Action Selection Through Affective States Modelling

Christopher J. Headleand
School of Computer Science,
University of Lincoln,
Lincoln, UK,
cheadleand@lincoln.ac.uk

William Teahan
Llyr Ap Cenydd
Computer Science Dept.
Bangor University,
Bangor, UK,
{llyr.ap.cenydd, w.j.teahan}@bangor.ac.uk

***Abstract*—We introduce an action selection framework for the advanced behavioural animation of virtual creatures. In modern creative media, the behavioural animation of characters which act in a believable fashion is an ongoing challenge. Traditional action selection approaches which attempt to make an agent act rationally often fall short of the believability required for the modern consumer. Often the most believable action is not the most rational one, and our judgement of an agent's behaviour may also be based on the perception of its personality. Our approach, Affective Spaces Modelling, addresses these issues by creating a multi-dimensional environment constructed of aspect dimensions, with each aspect dimension representing a linear scale of a single component of the agent's internal state. Affective states can then be modelled by placing them in a single point in this environment. As the agent's state changes within the affective state space, different affects trigger appropriate actions. We demonstrate through a case study how the technique can be used to simulate different types of agent behaviour, operating both individually and as part of a group. Our case studies focus on groups of agents, allowing for the direct comparison of different personalities and examples of behavioural phenomena.**

Keywords*—*Affection; Virtual Creatures; Action Selection; Computer Animation

I. Background and Motivation

In interactive media such as video games, the believability of autonomous agents or non-player characters (NPCs) can be crucial to user enjoyment and immersion. Suspension of disbelief that the agent can be accepted as 'real' can be key to player engagement.

Increasingly realistic, detailed and dynamic video game environments have changed the requirements of creative AI methodologies by requiring NPCs to react in an increasingly believable fashion. Control over behaviour is as crucial as ever in order to direct the agents and ensure the desired gameplay experiences. Furthermore, the recent resurgence of virtual and augmented reality magnifies the importance of believability, as these technologies imbue NPCs with much more perceptual volume and presence in the virtual world.

The film industry is another domain where being able to direct the behaviour of autonomous agents and digital doubles is a current issue. In the past, large crowd scenes have required the hiring of hundreds of human extras [1]. But, over the last 20 years, the majority of films which have featured large crowds have used simulation, the most noted example being the CGI battle scenes in the Lord of the Rings franchise, generated by the MASSIVE software [2]. This is partly due to the obvious cost advantages, but also because artificial agents can be placed into situations which would be dangerous, costly or unethical to place an actor in. However, using autonomous agents to tell a story requires a balance between fully directed actions, and accurately simulated behaviour.

Ultimately developers need to be able to structure an agent's 'brain' in such a way that their behaviour is autonomous, but also believable within the context of the story, it's environment and its perceived personality.

An agent may also have multiple objectives (as is often the case with game agents) and believable action selection must consider which objective is prioritized at each given moment. Task priority as a challenge is often approached from a principle of optimization, but this usually does not take into account the agent's personality or mood which may bias towards a non-optimal or irrational action.

As an example, consider a parent faced with the decision of going into a burning building to rescue their child. Clearly, this situation creates a conflict between the agent's internal states as while the parent wishes to save their child they would also be aware of the low chance of success, and their own mortality. We could argue that the rational action would be not to enter the building, but, faced with that decision how many parents would act rationally?

In this work, we will discuss an algorithmic approach where affective states (such as emotions) are designed and positioned within an n-dimensional affective space. We will show how this method can be applied to the problem of action selection with a specific focus on creative industry applications.

This paper is organized as follows. We will begin by defining what we mean by action selection. We will then provide a short overview of the bio-inspiration for our new approach. We will then describe our Affective States Modelling technique, and two additional components which allow for the simulation of social phenomena. We will then describe a case study which discusses how the method can be implemented in simulations. We will conclude by discussing some of the related work, highlighting that despite the large number of action selection techniques that exist, none have been universally accepted.

A. Objectives

The design of the Affective State Modelling method is motivated by a number of design objectives:

1) The method should allow for various 'personalities' to be exhibited by agents within the same underlying architecture.
2) The method should allow for *structured unpredictability* helping to add realism to the behaviour in a simulation. Agents should be able to routinely surprise an observer with a non-optimal or seemingly irrational behaviour that is still contextually appropriate based on their personality.
3) Tuning and validating animations is typically a task that an animator is responsible for, rather than a computer scientist or software developer. As such, any system designed for this audience should be able to be calibrated without the need to modify code.

II. Defining Action Selection

Action selection is often discussed in terms of Reynolds' model of animation. In his original paper [3], the action selection problem is only explored from the context of combinational steering behaviours. Reynolds stated that this happens in two ways, either as a sequential switch between behaviours or by blending them together. He goes on to argue that switching is appropriate for certain behaviours, such as a prey animal switching between eating to fleeing from a predator. In this example, all other concerns would be forgotten, as it is unlikely that the creature will change activity during its escape. This behavioural switching occurs at the *Action Selection* layer of his model.

However, more commonly steering behaviours are blended to produce aggregated behaviours such as the *Boids* flocking model [4]. An important distinction to make is that the blending of behaviours happens at the middle steering behaviour level, not the higher action selection level. While this may seem like a subtle distinction, the emphasis is on the *objective* of the agent. In the first predator-prey example, there is a clear switch in objectives, from eating to fleeing. In contrast, while the agent is blending between the *separation*, *alignment* and *cohesion* steering behaviours during flocking, the objective, flying within a flock, remains the same. However, this is a distinction which is often misunderstood, and it is why we require a working definition of action selection.

A popular definition from ecological research is that "Action selection is choosing the most appropriate action out of a set of possible candidates" [5]. Under this definition, an action is defined as the mutually exclusive signal sent to the effectors of the animal, and that only one can be executed at any one time (a muscle cannot simultaneously contract and expand).

However, Tu[6] argues that this definition leads to confusion, as it assumes that all action selection refers to movement. Thus, action selection would be no different to the overall control system of the agent in question. The alternative that Tu proposes is to differentiate motor control from action selection. The motor control should be considered a lower level process, simply combining the movement of actuators to form useful *actions*. These actions can then be selected by the 'higher level' action selection system. This understanding of action selection bears some similarity to the model proposed by Reynolds. The Locomotion Layer Reynolds describes is clearly analogous to the motor controller, as both are lower level processes than the Action Selection Layer. To ensure this differentiation is explicit, and to avoid the confusion described by Tu, we propose expanding Tyrell's definition for the purpose of this research to define action selection in the following way:

> Action selection involves choosing the most appropriate action out of a set of possible candidates, where an action represents a process or behaviour undertaken to achieve a goal.

III. The Affective Domain

The goal of this work is to create autonomous creatures that behave in a personality appropriate way within complex and dynamic environments. In the real world and in simulations, this is typically accomplished with agents making decisions based on their internal state and an interpretation of external stimuli [7], [8]. In psychology, these mental activities can be classified into three domains: the affective; the conative; and the cognitive.

As discussed in our background, we are looking for a way to produce behaviour that is believable for the given situation. This includes acting irrationally and emotionally if appropriate for the current situation that the agent is in.

The affective domain is concerned with feelings and emotion, and how an organism reacts to stimuli. An affective state is a psychological construct which describes a particular state along descriptive affective dimensions.

Experimental research and some clinical phenomena have shown that affective reactions are often the first reactions an organism will make to a stimuli, and that the affective reactions can occur without extensive reasoning, and with greater confidence than cognitive judgement. This has led some psychology researchers to conclude that the affective domain is largely independent, and precedes the cognitive in the processing of stimuli, and for lower organisms they are the dominant reactions [9].

Individual states are often categorized and described by their valence, the intrinsic attractiveness (positive) or adverseness (negative) [10] of the affect. This categorization is based on the agent's perspective rather than a more holistic viewpoint. Fear is said to have negative valence, even though fear of specific situations can be considered positive. For example, being scared of fire may prevent an agent from being burnt, but the feeling of fear itself is unpleasant, so it is described as having negative valence. An affect is also sometimes used to describe a facial, vocal or gestural display, which serves as an indicator of the underlying affective state.

One attempt to implement the affective domain in an agent uses a motivational action selection system, with goal-oriented behaviours Where an emotional layer introduces flexibility and believability to the behaviours [11], [12]. However, there has generally been little research in developing affective based action selection systems.

IV. Affective States Modelling

We now describe the Affective States Modelling (ASM) approach to action selection, which we have developed with the objective of producing behavioural animation which resembles personality driven responses to external stimuli.

The model is proposed primarily for the design of creative applications (rather than as a cognitive model), and as such we evaluate the design in the following sections in terms of whether it achieves our design objectives as stated in the motivation.

In our approach, the observable behaviour of the agent is derived from its affective state at the current time-step. Each of the affective states are associated with an action (or actions) which are designed to be appropriate for the specific affect being exhibited. For example, if a prey agent is currently "afraid" then an appropriate action could be to flee.

The affective states are modelled as fixed points, defined as coordinates within an n-dimensional space. This space is made up of a number of aspect dimensions, each describing an aspect or feature of the agent such as its senses or the resources it currently has available. These are maintained and monitored internally within the agent. Each dimension has an upper and lower limit, which represent the extremes of the particular aspect with the central point representing the neutral condition (see figure 1). Each agent's current position along each of these aspect dimensions is stored numerically in a variable.

As the agents interact with their environment, the variables change according to the agents' current condition. Using the example provided in the background and motivation section, the parent's internal state in this situation (of a parent deciding to go into a burning building to save their child) can be described using two dimensions, notably, "protectiveness" (for their children) and "fear".

These variables are used to map the agent's current state as a single position (the 'physiological state') within the affective space. This is illustrated in figure 1 where four affects have been positioned within the space. The agent's current physiological state (red) is mapped as being within the bottom right quadrant ($x+$ and $y-$) of the aspect dimensional space. This position is compared to each of the affective states (light or dark grey circles). The affective state which has the shortest (Euclidean) distance to the current physiological state is the affective state adopted by the agent (dark grey circle). In the simple example pictured in figure 1, the agent will exhibit the `affect4` affective state. For simplicity, this illustration only has 4 affective states in 2-dimensions; however, additional fidelity and complexity can be introduced through the modelling of additional affective states or aspect dimensions.

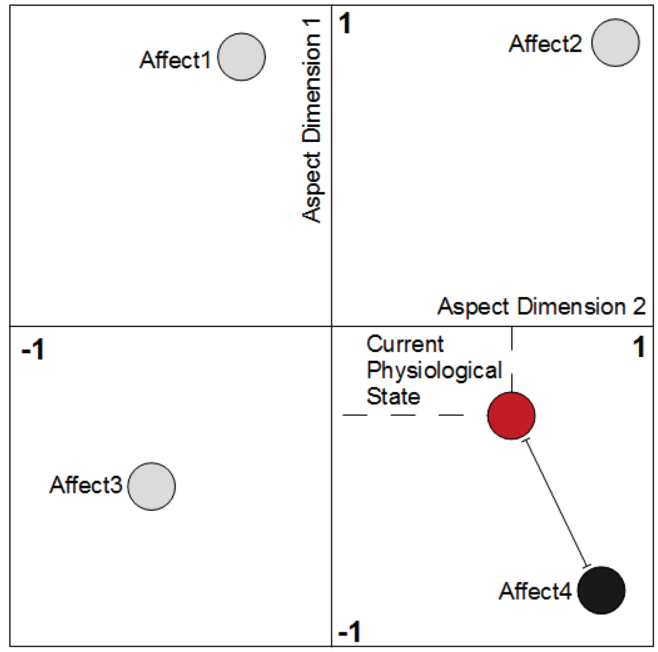

Fig. 1: A 2D example of four possible affective states (light gray or dark gray circles) expressed as positions within a 2D space. The agent's current physiological state (red circle) is compared to these predefined locations. In this example, the agent's current physiological state is closest to affective state 4, which is selected as the current affective state (indicated by the darker shade of gray).

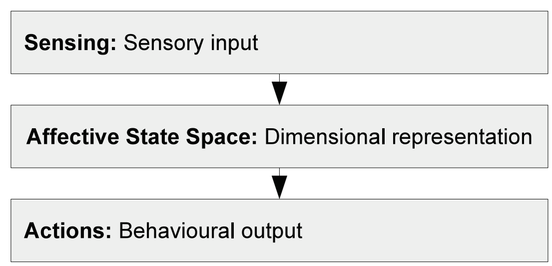

Fig. 2: The three components of the ASM action selection technique.

A. Components

To explore our approach in more detail, we break it down into three components (visualized in figure 2). The first (Sensing) refers to the environmental inputs and the internal drives of the agent. The second component (Affective State Space) is the aspect dimensions which structure the affective states space. This layer also involves the location of the individual affects within the space. The final component (Actions) is the design of each action associated with an affect.

1) Sensing: The sensing component is concerned with monitoring incoming information via the agent's sensors (internal and external) and storing it in a way that can be interpreted within the Affective State Space.

The sensing component of the agent is also responsible for any additional normalization or balancing functions on the individual aspect dimensions. For example, a decay may be required to return an agent to a normal state when a sensor is no longer activated (for example, how long it would take an agent to calm after being frightened).

2) Affective States Space: The Affective State Space is a multi-dimensional representation which stores the possible affects, and the agent's current physiological state as defined by the sensing component. The sensors provide values which are represented as a single position along an aspect dimension. These values collectively define the coordinates of a single point in the Affective States Space which represents the agent's physiological state at the current time-step. The physiological state is the raw state the agent is currently in as defined by it's sensors, before the affects are considered.

The agent's current state is then compared with the possible affects positioned within the Affective States Space. The closest affect to the current physiological state (calculated through Euclidean distance) is then set as the agent's current affect.

Designing an affect involves determining where in the Affective States space the affect should be positioned. The affect itself is a state which when active triggers a resultant action.

3) Actions: In our section defining action selection, we defined an action as *a process or behaviour undertaken to achieve a goal*. For the purpose of our model, the action is a function which produces behaviour that achieves a specific goal.

For the examples in this paper, we built actions as combinational steering behaviours [3] by changing the weightings of individual behaviours (such as seek, flee etc.). There are many other ways that an action can be designed (e.g. changing the target node of a path finding algorithm) and although steering behaviours are often a good way to construct an action, they may not be suitable for every simulation. For the purpose of this research, we are simply using them to demonstrate the technique.

In the case studies described in subsequent sections, each agent has a number of associated steering behaviours. The actions are designed by defining the weights that each of these steering behaviours will have when the agent is experiencing a specific affect.

For example, consider an agent that has two steering behaviours associated with it. The first behaviour (wander) influences the agent to explore the environment; the second (seek) influences the agent to move towards a home location. The agent also has two affects, each with associated actions. The first (curious) will cause the agent to explore the environment; the second (homesick) will cause the agent to return home.

Designing actions to suit the different affective states requires simply adjusting the weights of the two behaviours. For example, to achieve a 'curious' affect, a weight of 1.0 can be set for the wandering steering behaviour, and a weight of 0.0 for the seeking steering behaviour. To achieve a 'homesick' affect, the weights can be switched, with wandering now having a weight of 0.0, and seeking a weight of 1.0.

Immediate switching of a particular steering behaviour on or off in this fashion may be suitable for some situations such as simulating reactive or instinctive responses (like the fight or flight reflex). In other situations, a large number of possible intermediate behavioural states may be desirable in order to allow for emergent, intermediate or indecisive behaviour. For example, instead of a simple switch, we can interpolate between the steering behaviour weights of the old and current action when the affective state changes.

V. Crowds and Groups

When the same affective space is implemented on groups of agents, action selection based on Affective States Modelling (ASM) naturally produces behaviour which appears highly coordinated. This is simply because they react to the environment in the same way, but they are essentially acting independently.

However, in real life, groups of agents often do not act entirely independently, and are influenced by the state of the agents around them [13]. This form of influence can force agents to act irrationally, resulting in various group-behaviour phenomenons such as mass-hysteria [14], or stampedes [15]. Phenomena such as this has resulted in many devastating events in history including the Baghdad Bridge stampede in 2005 [16] and the Phnom Penh disaster in 2010 [17].

This social phenomena can be simulated using the ASM approach through the implementation of an 'empathy' function. In the following subsections, we will describe the implementation of an empathy function which we will refer to as 'intimidation'.

We work from the assumption that the agent is a member of a crowd. If the agent is separated from other agents, it is sensible to assume that they would not be influenced by distant neighbours. Before either of the empathy functions are used, an agent must establish its local neighbours.

A. Intimidation

In the intimidation method, an agent is influenced by a subset of agents (the 'neighbours') within a predefined range (their 'vision').

In the intimidation method, the agents have an awareness of the current affective states of their neighbours, not their internal physiological state The agent only has knowledge of the affective state based on its own internal representation. So, if the agent is being influenced by a neighbour that is scared, it would assume this was the result of a physiological state matching the coordinates of the agent's own scared affect. Although the positions of the internal states may be different, this is the reference point the agent would work from.

At each time step, the agent observes the affective states of all its neighbours, and converts these into coordinates by comparing the states it sees to its own representation (see figure 3.A). A new centre-of-mass location is calculated by generating the sum of these locations and dividing by the number of neighbours (see figure 3.B).

The next stage is to calculate the group authority, the amount of influence that the neighbours have on the agent's own affective state. The group authority is ascertained by generating the sum of the combined authorities of all the agent's neighbours, and dividing by the number of neighbours.

Once the group authority and affective state influence have been calculated, a weighted average is used to calculate the final state, with the group authority replacing both the neighbour state and the neighbour authority (see figure 3.A).

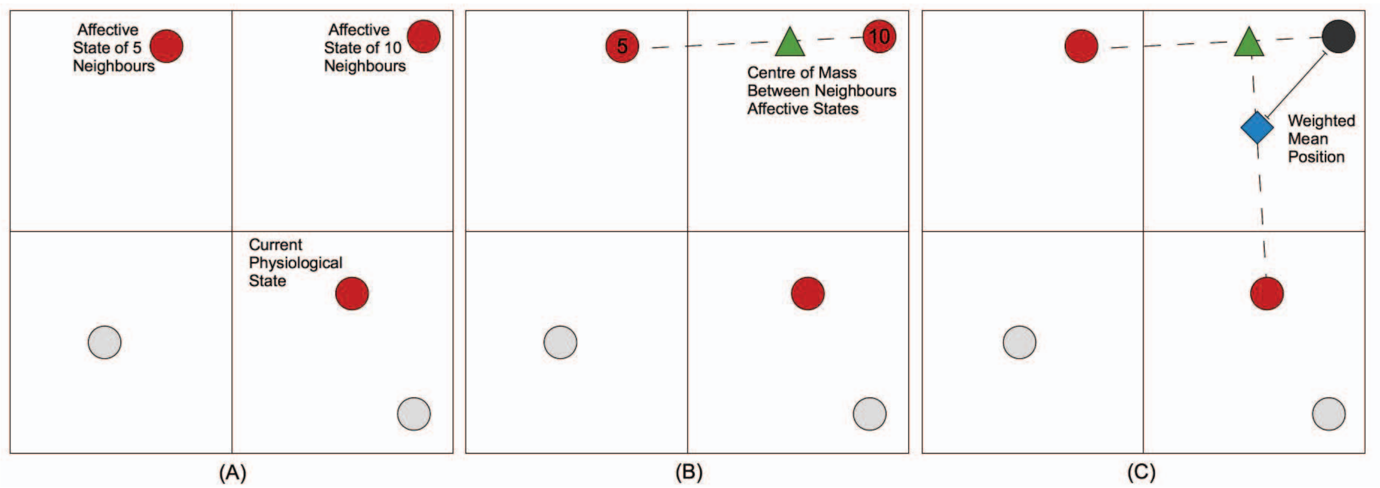

Fig. 3: A graphical representation of the stages of the intimidation algorithm, using the same affective space as figure 1. In (A), the affective states of local neighbours are evaluated and plotted onto the agent's own representation (red circles). In (B), a "centre-of-mass" position (green triangle) between the two states is calculated, with the weights taken from the number of agents in each state. In (C), a weighted mean position (blue diamond) is established between the agent's physiological state and the centre-of-mass position. The weights are generated through the group authority and the agent's independence. The weighted mean is then used to establish the current affective state (dark grey).

VI. Simulation of Fish

To evaluate the ASM method, we simulated three separate species of fish interacting with each other and a predator. The aim of the study was two-fold: firstly, we wanted to see whether distinctive species personalities could be designed into the fish. Secondly we wanted to test the scalability of the algorithm, increasing the number of affective states and aspect dimensions while moving away from 2D simulations and into 3D game environments generated using the Unity Game Engine [18].

Each fish species was given a virtual brain with varying levels of complexity. Affective states were constructed based on information freely available through online sources; for example, sources such as web articles, aquarium forums, and the 'species' Wikipedia page. Additionally, videos of the fish species were observed to provide visual insight. A corpus of the videos and web sources which informed the design of affects and actions can be found on the AMBER project website [19].

Once the simulation was run, a validation and tuning exercise was conducted comparing the virtual species behaviour in the simulation to videos of the real counterpart. If required, the positions of the affective states within the dimensional space was then tuned to produce behaviour which resembled the virtual creature's real-world counterpart as closely as possible.

The primary intention of this model was not to accurately simulate fish from a biological perspective, but rather to create a simulation which was believable to a human observer that would be suitable for use in the creative sector. It also serves as an example of how different personalities and levels of perceived intelligence can be easily constructed and modified by an animator using our approach.

(a) (b) (c)

Fig. 4: Fish Models: (a) The virtual model of the Angel fish. (b) The virtual model of the Blue Tang fish. (c) The virtual model of the Clown fish.

A. Parameters and Weights

For the each of the three species, each agent within the population had an independence and authority randomly generated within a set range. The Angel fish were given a low authority and a high independence as they are generally solitary fish. The more social Blue Tang were given a low independence and a high authority, to make the fish react with the group. Finally, the Clown fish were given entirely randomly generated independence and authority as there are both social and solitary examples.

B. Angel fish

Various online sources suggested that Angel fish are bold, almost fearless. Also, typically they are either solitary or form harems with a single male leading several females. Various sources (such as fish keeper guides) also suggested that the fish were fairly intelligent and would aggressively defend a small territory. (The Angel fish model we used in our simulation can be seen in figure 4.a).

1) Virtual Brain: We gave the fish a four dimensional virtual brain, representing hunger, fear, anger and sociality. The agent's current state moved around within the four dimensional

TABLE I: The Affective States Weightings table used to define the virtual 'brain' of the Angel fish. Similar weightings tables were used to define Clown and Blue Tang fish.

	Aspect dimensions			
Affective State	**Hunger**	**Fear**	**Anger**	**Sociality**
Lonely	0	0	0	-0.9
Scared	0.1	0.6	0.3	0.1
Curious	1	1	0.4	0.5
Hungry	0.6	0	0	0
Normal	0	0	0	0
Exposed	0.2	-0.9	0	-0.9
Territorial	0.9	0.3	-1	0.3

states space at each time-step. On the hunger dimension, the value went to +1.0 whenever the fish encountered a food source and decayed at each time-step (to a maximum of -1.0). The fear went from positive (not scared) to negative (scared) in a linear fashion based on the proximity of a predator. The agent's anger value increased (angry) and decreased (calm) in a similar way based on the proximity of fish from another species. The sociality aspect value decayed the further the fish went from the rest of the school, and increased the closer it came to another school member.

Seven affective states were devised based on the insights we gained from the behavioural notes and video observations. These were: Lonely, Scared, Curious, Hungry, Normal, Exposed and Territorial. These were initially placed in the locations described in table I within the 4D representation.

As with the 2D examples, each one of these affective states adjusted weightings on steering behaviours to vary the direction and speed of the fish.

C. Blue Tang Surgeonfish

The Blue Tang fish (see figure 4.b) is described as not being migratory, and instead stays within a range of a home position. They are also described as having three distinct social modes, territorial, wandering and schooling. The territorial behaviour involves remaining close to the home location, wandering is the process of independently exploring the environment and schooling involves coordinated group movements.

1) Virtual Brain: The Blue Tang were given a simpler, 3 dimensional virtual brain, removing the anger dimension. None of the behavioural guidelines we found indicated that the fish demonstrated any aggressive tendencies. The territorial behaviour observed in videos seemed more submissive, staying close to a territory or school rather than actively defending it. The fish had the same basic affects as the angelfish in slightly different positions to better represent their perceived personality.

The adjustments made the Blue Tang easily startled, and more likely to stay close to their school than the Angel fish. This is demonstrated in figure 5, where a small school of Blue Tang have wandered into the territory of an Angel fish. The Angel fish responded aggressively, chasing away the smaller fish, who remained in a tight school.

D. Clown Fish

The third fish included in this study were Clown Fish (see figure 4.c), a brightly coloured species made famous by Disney's animation film 'Finding Nemo'. They are an interesting species for several reasons. Firstly, each individual forms a symbiotic relationship with an anemone they reside within. This provides the Clown Fish with protection, as outside the anemone it would be at greater risk of predation.

The species live in a group called a queue, which describes their social hierarchy, with dominance being achieved by outliving more dominant members. The current dominant pairing within the group breed more than the subordinate members of the group. Due to this social structure, the fish adopt behaviour which will increase their probability of achieving social dominance. The size of the group correlates with the home range of the fish (the patch), with the fish preferring to say close to their anemone.

The behavioural discussion available noted that the species stayed close to their home territory and would retreat back there whenever scared. It also indicated that the fish would retreat quickly if they encountered a predator, increasing their chance of moving up the social hierarchy within the queue. There was also evidence that Clown fish were aggressive within their own social group, and competitive over food.

1) Virtual Brain: We gave the Clown fish a similar virtual brain as the Blue Tang including the hunger and fear dimensions. A third dimension (Isolation) represented how far the fish had roamed from its home location.

Fear (as with the Angelfish and Blue Tang) increased with the proximity of the predator. However, this was at twice the rate of the Blue Tang, the intention being to make the fish more responsive to danger. The hunger state changed in the same fashion as the Angel fish and the Blue Tang. On the Isolation dimension, the value started at +1.0 when the fish was within its home location (the simulated anemone) and decayed in relation to the distance between itself and its home location.

The Clown Fish was given the same basic affective states as the Blue Tang fish with the exclusion of the Lonely state which was excluded.

To make the personality of the Clown fish distinct from that of the Blue Tang, the weightings of the actions were adjusted. This was to demonstrate a different way the behaviour could be modified while maintaining the same "virtual brain" prototype as the Blue Tang. For example, when scared, instead of simply running from the predator, they would try to return to their home position (an anemone). This is demonstrated in figure 6. Additionally, when searching for food (when in the hungry affective state), they would attempt to avoid each other. This was achieved by biasing a "separation" weighting, causing them to veer away from each other if they came too close.

E. Verification and Validation

After each fish species was designed, a process of verification and validation was undertaken. This step was inspired by the model promoted by Naylor and Finger [20] who proposed a three phase approach to validating simulation models.

978-1-4673-8460-5/16/$31.00 2016 IEEE

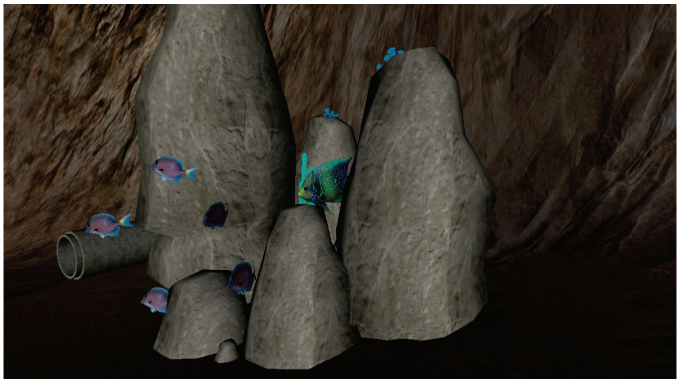

Fig. 5: A school of Blue Tangs wander into an Angel fish's territory. The Angel fish responds aggressively, chasing the Blue Tangs back out into the open water.

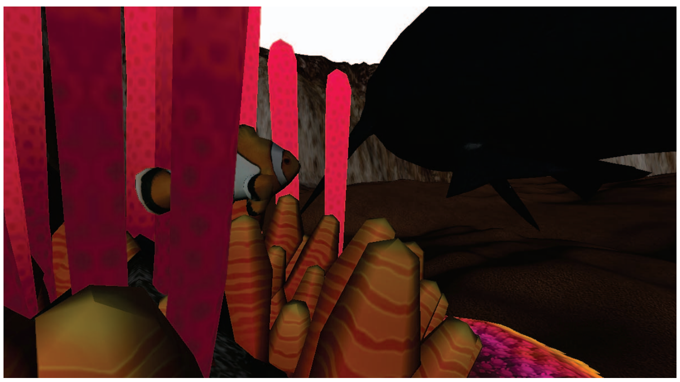

Fig. 6: A Clown fish returns to its home anemone after being startled by the presence of a predator.

1) Build a model that has high face validity.
2) Validate model assumptions.
3) Compare the model input-output transformations to corresponding input-output transformations for the real system.

Our final stage of designing each fish relates to the third point (comparing the model input-output transformations to corresponding input-output transformations for the real system). This involved running a simulation while observing publicly available videos of the fish in their natural habitat (available on-line through services such as YouTube).

While watching the videos, the observer was tasked with noting down situations which consistently produced recognizable behaviour. For example, a startled Clown Fish returning to its anemone. The same situation was then reproduced in the simulation model, and tuned until the fish produced situationally appropriate actions based on the observations (via a reverse-engineering process).

With our system, there are three distinct points in the pipeline where adjustments can be made, the sensing, the affective states space, and the action weightings.

- **Sensing** The adjustments here concern how items are sensed in the environment. This includes how much effect a sense has on the internal state, and how quickly the state will decay back to normal levels once a particular sense is no longer activated.
- **Affective States Space** Tuning at the affective states component involved moving the designed affective states within the n-dimensional space to change when they activated. This was simply achieved by changing the encoded positions of each emotion as defined in the Affective States Space Weightings Table. Generally speaking, moving an affective state closer towards the normal position (0 in all dimensions) would make it more likely to activate, whereas moving it further from this position would make it harder to activate. This simulates the normal and extremes of personalities.
- **Action** In this paper, we have focused on the action selection rather than the actions themselves. Clearly the way you would tune actions would depend entirely on how you have chosen to encode them. For the purpose of this case study, we used a combination of steering behaviours, with the resultant observed action emerging from the blending of each steering behaviour depending on the weights used as defined by the Action Selection Weightings table. To tune the actions with this method of encoding, we simply need to modify the weightings as defined in the Action Selection Weightings table, increasing or decreasing a bias to produce the desired emergent behaviour.

For the purpose of this process, the validation was conducted in two stages. Each fish type was tuned individually before being placed back in the main environment with the other creatures. This allowed us to first make sure that the individual fish behaved in a way which seemed appropriate in a sterile environment. The second stage allowed us to identify and debug any anomalous behaviour which arose from placing the fish in the same environment.

F. 3D Case Study Discussion

The purpose of this case study was to see if realistic fish behavioural animations could be generated using our approach. As a secondary objective, we wanted to see how different personalities of fish could be generated using the approach.

One interesting aspect which we discovered through this case study is that the perceived intelligence of a creature can by dramatically changed through the addition and removal of dimensions. Specifically, as a creature has less dimensions in its virtual brain, it appears more instinctive and reactive to its environment. However, a larger number of dimensions (the Angel fish) generated a virtual personality with more subtlety and complexity.

As with any behavioural AI or procedural animation technique, fine tuning can be time consuming. However, breaking the animation into three distinct points along the animation pipeline has a logical advantage for the animator. Notably, the clear separation allows issues in the animation to be easily identified and managed. For example, any point where the sensing results in aspect values that are held at the extremities of the ranges being used is identified as a problem with the sensing component. Conversely, if the sensing appears to be correct but inappropriate actions are being triggered, this can be isolated as a problem with the affective states component.

VII. Related Work

In this section, we discuss some of the work which relates to our method. First we will explore how action selection has been explored in games based research. It is important to note that with games AI, identifying the current state of the art is tricky, as some of the most celebrated examples are commercial products and not openly published. For this reason we will focus on academic research in this area, rather than discussing commercial examples. We will then explore research into how personalities and moods have been simulated in other examples. We will conclude with a brief discussion on dimensional representation of knowledge.

A. Action Selection In Games

In this review, we have purposely focused on reactive approaches which have been applied to the creative industry and excluded techniques which require significantly more computation (therefore limiting their ability to operate in real-time). Action selection is a foundational problem in A.I. research, so this review is not intended to be a comprehensive survey of the field, but rather to identify the spectrum of approaches.

The reactive AI approach underpins many AI systems in the creative industry. They are advantageous for computer games due to their ability to operate in real-time, and they typically produce more natural looking movement than more computationally-intensive approaches.

One of the principle methods associated with the reactive approach is the Subsumption Architecture, introduced by Brooks in 1986 [21], [22]. The subsumption architecture couples sensory information to action selection, by decomposing behaviour into a series of layers. Each layer represents a different level of behavioural competence, responsible for a single behavioural goal, such as ‘search’ or ‘avoid obstacles’. Each layer runs continuously, and independently, the only communication between the layers is inhibition or suppression signals, which determine which module should output behaviour. The individual modules are typically constructed as Augmented Finite State Machines [23], [24]. There are examples of the architecture being applied to games, including soccer simulations [25], but its main application has been in the field of robotics [26].

Within the computer games industry specifically, the approach which has arguably made the greatest impact is the Finite State Machine (FSM). An FSM describes the possible states the agent can find itself in, and the events and conditions by which an agent’s changes its current state [27]. While this technique is computationally efficient, it has been criticized as being “unwieldy” [28]. Additionally, the approach can easily become cluttered in large implementations, requiring the duplication of states (the situational reuse problem) and suffer from transition bloat, which (from a programming perspective) can result in a system which is messy and lacking in scalability [29].

Because of these criticisms, most games now use behaviour trees, first implemented in the robot Shakey [30]. The main advantage of this approach is that action selection is re-evaluated at each decision cycle [31]. If a previous condition becomes invalid, the system will automatically return to the root node and re-evaluate the possible actions.

More recently in the games industry, there has been a plethora of approaches that are loosely based on these early AI systems / methods, but employ a range of techniques that are not easily categorized under a single heading. Clearly, the problem of how to design Action Selection for an autonomous Agent is a difficult issue.

There have been many methods applied to Action Selection and this brief survey has barely scratched the surface. The purpose of this section was to show the variety of approaches within the reactive family of algorithms, and emphasize that no one approach has yet been universely adopted, or has become the de facto standard for action selection. For this reason, developing a new approach is a worthy exercise as the hope is that it will lead to further and different insights into the existing approaches in the literature.

B. Virtual Personalities, Moods and Emotions

The simulation of virtual personalities is a research area which combines traditional behavioural animation with psychology. The aim is to personify virtual humans, providing them with the ability to respond in a believable way based on our knowledge of their personality. This is one of the principle objectives of the ASM method.

In [32], the Autonomous Interactive Reaction (AIR) model is proposed as a simplified method of behaviour selection based on emotion modelling. In this approach, NPCs have parameterized emotions and personalities, and the behaviour they exhibit at any given time is determined based on rules which refer to these parameters. The values of each parameter are mutually adjusted between communicating agents, essentially adjusting the mood of each NPC based on its neighbors. One interesting component of this approach is the way the agent is designed as a series of three layers: *Personality*, the governing parameters of the agent which do not change; *Emotion*, which is the expression of the agent’s changing internal state; and *Knowledge*, which is a parameter which represents preferences towards objects in the environment. The authors demonstrate that appropriate behaviours can be produced when interacting with AIR agents.

Another area where researchers have included visible personalities into virtual humans is conversational agents, for example where personalities can be reflected visually through facial animation. One approach to this problem has been to use a Bayesian Belief Network (BNN) [33] to represent emotion. This has been extended by further de-constructing the agent’s emotion into a series of layers, the highest being the personality, the next being a mood, and finally an associated facial expression [34], [35]. A clear distinction is made between personality and mood. The personality (defined using the Five Factor Model) specifies reactions to stimuli which in turn causes the mood to change. The authors acknowledge that mood and personality only differ in terms of arousal, but they have chosen to define the mood layer as a longer term state

TABLE II: Components of the Five Factor Model (FFM)

Factor	Description	Adjectives
Extroversion	Preference for and behaviour in social situations	Talkative, Energetic, Social
Agreeableness	Interactions with others	Trusting, Friendly, Cooperative
Conscientiousness	Organized and persistent in achieving goals	Methodical, Organized, Dutiful,
Neuroticism	Tendency to experience negative thoughts	Insecure, Distressed
Openness	Open minded	Imaginative, Creative, Exploitative

of mind which directly influences emotions (and the facial expressions). The authors claim that rule-based systems are unlikely to be successful in the believable simulation of human behaviour, as uncertainty is an important facet. A strength of this approach is that it allows personalities to be designed by changing parameter values or weightings.

A common theme between these approaches is that computational models of emotion must include two parts[36]:

1) mechanisms eliciting emotions from external and internal stimuli, including potentially the agent's own goals, beliefs and standards;
2) emotion representations keeping track of the emotional states and their changes over time.

From a design perspective, this is somewhat self evident; an agent requires stimulus of some form in order to react emotionally and these emotions need to be represented in a way which allows the agent to keep track of its emotional state. Our ASM technique covers these two areas with the sensing and Affective State Space layers.

C. Dimensional Representation of Personalities

A dimensional approach which is well respected in the psychology community is the Five Factor Model (FFM) [37], [38]. This model is used to describe personality and was proposed not only as a tool to aid in the understanding and description of personality traits, but also as a model for diagnosis and treatment of personality disorders. The five factors which make up the model are considered to be the dimensions which construct the personality space [34]. The theory states that any personality can be represented within this space by adjusting the position along each of the dimensions. The five factors which make up this model can be seen in table II.

As a diagnostic tool, the FFM has outperformed other personality modelling tools in the prediction of typical personality disorder symptoms [39]. However, the approach has attracted some criticism, notably that it only considers a subset of personality traits and ignores others (such as honesty, conservativeness and seductiveness) [40]. The ASM method, rather than being restricted to a limited number of factors, can be adapted to suit the simulation.

VIII. Conclusion

In this research, we introduced a novel approach to action selection where affective states are modelled as positions within an n-dimensional space constructed from aspect dimensions. We discussed two example implementations and a case study demonstrating the flexibility of the algorithm. Having extensively tested the approach, we have developed the following conclusions.

A. Ease of Calibration

One of the main strengths of the approach is the ease in which it can be calibrated. The animator is able to adjust the personality of each agent simply by adjusting the relative location of each affective state within the n-dimensional space. This provides an animator with the ability to design complex personalities for the simulated agent without the need to code. It also allows the increasing or decreasing of an agent's perceived complexity simply by adding additional dimensions, or removing existing dimensions from the affective states space.

B. Personalities

One of the principle motivations behind this work was the development of an approach suitable for the animation of group behaviour. One of the challenges we identified was giving agents within a group distinct, individual personalities while remaining believable to the user. This is important in games as it helps prevent the user from "learning the AI" and predicting the exact responses of NPCs. The simulation of individual personalities is also important to prevent agents resembling drones or zombies with no unique facets.

The affective states n-dimensional approach provides a simple solution to this problem. An animator is able to define a single, behaviourally appropriate agent, then, by randomly varying (within a limited range) the location of affective states, distinct personalities can be procedurally generated around this original prototype. This can have the effect of making characters appear to have different personalities, while still keeping their core personality within an acceptable range of a realistic prototype.

C. Future work

Our research has shown that our Affective State Modelling approach is a suitable action selection technique. However, while we have had great success simulating relatively primal behaviour, as we move up the behavioural hierarchy, we believe that we will begin to encounter limitations.

The next stage of this work will involve building models which use our ASM technique to influence a cognitive process. We believe that the interplay between feeling and thinking is the next important area to investigate in order to produce more realistic agents. Studies in psychology have proven that an affective state can directly affect judgment and decision making [41]. For example, people are often advised to not make important decisions when they are angry.

Aside from producing believable agents, this work has wider implications to the general field of decision making. In theory, we can vary the quality of an agent's decisions based on their affective state generating predictable, irrational behaviour, within the framework of bounded rationality.

References

[1] M. Praºák, "Locomotion for crowd animation," Ph.D. dissertation, University of Dublin, Trinity College, 2012.

[2] D. Thalmann, C. Hery, S. Lippman, H. Ono, S. Regelous, and D. Sutton, "Crowd and group animation," ACM, p. 34, 2004.

[3] C. W. Reynolds, "Steering behaviors for autonomous characters," in *Game developers conference*, vol. 1999, 1999, pp. 763–782.

[4] ——, "Flocks, herds and schools: A distributed behavioral model," *ACM SIGGRAPH Computer Graphics*, vol. 21, no. 4, pp. 25–34, 1987.

[5] T. Tyrell, *Defining the action selection problem*. Centre for Cognitive Science, University of Edinburgh, 1994.

[6] X. Tu, *Artificial animals for computer animation: biomechanics, locomotion, perception, and behavior*. Springer, 1999, no. 1635.

[7] B. Blumberg, M. Downie, Y. Ivanov, M. Berlin, M. P. Johnson, and B. Tomlinson, "Integrated learning for interactive synthetic characters," in *ACM Transactions on Graphics (TOG)*, vol. 21, no. 3. ACM, 2002, pp. 417–426.

[8] X. Tu and D. Terzopoulos, "Artificial fishes: Physics, locomotion, perception, behavior," in *Proceedings of the 21st annual conference on Computer graphics and interactive techniques*. ACM, 1994, pp. 43–50.

[9] R. B. Zajonc, "Feeling and thinking: Preferences need no inferences." *American psychologist*, vol. 35, no. 2, p. 151, 1980.

[10] D. McDuff, R. El Kaliouby, K. Kassam, and R. Picard, "Affect valence inference from facial action unit spectrograms," in *Computer Vision and Pattern Recognition Workshops (CVPRW), 2010 IEEE Computer Society Conference on*. IEEE, 2010, pp. 17–24.

[11] E. de Sevin and D. Thalmann, "A motivational model of action selection for virtual humans," in *Computer Graphics International 2005*. IEEE, 2005, pp. 213–220.

[12] E. De Sevin and D. Thalmann, "An affective model of action selection for virtual humans," in *Proceedings of Agents that Want and Like: Motivational and Emotional Roots of Cognition and Action symposium at the Artificial Intelligence and Social Behaviors 2005 Conference (AISB" 05)*, no. VRLAB-CONF-2005-014, 2005.

[13] S. Schneider and H. Weinberg, *The large group re-visited: the herd, primal horde, crowds and masses*. Jessica Kingsley Publishers, 2003, vol. 25.

[14] R. E. Bartholomew, *Little green men, meowing nuns and head-hunting panics: A study of mass psychogenic illness and social delusion*. McFarland, 2001.

[15] F. M. Burkle and E. B. Hsu, "Ram Janki Temple: understanding human stampedes," *The Lancet*, vol. 377, no. 9760, pp. 106–107, 2011.

[16] R. F. Worth, "More than 950 iraqis die in stampede on Baghdad bridge," *New York Times*, vol. 31, 2005.

[17] E. B. Hsu and F. M. Burkle, "Cambodian bon om touk stampede highlights preventable tragedy," *Prehospital and disaster medicine*, vol. 27, no. 05, pp. 481–482, 2012.

[18] U. Technologies, "Unity-3d game engine," 2010.

[19] P. AMBER, "Fish Behaviour Corpus," Online: http://www.project-amber.co.uk/2014/06/fish-behaviour-corpus/, jun 2014.

[20] T. H. Naylor and J. M. Finger, "Verification of computer simulation models," *Management Science*, vol. 14, no. 2, pp. B–92, 1967.

[21] R. A. Brooks and J. H. Connell, "Asynchronous distributed control system for a mobile robot," in *Cambridge Symposium_Intelligent Robotics Systems*. International Society for Optics and Photonics, 1987, pp. 77–84.

[22] R. A. Brooks, "A robust layered control system for a mobile robot," *Robotics and Automation, IEEE Journal of*, vol. 2, no. 1, pp. 14–23, 1986.

[23] ——, "A robot that walks; emergent behaviors from a carefully evolved network," *Neural computation*, vol. 1, no. 2, pp. 253–262, 1989.

[24] R. A. Brooks and A. M. Flynn, "Fast, cheap and out of control," DTIC Document, Tech. Rep., 1989.

[25] H. Nakashima and I. Noda, "Dynamic subsumption architecture for programming intelligent agents," in *Multi-Agent Systems, International Conference on*. IEEE Computer Society, 1998, pp. 190–190.

[26] R. A. Brooks, "Elephants don't play chess," *Robotics and autonomous systems*, vol. 6, no. 1, pp. 3–15, 1990.

[27] A. Nareyek, "Ai in computer games," *Queue*, vol. 1, no. 10, p. 58, 2004.

[28] A. Khoo, G. Dunham, N. Trienens, and S. Sood, "Efficient, realistic npc control systems using behavior-based techniques," in *AAAI Spring Symposium on Artificial Intelligence and Computer Games*, 2002, pp. 46–51.

[29] S. Rabin, *AI game programming wisdom*. Charles River Media, Inc., 2002.

[30] R. E. Fikes, P. E. Hart, and N. J. Nilsson, "Learning and executing generalized robot plans," *Artificial intelligence*, vol. 3, pp. 251–288, 1972.

[31] I. Horswill, "Very fast action selection for parameterized behaviors," in *Proceedings of the 4th International Conference on Foundations of Digital Games*. ACM, 2009, pp. 91–98.

[32] J. Sato and T. Miyasato, "Autonomous behavior control of virtual actors based on the air model," in *Computer Animation'97*. IEEE, 1997, pp. 113–118.

[33] G. Ball and J. Breese, "Emotion and personality in a conversational agent," *Embodied conversational agents*, pp. 189–219, 2000.

[34] S. Kshirsagar, "A multilayer personality model," in *Proceedings of the 2nd international symposium on Smart graphics*. ACM, 2002, pp. 107–115.

[35] A. Egges, S. Kshirsagar, X. Zhang, and N. Magnenat-Thalmann, "Emotional communication with virtual humans." in *MMM*, 2003, pp. 243–263.

[36] E. Tanguy, P. J. Willis, and J. Bryson, "Emotions as durative dynamic state for action selection." in *IJCAI*, vol. 7, 2007, pp. 1537–1542.

[37] J. M. Digman, "Personality structure: Emergence of the five-factor model," *Annual review of psychology*, vol. 41, no. 1, pp. 417–440, 1990.

[38] R. R. McCrae and O. P. John, "An introduction to the five-factor model and its applications," *Journal of personality*, vol. 60, no. 2, pp. 175–215, 1992.

[39] R. M. Bagby, M. Sellbom, P. T. Costa, and T. A. Widiger, "Predicting diagnostic and statistical manual of mental disorders-iv personality disorders with the five-factor model of personality and the personality psychopathology five," *Personality and Mental Health*, vol. 2, no. 2, pp. 55–69, 2008.

[40] S. V. Paunonen and D. N. Jackson, "What is beyond the big five? plenty!" *Journal of personality*, vol. 68, no. 5, pp. 821–835, 2000.

[41] N. Schwarz and H. Bless, "Happy and mindless, but sad and smart? the impact of affective states on analytic reasoning," *Emotion and social judgments*, pp. 55–71, 1991.