The μ-SIC System:
A Connectionist Driven Simulation of Socially Interactive Agents

Brian Mac Namee & Pádraig Cunningham

August 20, 2002

1 Introduction

The success of games such as *The Sims* (thesims.ea.com) and *Black & White* (www.bwgame.com) have shown that there is a demand for the personalities, moods, and relationships of Non Player Characters' (NPCs) to be made the focus of game-play. In order for this shift of focus to take place, agent architectures used to create NPCs must be augmented with models of these aspects of a character's persona which must then be used to drive characters' behaviour.

Psychology offers a number of quantitive models of personality, mood and inter-personal relationships which can be used to capture these important aspects of a character's persona. In order to use these models to drive character behaviour we can turn to connectionist AI techniques, and in particular Artificial Neural Networks (ANNs). This paper will describe the μ-SIC system which does just this.

The purpose of the μ-SIC system is to choose which social interactions characters should engage in when placed within a virtual environment with other characters. When a moment within a simulation arises where a character is free to engage in an interaction, the μ-SIC system is queried with the character's personality and mood details, and their relationship details to each of the other characters in the same location who are also available for interaction. From these queries a particular interaction with a particular character is chosen.

This paper will begin with a short overview of a larger project of which the μ-SIC system is a part. Following this, a description of the psychological models used by μ-SIC will be given. The actual implementation details of the system will be described next, along with a short description of a simulated situation which uses the μ-SIC system. Finally, a discussion of the benefits and drawbacks of μ-SIC will be given, along with some pointers as to how the system can be improved.

2 Project Overview

Although games are becoming ever more engaging, there is a trend in current adventure and role-playing games for the behaviour of computer controlled NPCs to be very simplistic. Usually, no modeling of NPCs is performed until the player reaches the location at which an NPC is based. When the player arrives at this location, NPCs typically wait to be involved in some interaction, or play through a pre-defined script which can lead to very predictable, and often jarring behaviour. In order to overcome these limitations new models are required for implementing game characters.

Although such models have not been widely used in computer games, a number of architectures for creating realistic characters have been developed. For example, work led by Thalmann [1] and the Oz project [8] based on interactive drama have both developed virtual human architectures. As part of the TCD Game AI Project [3] the Proactive Persistent Agent (PPA) [7] architecture is being developed for the creation of NPCs which overcome the limitations typically associated with computer game characters.

Agents based on the PPA architecture are *proactive* in the sense that they can take the initiative and follow their own goals, irrespective of the actions of the player. *Persistence* refers to the fact that at all times, all NPCs in a virtual world are modeled at least to some extent, regardless of their location relative to that of the player.

This paper will focus on the PPA architecture's social unit (implemented using the μ-SIC system) which is used to drive characters' social behaviour and maintain their relationships with both players and other NPCs.

3 Using Psychology to Model NPCs' Personas

This section will describe quantitive models taken from psychology which are used to model NPCs' personalities, moods and relationships. However, before discussing the models used, it is worth taking a moment to discuss the criteria used for selecting suitable models.

The first selection criterion worthy of note is that the models chosen need not necessarily represent the current state of knowledge in cognitive science in all its aspects. Our goal is to create characters which behave plausibly at all times within a simulation, so models which achieve this are enough.

The second important criterion for model selection is that the models used should be as simple as possible. In order for game designers to successfully use the PPA architecture to place characters within their games, the models involved must be simple enough so that the designer can understand how they work, and more importantly how changing a model's parameters might affect a character's behaviour.

In addition to the concern for usability, any system for use in games must be efficient both in terms of memory usage and computation required.

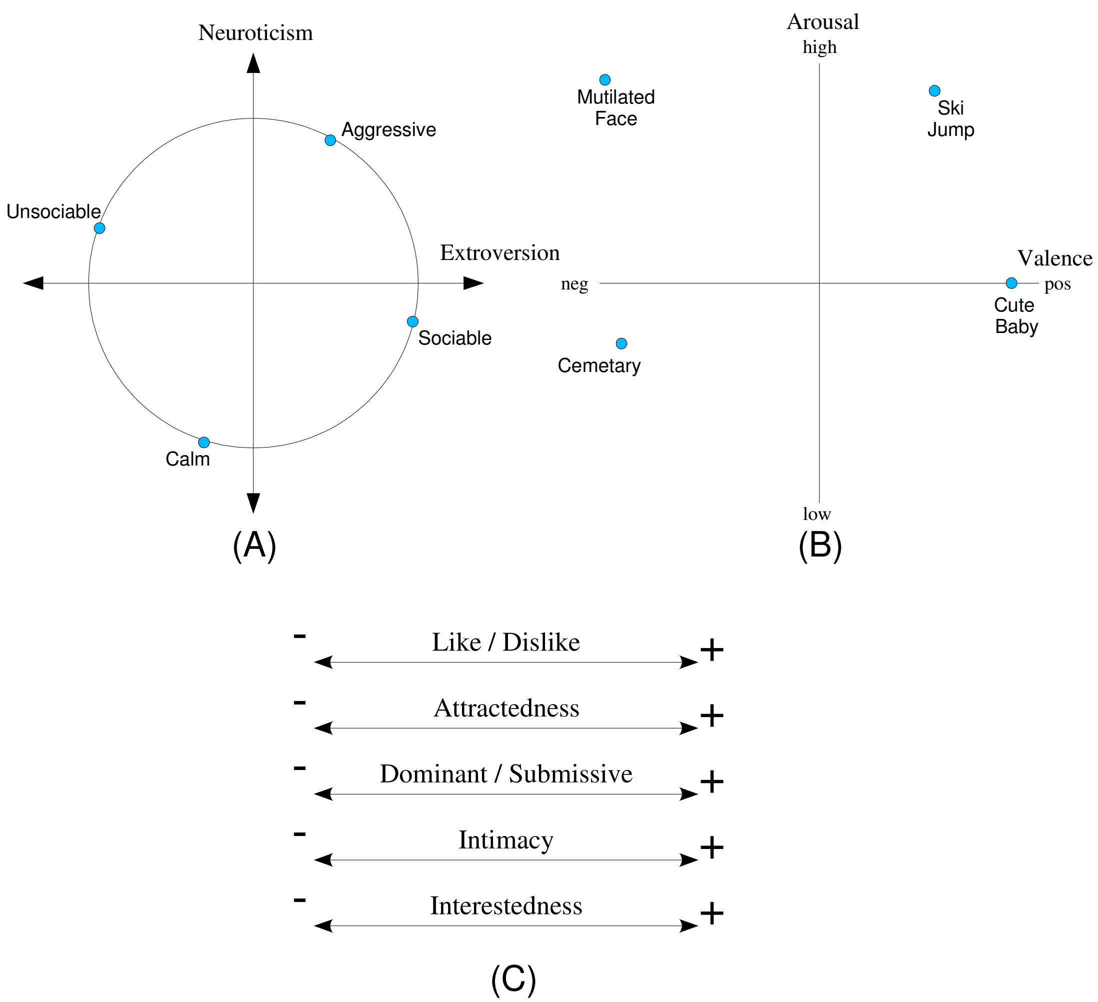

Figure 1: The psychological models used within the μ-SIC system. (A) The Eysenck personality model which measures personality across the Introvert-Extrovert and Neuroticism-Stability axes. (B) The Lang mood model which plots mood according to valance and arousal. (C) The relationship model used which plots a character's relationship to another character.

3.1 Personality Model

The first important factor of an NPC's persona which needs to be modeled is personality, which will allow the creation of characters with personality types, such as aggressive, sociable, moody etc. From the whole myriad of psychological models of personality available we have chosen Eysenck's two dimensional classification of personality [2].

The Eysenck model plots a character's personality across two orthogonal axes, *introversion-extroversion* and *neuroticism-stability*. From [6], the extrovert is said to be sociable, impulsive and open to new experiences, while the introvert is quiet, serious and prefers solitary experiences. The neurotic is contrasted with a stable person by suffering from tension and interpersonal difficulties. An illustration of the model, which shows the positions of a number of the possible personality types, is shown in figure 1 (A).

It is worth noting that psychologists generally accept that two axes is not enough to accurately model the whole gamut of human personality types. Currently the most sophisticated models, such as the OCEAN [9] personality model, operate across five axes. However, the use of more axes was deemed overly complex for the purposes of game simulation, and the Eysenck model was chosen as it remains one of the most respected and well established personality models in psychological theory [6].

3.2 Mood Model

The second psychological model used, simulates a character's mood as it changes over time through interactions with other characters or players. Again, simplicity is key and a model (shown in figure 1 (B)) which works across two axes has been chosen. An agent's mood is measured according to *valance* and *arousal*, where valance refers to whether the mood is positive or negative, and arousal refers to the intensity of the mood.

This model has been used in computing applications before [11], and is originally due to Lang [4]. Over the course of Lang's work, this model was used in experiments wherein subjects were shown a number of pictures with their reactions to these pictures plotted according to the two axes. Some of these reactions are shown in figure 1 (B).

3.3 Relationship Model

The third model we use (shown in figure 1 (C)) simulates agents' relationships with each other and players. The model has been used in a number of other entertainment projects, namely the Oz Project [13], TALE SPIN [10], and UNIVERSE [5], and has its psychological basis in the work of Wish et al [14]. Traditionally, four values are used to characterise the relationship of one character to another. These are the amount that a particular character likes another character, how physically attracted one character is to another, whether the characters are dominant or submissive towards each other and how intimate the characters are.

To facilitate conversation, we have augmented this model with a value indicating how interested one character is in another. Conversation within the μ-SIC system is based on a very simple model in which each character has a list of subjects in which they are interested. When characters engage in a conversation they simply pass these subjects back and forth. Thus, characters are interested in one another if they share a number of common subjects of interest.

4 Implementing the μ-SIC System

In order to use the psychological models just described to drive social behaviour, we need a technique which can take the current values of these models, and determine whether an interaction should be started, and if so which one. An ANN has been chosen to perform this task.

ANNs [12] are a class of machine learning technique which is based on the manner in which neurons in biological brains operate. ANNs can be used to perform classification tasks in which a set of inputs describing a particular problem case are presented to the network, which then outputs its class.

The structure of the ANN used within the μ-SIC system is shown in figure 2. The network's input layer has nodes for the personality and mood of the character who is attempting to instigate an interaction, and their relationship to the current character being considered for interaction. The output layer has nodes for each of the possible interactions which the characters can engage in.

Before an ANN can be used to perform classification, it must be trained to recognise the different classes involved. Training a network involves presenting a number of known examples of the problem case to the network and adjusting the network's internals based on how well the network can recognise these training examples.

For training, a data set describing the problem space must be acquired. Data acquisition is often a difficult problem, and is particularly so for the μ-SIC system, as there are no databases available which contain information on how people interact. For this reason, an artificial data set was created. A number of simulation situations were created and populated with characters whose personalities were set using the Eysenck model. Relationships between these characters were then initialised and a group of people determined which interactions these characters would engage in as their moods changed over time.

Based on this initial data set (consisting of approximately 100 data elements) a number of interaction exemplars (data items considered to be particularly fine examples of when an interaction would take place) were identified. These exemplars were used to determine the ranges of each input value which would cause each possible interaction. Using these ranges, a set of 2000 random data points covering the set of possible interactions was created.

To determine the accuracy of the network a five-fold cross validation was performed in which the network achieved an accuracy of 85%, indicating that the output of the network was consistent and coherent. Further to

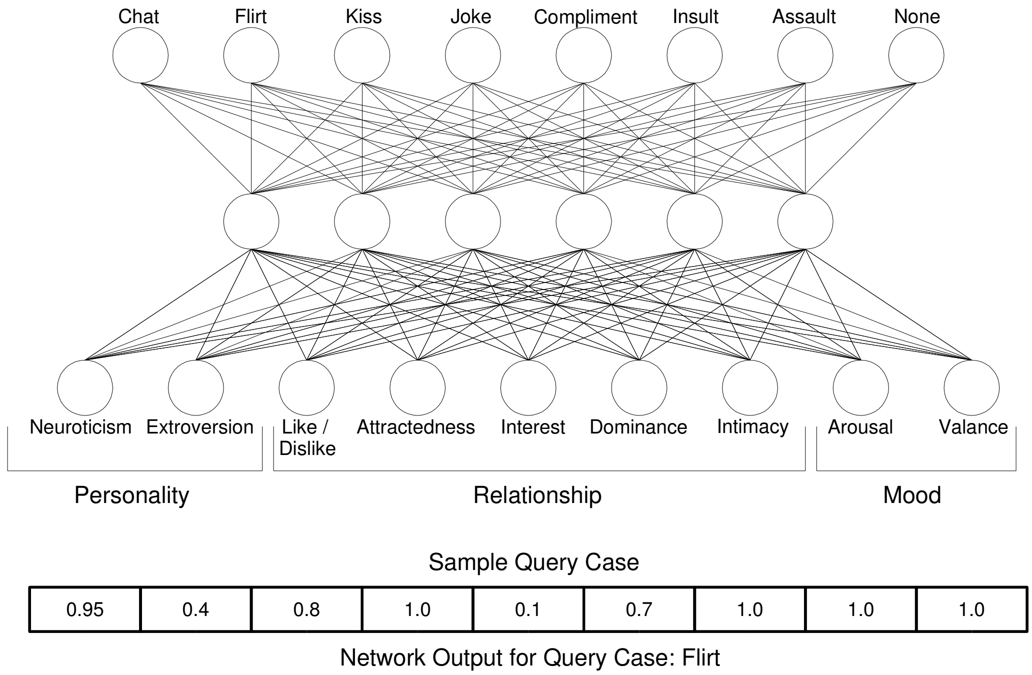

Figure 2: The structure of the ANN used to drive NPCs' social behaviour.

this high accuracy, when the system produces incorrect predictions these are rarely significantly incorrect. For example, the system may produce a *chat* interaction rather than a *joke* interaction, but will never produce an *assault* interaction instead of a *kiss* interaction.

Only one copy of the μ-SIC system is stored within the game engine, with NPCs querying this each time they are free to begin a new interaction. For this reason the system can be considered an oracle that advises NPCs on how to behave (see figure 3).

To determine the success of the μ-SIC system a simulation example has been constructed in which a number of characters have been placed within a bar environment, free to interact with one another. A screen-shot of this simulation is shown in figure 3. The simulation successfully demonstrates the full range of possible interactions and how relationships between the characters within the simulation evolve as the simulation progresses.

5 Conclusions

The purpose of this work is to develop a system which can be used within a larger agent architecture to allow NPCs within computer games perform social interactions with other NPCs or players, based on their personalities, moods and inter-personal relationships. The system achieves this by simulating these aspects of a character's persona using quantitive models from psychology. These models are used as inputs to an ANN which determines which interactions the character should engage in, with which other characters. This ANN has been trained with a data set generated from a small set of hand coded interactions. The μ-SIC system successfully performs a comprehensive range of social interactions based on the data set produced, and a simulation example has been created to demonstrate this.

Although the system is quite successful in its present state, one addition to the system has been identified which could improve the system considerably. At present characters engage in simple interactions wherein one character performs an interaction, and the other character reacts, thus ending the interaction. In order to more accurately model the cut and thrust of conversation, an extra input node indicating the previous interaction which the characters were involved in could be added to the network.

In this way context would be explicitly added to the interaction model, allowing interaction sessions to evolve through different interaction modes. So, for example, two characters might start by chatting, find they have little in common and so start to insult each other, and finally end their interaction by assaulting one another. Although this is possible in the current system it would be spread across a number of interaction sessions. The major drawback to this extension to the model is that an order of magnitude more data would be required for training. As previously discussed, data acquisition is difficult although the techniques discussed in section 4 could be used.

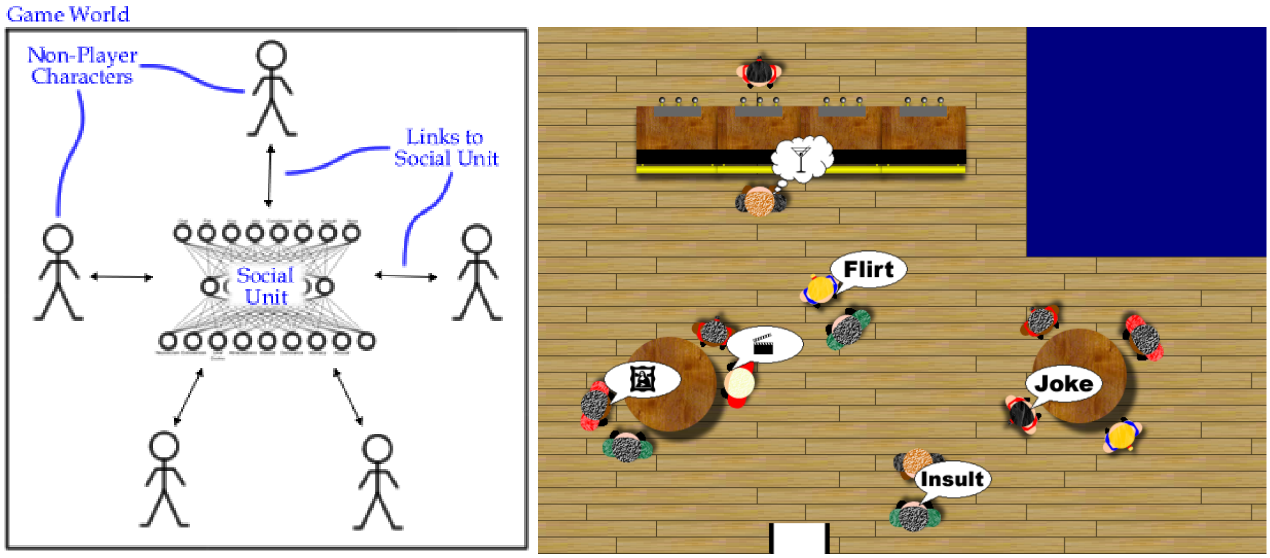

Figure 3: An illustration of how the μ-SIC system is incorporated into a virtual world, and a screen-shot of the simulation situation.

References

[1] A. Caicedo & D. Thalmann, "Virtual Humanoids: Let them be autonomous without losing control", The Fourth International Conference on Computer Graphics and Artificial Intelligence, 2000.

[2] H.J. Eysenck & S. Rachman, "The Causes and Cures of Neuroses", London: Routledge and Kegan Paul. 1965.

[3] C. Fairclough, M. Fagan, B. Mac Namee & P. Cunningham, "Research Directions for AI in Computer Games", Proc. of the 12^{th} Irish Conference on AI and Cognitive Science, 2001.

[4] P. J. Lang, "The Emotion Probe: Studies of motivation and attention", A study in the Neuroscience of Love and Hate. Hillsdale, NJ: Lawrence Erlbaum Associates, Publishers, 1995.

[5] M. Lebowitz, "StoryTelling as Planning and Learning", In Poetics. Vol. 14. No. 6. December 1985.

[6] P. Lloyd, A. Mayes, A.S.R. Manstead, P.R. Meudell & H.L. Wagner, "Introduction to Psychology: An Integrated Approach", Fontana Paperbacks, 1984.

[7] B. Mac Namee & P. Cunningham, "Proposal for an Agent Architecture for Proactive Persistent Non Player Characters", Proc. of the 12^{th} Irish Conference on AI and Cognitive Science, 2001.

[8] M. Mateas, "An Oz-Centric Review of Interactive Drama and Believable Agents", Technical Report CMU-CS-97-156, School of Computer Science, Carnegie Mellon University, Pittsburgh, PA. 1997.

[9] R. R. McCrae, & P. T. Costa Jr., "Toward a new generation of personality theories: Theoretical contexts for the five-factor model" In J. S. Wiggins (Ed.), The five-factor model of personality: Theoretical perspectives (pp. 51-87). New York: Guilford. 1996.

[10] J. Meehan, "The Metanovel: Writing Stories by Computer", Research Report #74. Computer Science Department, Yale University. New Haven, CT. 1976.

[11] R. W. Picard, "Affective computing", Perceptual Computing TR 321, MIT Media Lab, 1995.

[12] S. Russell & P. Norvig, "Artificial Intelligence: A Modern Approach", Prentice Hall, 1995.

[13] W. Scott Neal Reilly, "Believable Social and Emotional Agents", Ph.D. Thesis. Technical Report CMU-CS-96-138, School of Computer Science, Carnegie Mellon University, Pittsburgh, PA. 1996.

[14] M. Wish, M. Deutsch & S. Kaplan, "Perceived Dimensions of Interpersonal Relations". In Journal of Personality and Social Psychology. Vol. 33. No. 6. April 1976.