward The *reward* for a single *soft goal* based on the *soft goal achievement level* and the distance to the ideal soft goal value from an agent's *soft goal personality*, see Section 3.2.2.1 (page 91).

Individuality In the general sense, individuality is what makes each of us unique and different from other people. To test our model we needed a quantitative measure of individuality to compare characters to each other based on their patterns of behaviour over the entire running time. The quantitative measure of individuality is a count of the number of differences between characters based on whether the *behaviour* (action choices) are significantly different for the *top-level activities* (see Section 4.2.3.3, page 126). Individuality is used as a *measure of*

GLOSSARY

effectiveness of the model to test *criteria for success*. Individuality can relate to a specific character (number of characters that character is different from) or to an entire run (total number of characters different from each other).

Learning Feedback Loop Also known as the *adaptation* process. After completing an *activity* the agent evaluates the *personal reward* for the activity and then feeds this back into the point where the decision was made to do that activity, see Section 3.2 (page 82). In our model, feedback is based on *personal reward* and updates *somatic marker* preferences.

Measure of Effectiveness Observed data from testing that is used to determine whether *criteria for success* are satisfied. Three measures are used: *behaviour* (based on a count of the number of times characters chose *actions*), *reward* (based on *personal reward* calculations for characters), individuality (a count of differences between characters based on *behaviour*).

Mode A scenario used for testing. The different modes cause the characters to use random choice when making decisions ('adaptation off'); or do not allow the characters to distinguish between *contexts* when they are learning ('context off'); or using the full model (from Chapter 3) where characters use the methods specified ('normal'). See Section 4.2.4.1 (page 134) for further explanation of modes used for testing.

Opinion A *belief* that has a value judgement attached. For example, a character can store "I like Anna a lot". The values on opinions can be changed by the character, if so desired.

Past Experience A lookup table of preferences based on past rewards. See *somatic markers*.

Personal Reward Also known as self-reinforcement value. The agent's personal evaluation of how "good" it thinks the last *activity* was. This represents how close the agent is to achieving all of its *soft goals*, with more importance placed on different goals according to the agent's *soft goal personality* (see calculation step in Section 3.2.2.2, page 93). Reward is used as a *measure of effectiveness* of the model to test *criteria for success*.

Personality Personality is the set of observable characteristics that make an individual themselves. In our model, we restrict the term to relate to observable behaviour. A character's final personality is a combination of their initial *personality template* as well as their learnt preferences, or *somatic markers*.

Personality Template In the general sense a personality template represents the basic genetic set-up of an individual, actual personality emerges through life experience. In our model, a personality template is made of three components: a domain-dependent *goal/plan hierarchy*, a *soft goal personality* and *emotionality* values. See Section 3.1.3 (page 78).

Plan A set of instructions or recipe that the agent can execute. A plan can result in observable *actions* in the virtual world or can change its *beliefs*.

Player A human participant in the game or simulation.

Preference Value See *somatic marker* preference.

Reference Reward A value that is representative of all past *personal rewards*. Used to determine whether a personal reward for a particular activity was "good" or "bad" compared to all other activities that have been executed. To see how the reference reward is used see Section 3.2.2.3 (page 94); to see how reference reward is updated see Section 3.2.2.5 (page 97).

Reinforcement Comparison Technique A simple *reinforcement learning* technique (from Sutton & Barto (1998)) that compares the current *reward* received with all other rewards that the agent has received, using a *reference reward*, see literature survey Section 2.1.4.2, page 40). The technique specifies how to update decision-making selection policy. In our model, it is used to update somatic markers, see Section 3.2.2.3 (page 94).

Reinforcement learning Learning that is based on maximising reward from an external agent based on trial and error, i.e. punishment and reward, see literature survey Section 2.1.4.2 (page 38). In our model, we use learning based on self-evaluation (*personal reward*), not an external agent, and use the reinforcement comparison technique to update selection policy (*somatic markers*).

Research Questions For this thesis, the three research questions relate to developing, implementing and testing a model of personality that is *adaptive*, *context-aware* and *individual*. Introduced in Section 1.2 (page 14).

Research Sub-questions Detailed questions that breakdown the *research questions* into smaller parts. The sub-questions are also divided into model-based (relating to the development of the personality model) and testing-based (relating to how to determine success of the implemented model). The testing-based sub-questions are considered answered when the *criteria for success* have been satisfied. Introduced in Section 1.2 (page 14).

Reward A measure of merit of an activity, i.e. “good” or “bad”. In our model there are three types of reward: *individual soft goal reward* which is reward for a specific soft goal based on its achievement level; *personal reward* which is based on combining all individual soft goal rewards based on an agent’s *soft goal personality* template (used as a *measure of effectiveness* of the model); and *reference reward* which is a running average of all previous personal rewards.

Soft goal equations The domain-dependent functions used to determine the *achievement levels* of the soft goals. For a description of how they are used generically see Section 3.1.1.1 (page 74); for the specific soft goal equations used in our implementation see Section 4.1.2.3 (page 115). For example, if the soft goal is “have friends”, the equation to determine the quantitative achievement level could be based on such beliefs as: number of people the agent likes, number of people who like the agent, or a combination of these beliefs.

Soft goal personality The soft goals this agent is trying to achieve, the *importance* levels it places on the goals, and the ideal value of each goal, see Section 3.1.3.2 (page 80). Once set for an agent, this will not change. For example, in one domain the agents may be able to have the soft goals “not be hungry” and “have money”. One possible soft goal personality is that the agent places a high *importance* on “not be hungry” and a medium importance on “have money”. Importance is a number on a scale of $[0, 1]$, where 1 represents a soft goal that the agent really wants to achieve, and 0 is one they do not care whether they achieve or not. For each of the soft goals that the agent is trying to achieve, an ideal or maximum value is specified. For example, one agent may consider “have money” achieved when they have \$100,000, another may believe they need \$1 million.

Soft goals Soft goals are a set of potentially *conflicting goals* that the agent is attempting to achieve at every step, see Section 3.1.1 (page 73). Plans may contribute partially to achieving a number of different soft goals. Some examples of soft goals are: have friends, have money and not being hungry. Agents do not initially know how to achieve their soft goals, they must learn.

GLOSSARY

For example, if the soft goal is to “have friends”, the designer developed domain-dependent *goal/plan hierarchy* does not need to have a plan that can directly achieve this, i.e. there is no need for a plan called “make friends”. Plans such as “interact” or “give away something” may improve the achievement of the soft goal, but this can only be learnt based on feedback from trial and error.

Sub-plan Some plans require more *hard goals* to be posted to finish the plan. These goals will be handled by sub-plans. These are plans that are beneath the originating plan in the *goal/plan hierarchy*. For example, to achieve “make something” the goal “choose what to make” is achieved by implementing a sub-plan such as “make bread”, see Figure 3.2 (page 78).

Top-level Activity A *plan* that is very high up in the *goal/plan hierarchy*. It should be something that has a long enough duration that the agent’s domain-dependent beliefs will have changed and the agent can perform an *evaluation* of what has happened. For example, it would be difficult to perform an evaluation after a small step, such as choosing who to talk to. A suitable top-level activity would be a longer interaction, such as actually having an entire conversation with a character. In our implementation, the agents have three top-level activities: “move”, “insult” and “wait”, see Figure 4.2 (page 106).

Weight Used for *soft goals*, also known as *importance*. This is a number on a scale of $[0, 1]$, where 1 represents a soft goal that the agent really wants to achieve, and 0 is one they do not care whether they achieve or not. It is part of soft goal personality, see Section 3.1.3.2 (page 80).

Abstract

Personality is a key component of characters that inhabit immersive virtual environments, such as games and virtual agent applications. In order to be distinguishable from other characters in the environment, each character should have its own personality in the form of different observable behaviour, not solely in its physical appearance or animation. Previous work in this field has mostly relied on time-consuming, handcrafted characters and static, trait-based approaches to personality. Our goal is a method to develop complex, individual personalities without handcrafting every behaviour. Unlike most implemented versions of personality theories, cognitive-social theories of personality address how personality is developed and adapts throughout childhood and over our lifetimes. Cognitive-social theories also emphasise the importance of situations in determining how we behave. From this basis, we believe that personality should be individual, adaptive, and based on context. Characters in current state-of-the-art games and virtual environments do not demonstrate all of these features without extensive handcrafting.

We propose a model where personality influences both decision-making and evaluation of reward. Characters use their past experiences in the form of simple somatic markers, or gut-instinct, to make decisions; and determine rewards based on their own personal goals, rather than via external feedback. We evaluated the model by implementation of a simple game and tested it using quantitative criteria, including a purpose-designed individuality measure. Results indicate that, although characters are given the same initial personality template, it is possible to develop different personalities (in the form of behaviour) based on their unique experiences in the environment and relationships with other characters. This work shows a way forward for more automated development of personalities that are individual, context-aware and adapt to users and the environment.