

The Impact of Durative State on Action Selection*

Joanna J. Bryson

Konrad Lorenz Institute for Evolution and Cognition Research, A-3422 Altenberg, Austria
Artificial models of natural Intelligence, University of Bath, Bath BA2 7AY, United Kingdom

Abstract

Chemical / hormonal diffusion is the phylogenetically oldest form of biological action selection. In this paper we argue its persistence in higher animals is a consequence of its utility in solving problems of dithering between high-level goals. Chemical state underlying emotions and drives provides greater persistence more easily than the electrical action potential systems underlying the fine details of action sequencing, while also providing periodicity and transience not easily afforded in longer-term learning systems such as synaptic plasticity. We argue that artificial real-time autonomous systems require similar systems, and review our own efforts and approaches to providing these.

Introduction

Action selection is one of the fundamental problems of intelligence (Prescott, Bryson, & Seth 2007). For an agent (whether biological or artificial), *action selection* is the ongoing problem of choosing what to do next. For a developer, action selection presents two problems — designing the agent's action-selection process, *and* determining the level of abstraction at which the process will operate. A physical agent must ultimately perform precise control of motors or muscles, but this is almost certainly not the level at which decision making should occur.

In previous work we have described Behavior Oriented Design (BOD) (Bryson & Thórisson 2000; Bryson & Stein 2001b; Bryson, Caulfield, & Drugowitsch 2005), a methodology for determining through iterative development the appropriate decomposition of intelligent behaviour into three categories:

- modules providing set primitive behaviours, which are coded using standard object-oriented languages,
- adaptive components of these modules, which are the mechanism through which the artificial agents can learn, and

*Research funded by the British Engineering and Physical Sciences Research Council (EPSRC) under the Adaptive and Interactive Behaviour of Animal & Computational Systems programme, grant number GR/S79299/01.
Copyright © 2008, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

- structured dynamic plans, which arbitrate between currently-appropriate actions, providing real-time action selection where there is immediate contention for an agent's resources (e.g. its physical location).

While one could theoretically imagine an intelligent agent being controlled entirely by a single system, e.g. an adaptive neural network or an elaborate dynamic plan, we have both demonstrated empirically (Bryson 2000b; 2001; Partington & Bryson 2005; Brom *et al.* 2006; Bryson & Leong 2006) and argued from the history of AI agent construction (Bryson 2000a) that the above sort of decomposition makes the successful construction of an agent far more probable, and makes its code easier to maintain.

We now believe that this decomposition, while well-supported in the AI literature, is not entirely adequate. In recent years we have engaged in scientific modelling of non-human primates using artificial life (Bryson & Leong 2006; Bryson, Ando, & Lehmann 2007), and creating VR humanoid avatars (Tanguy, Willis, & Bryson 2006; 2007) and intelligent game characters (Partington & Bryson 2005; Brom *et al.* 2006). This has led us to improving our coding / description of long-term complete animal behaviour. This improvement has required developing representations underlying emotional and other goal-directed behaviour expression.

After a brief review, we describe an established AI ontology. We then examine problems encountered, and describe expanding the system to solve them. These problems and solutions may be salient for explaining the role of emotional expression in natural action selection, a question we visit throughout.

Durative State in Action-Selection Architectures

We believe *durative state* (in contrast to permanent or transient) has not been adequately incorporated in standard action-selection architectures. By this we mean primarily emotions and drives, but we will also touch on consciousness towards the end of the paper.

Certainly many (if not most) current AI architectures do address emotions and/or drives in some way. However, we do not think that the best-practice techniques of action selection have been fully integrated with an emotional systems. Some systems use emotions as relatively isolated systems, essentially social effectors for human-robot interac-

tion (Breazeal & Scassellati 1999; De Rosis *et al.* 2003; Velásquez & Maes 1997). Outstanding in this group is Marcella & Gratch 2002, whose emotional system also affects action selection by changing their military training system's basic reasoning in order to better simulate civilian reactions to troops. Most AI research postulating agent-centric utility for affect have focussed on it as a sort of reward accumulator. Emotions serve as additional pieces of internal state to assist in learning and applying action selection policies (Hiller 1995; Gadanho 1999; Zadeh, Shouraki, & Halavati 2006; Broekens, Kusters, & Verbeek 2007). Some of this research has been inspired partly by the controversial somatic marker hypothesis (Tomb *et al.* 2002). Several elaborate architectures have been proposed but not yet constructed which postulate similar referential or marker roles for emotional systems, but operating at a higher self-referential level (Sloman & Croucher 1981; Norman, Ortony, & Russell 2003; Minsky, Singh, & Sloman 2004). These latter theories are beyond the scope of this paper, which deals with basic action selection, not reflective reasoning.

We believe emotions and drives are absolutely integral to action selection, determining the current focus of behaviour attention. The nearest correlates to our ideas include the work of Morgado & Gaspar (2005), but these mechanisms are not incorporated into a full architecture. At the other extreme, Breazeal (2003) also treats both moods and drives as mechanisms for maintaining homeostatic goals, and has an extremely elaborate robotic system built around them. The system we present here emphasises generalisation and usability, and is cast within a well-established AI architectural framework: the hybrid or multi-layer architecture, which features both modules for programmability and hierarchical plan structures for coding agent priorities or personality. Below we describe emotions and drives incorporated into a general-purpose, publicly-available open-source modular AI control system.

We originally believed durative state should be represented simply in behaviour modules like other variable state. We were wrong. We were wrong. We now believe that a system more like spreading-activation networks is needed for influencing some of the selection of the highest-level goals. Of course, spreading activation has been shown intractable for the problem of action selection as a whole due to poor scaling (Tyrrell 1994). But over a small, strictly limited set of options (such as a list of emotions / goals) it is a feasible system (Shanahan 2005; Maes 1991). We describe our new action-selection systems immediately after first describing the basic components of our existing action-selection mechanism.

Currently Conventional Components of Systems AI Action Selection

This section gives a quick overview of a widely accepted state-of-the-art architecture for action selection. While describing an ontology of intelligent control in ordinary AI, we also provide an analogous ontology of well-understood natural intelligence mechanisms. This will help illustrate our account of what is currently missing.

To summarise our argument in advance, it is this. The components that ordinarily account for learning and executing primitive actions are equivalent to synaptic-level connectivity and adaptation in the brain. These change at a relatively slow pace, either through evolution (learning by the species) or through development and lifelong-learning by the agent. The components that account for the context-appropriate expression of actions — that determine action ordering — are analogous to the electrical action potential in the brain. That is, dynamic planning systems correlate to what cells are “firing” at any instant. This accounts for changes of behaviour on the order of tens of milliseconds. These two types of components are the focus of this section.

Action selection in nature also depends on a chemical / hormonal systems which provides a third temporal duration for intelligent control. Like the electrical system, the chemical control is temporary, (although both chemical and electrical selection can have lasting secondary impact by influencing long-term learning.) However, the chemical systems have durations measured in minutes or longer — across entire episodes of behaviour. During this period they influence metabolic shifts, thus affecting what actions are available to an agent and how they are expressed, and provide additional context for the instantaneous / electrical action selection system, influencing its ‘decisions’. AI equivalents of this sort of control will be described in the section after this one, titled Durative State.

Action

Every AI (and natural intelligence, NI) system must start from some set of primitive affordances — control the agent's intelligence can access. At the most basic level this may be seen as degrees of freedom for the agent. For both NI and AI, generally ‘intelligent’ control is exercised at a more abstract level. In NI, *what* intelligence controls co-evolves and co-develops along with *how* it is controlled. Skill compilation is also a popular idea in AI (see e.g. Newell 1990) although to date it is not often used in commercial / production-quality systems due to issues of combinatorics and the fact that systems that learn in unregulated ways can often break themselves (Bryson & Hauser 2002)¹.

One widely-utilised methodology for creating action primitives in real-time, production-quality systems is to use modularity. This approach was championed as behavior-based AI by Brooks; Brooks (1986; 1991). There have been a number of different approaches to building such modules. In BOD, the ultimate purpose of a module is to produce an action, but the module must also support any perception necessary to do produce its actions (Bryson & Stein 2001b). This is in common with Brook's behavior-based approach. However, perception and action both require memory and learning. Memory is required for calibration, category formation, and the prior expectation gathering necessary for discriminating contexts which if sampled as a snapshot would be perceptually ambiguous. Learning is neces-

¹Although see Whiteson, Taylor, & Stone (2007) for a recent example of a system that simultaneously learns both abstraction and control using two different learning systems.

sary for updating memory. Thus in BOD we propose that *all* variable memory and therefore learning should be contained within modules as well. The purpose of memory is to facilitate perception or action, thus it makes sense to couple the memory inside the modules that require the memorised state to perceive and act.²

One advantage of all learning being modular is that it is more combinatorially tractable and therefore more likely to succeed. Representations can be biased per required learning situation, as seems to be the case in nature (Bryson & Hauser 2002; Gallistel *et al.* 1991). On the other hand, the cost of specialised learning is being context or domain specific (Carruthers 2003). As we mentioned in the introduction to this section, this is the place in our ontology equivalent to the fundamental hardware of neural synapse connections and weights — both what an animal *can* learn, and what it *has* so far in its life. Note that *learning* in this context accounts for everything from very short-term perceptual memory to long-term learning of content such as semantic concepts or navigation maps (Bryson 2001).

Sequencing

Although many AI researchers have attempted to use a single, homogeneous system for the actual arbitration or ordering of expressed actions (Brooks 1986; Maes 1991; Schaal & Atkeson 1994), there is a good deal of empirical evidence that having dedicated representations for sequencing actions is useful, and perhaps in some combinatorial sense, necessary (Georgeff & Lansky 1987; Nilsson 1994; Laird & Rosenbloom 1996; Bonasso *et al.* 1997; Bryson 2000b). Most commercial AI systems use finite state machines (FSM) for this purpose. Brooks (1986) also uses FSM at a low level and linear hierarchy at a high level to help order his control, even though his “Subsumption Architecture” is better known for fully distributed control.

Note that the issue of dedicated sequence representation is no longer debated in the NI literature, and has not been for some time (Lashley 1951; Glasspool 2005; Davelaar 2007). Again this strategy may be due to the combinatorial constraints on finding the appropriate ordering of actions.

As mentioned earlier, we consider this aspect of action selection analogous to the *electrical* aspect of NI behaviour generation — which neurons are firing *now*, what are we currently doing? “Which neurons are firing” is determined by the long-term configuration of the intelligent system (which we have discussed under “Action” above) combined with the the current environmental context (as the agent perceives it) and by the agent’s own motivational context. The external environment may vary extremely rapidly, and thus demands rapid action control as well as perception. However, successful execution of complex goals (including flight from predators or reproductive acts) often requires greater persistence than is easily provided by reactive action selection. This is provided by the motivation or global attention of the system, which we discuss next.

²In fact, variable state (as in OOD) serves as the basis of behaviour decomposition in BOD. (Bryson 2001).

Motivation

Most intelligent systems can be most easily described (and built) by accounting for action selection at two different levels. The first level pertains to action selection particular to individual tasks, and the second to more global motivational state. This simplifies design for AI, and apparently also facilitates evolution, since we see such separate systems throughout animal intelligence (Carlson 2000; Prescott 2007).

Most commercial AI systems (that is, ones that really need to work) have a completely separate mechanism for determining high-level goals than from executing actions. This is also true of some academic AI systems (e.g. Georgeff & Lansky 1987); however in general academics have tried to create more elegant or uniform control systems.

In developing BOD, we called the main aggregate components of the task hierarchy (besides simple primitives and sequences, already mentioned) Basic Reactive Plans (Bryson & Stein 2001a). At each flexible layer of the hierarchy is essentially a cond, switch or case statement (depending on your programming language of choice). Each component of that level and module of a plan has a set of (possibly composite, otherwise primitive) action elements which are utilised for bringing about the subgoal that component is built to achieve. These elements are prioritised, with the highest priority going to the consummatory action (e.g. eating), and the next highest priority to a step that enables the consummatory action (e.g. holding food), and so on (see also Nilsson 1994). Each element is paired with a perceptual precondition, which determines whether the action can be executed in the current environmental context. For example, if the consummatory act of a plan segment is to eat a boiled egg, a peeled hard-boiled egg must be present. If no such peeled egg is available, a hard-boiled egg must be peeled; if there is no hard-boiled egg one must be acquired. Such a dynamic plan is both robust to unexpected action failure (dropping an egg) and opportunistic in exploiting available resources (finding cooked ones in the refrigerator).

BOD uses the Parallel-rooted, Ordered, Slip-stack Hierarchical (POSH) action-selection representation for dynamic planning (Bryson & McGonigle 1998; Bryson 2000b; Bryson & Stein 2001b). The bulk of this system is as just described. However, the *top* level of a POSH plan hierarchy, is a different but related construction to the basic reactive plan. Here the elements are the global goals for the agent. POSH specifies only a very few differences between this top level of the hierarchy and any of its similar lower-level plan components.

- The top level is revisited at every action selection cycle (many times a second). This is not true of the rest of the hierarchy or there would be no focus, no concept of attention, and the benefits of hierarchy would be lost. On every plan cycle, *only* the top elements and the current plan elements are visited. This mechanism entails an interesting assumption about the sorts of things that should be considered high-level goals — they are the sorts of things that you need to be able to switch your attention to even at the cost of disrupting your progress on another task.

- The top level has a third sort of precondition available (besides priority and perceptual context): frequency of expression. This allows coding behaviours that are very important but should not completely monopolise the agent’s time. For example, looking around to check for predators or updating your beliefs about nearby obstacles based on your sonar sensing. These can be placed at a high priority but a low frequency, so that other lower-priority but useful goals can be pursued.

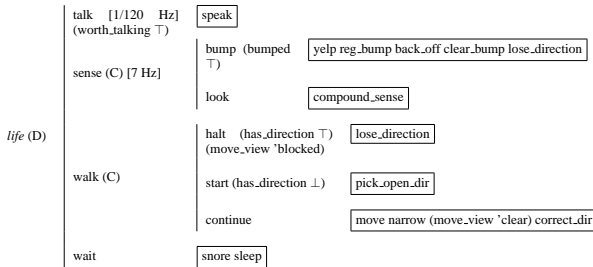


Figure 1: Illustrating the use of frequency-dependent top-level goals — A POSH dynamic plan for keeping a robot moving in a crowded environment. The robot only talks if it is running short on battery — an important but rare and gradual event. *compound_sense* reports fused sonar, IR and bump sensor history input; the fusion occurs in a behaviour module. From Bryson & McGonigle (1998).

This system worked reasonably well on mobile robots (see Figure) and real-time games (Partington & Bryson 2005), and very well in discrete-time artificial life simulations (Bryson 2000b). However, scheduling via frequencies of action-selection attention is difficult, particularly for any action element more complex than a single primitive act. Also, these two systems are not really sufficient to account for a full range of life-long agent behaviour, where agents oscillate over the day between pursuit of widely divergent goals (Dunbar 1997). We now turn to describing our more recent work extending and improving this system.

Durative State

As we introduced at the beginning of the last section, action selection in nature always includes a chemical / hormonal component, which provides additional action-selection context, presumably for reducing dithering and focusing attention while an episode of behaviour is completed. In biological agents, this chemical internal context is controlled either by dedicated organs (such as the adrenal gland) or as part of broader metabolic processes, such as digestion or ovulation. Some such “neurotransmitter” controls for excitation and inhibition actually predate neurons phylogenetically (Carlson 2000). As such, this area which is broadly (though not entirely) neglected in AI may be the most essential form of NI action selection.

Drives

Drives (like hunger and procreation) are obviously central to intelligent control. Our original idea was: Since drives fluctuate in level, they are variable state, and as such under BOD

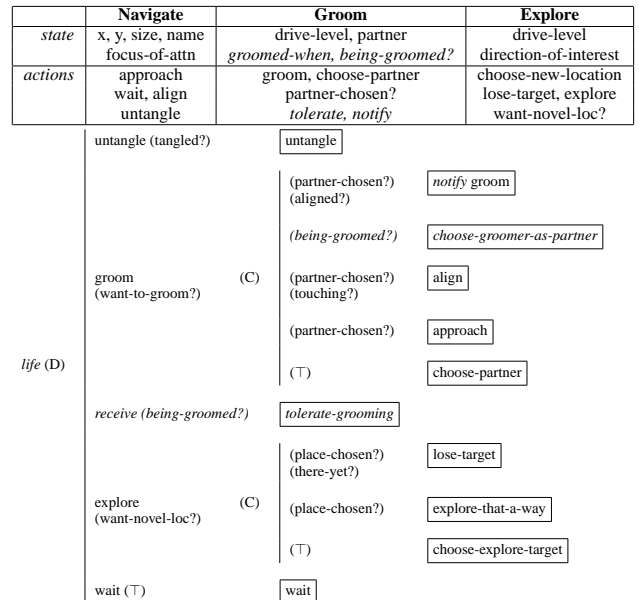


Figure 2: Plan using drive levels to force slow behaviour oscillations. Time between bouts of grooming and wandering determined by tuning of thresholds, amount of consummation from actions and the rate the drive increases over time when unattended (Bryson 2003).

should be represented as an action-generating behaviour module. We have had some success with this approach (Bryson 2003). We created a dedicated VariableDrive subclass of the behaviour modules, with these attributes:

- a *level*. For drives, this increases gradually with time, while consummatory actions can reduce it. Emotions are the inverse — the passage of time gradually reduces the level, while perceptual events can increase it dramatically. These are sufficiently similar that in all probability one AI representation can suit both needs, but in 2003 we focus on drives.
- a *latch*. A latch consists of two thresholds and a single bit (true-false) of information. The first threshold determines when the level of a drive has reached a level such that, if the agent had not been currently influenced by the drive, it now will be; while the second threshold represents the point at which the drive will stop influencing behaviour if it previously had been. The bit simply records whether the drive’s behaviour is active when the level is between the two thresholds. Without latching an intelligent agent will suffer from dithering — a lack of persistence that results from trivially satisfying a need (e.g. eating) the small amount required to make some other need (e.g. drinking) seem more important.

We use these special drive modules to provide context for our top-level global-attention-selecting POSH component described earlier. Thus if the hunger module says it is hungry, the agent will eat, unless some other higher-priority element is triggered (e.g. fleeing from predators). Figure 2

shows a plan that, coupled with drives from “grooming” and “wandering” (foraging) lead simulated monkeys to alternate smoothly between social and individual behaviour.

Emotions

As we said earlier, from a control-theoretic standpoint, emotions are similar to drives. Each has a characteristic set of primitive actions or perceptions which affect its level, and a set of behaviours that tend to be expressed when the level reaches a critical threshold. Also, the passage of time affects the levels as well, although here emotions are typically the inverse of drives. Drives typically build gradually in urgency (although some are facilitated by perceptual stimuli, e.g. the presence of a courting male) and are placated relatively quickly by consummatory actions (e.g. eating). Emotions are triggered relatively quickly by perceptual events (including proprioceptive ones), and decay gradually over time, although here again the gradual dissipation can be facilitated by new information / perceptual events.

Emotions are also relatively easy to observe, and humans are biologically tuned to be good at observing them, because they tend to be displayed on the face. The problem of making photo-realistic talking-head avatars motivated us to examine more closely the problem of accurately representing emotional state. This work was done relatively *ad hoc* (though inspired by psychological research), and as such does not represent scientific evidence, but it does give us a better feel for some of the apparent problems and dependencies for building biologically accurate emotions system. This in turn will lead us back to a discussion of drives and action selection.

Tanguy (2006) proposes an architecture where durative state (e.g. emotions and moods) is stored in a series of modules. Each module includes three *dimensions*, one each for governing the onset, sustain and decay shapes of that entity’s state dynamics. These modules receive their input through a set of dynamic filters, which are themselves influenced by both that and other modules’ current state. Raw input to the system comes from brief *emotional impulses* which are tagged by an appraisal mechanism to indicate the type of events the agent has just experienced. These influence the agent’s reaction, through the combined basis of a hard-wired network (e.g. angry events increase anger and decrease happiness) and the dynamic filtering (the angry event’s have less impact of the agent’s mood is good. The mood is altered negatively if the emotions change rapidly.) This results in a system whereby the same set of events can have very different impact on an agent’s emotional state, given its *previous* emotional state (see Figure 3).

This systems is significantly more complicated than most AI emotion systems, involving a good deal of hand coding for the different emotions. Tanguy generated a large variety of facial expressions by allowing compatible emotions to be active simultaneously and give mixed facial expression, including intentional (or “fake”) emotional gestures overlaying real emotional state. This is in contrast to winner-take-all emotion systems typically used for agent control (Breazeal 2003). Biological evidence indicates that each emotion / behaviour system of intelligent control evolved more-or-less

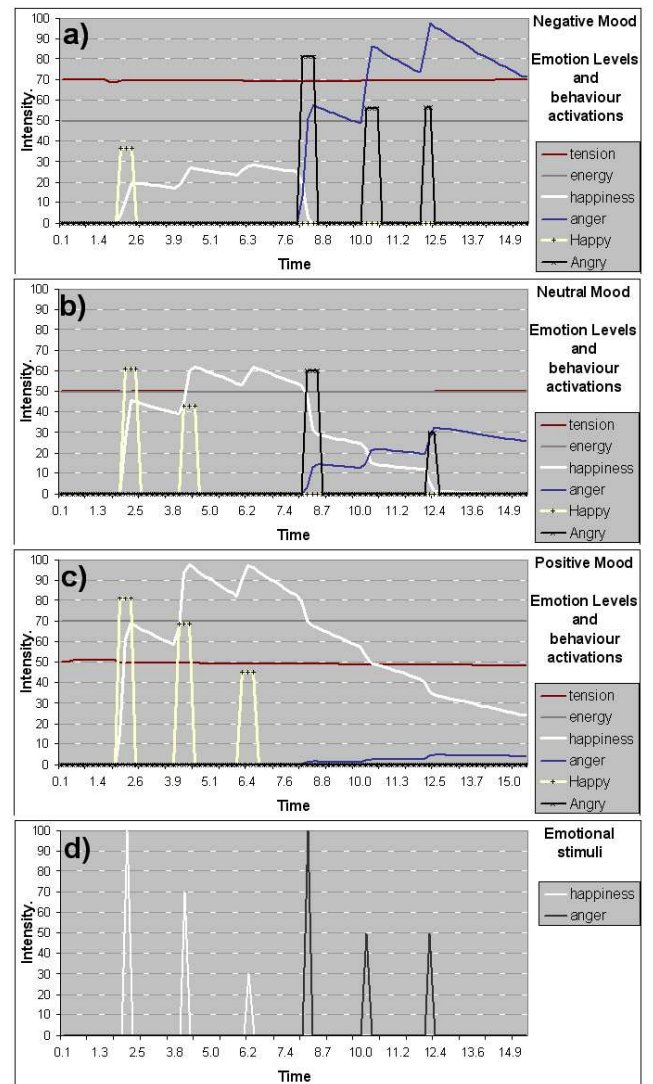


Figure 3: Changes of emotional states at three different durations (expressed behaviour, classical emotion and mood). The same six perceived events have very different consequences to the system depending on its prior state. Graphs a, b and c show state changes in the contexts of *Negative Mood*, *Neutral Mood* and *Positive Mood*, respectively. Graph d shows the time course and strength of the emotionally-salient events (Tanguy, Willis, & Bryson 2007).

independently (LeDoux 1996), and presumably also have a wide variety of interacting excitation and inhibition systems. This implies a high level of heterogeneity which may make close replication of human-like systems necessarily messy.

Overall, Tanguy’s system is useful for generating believable and varied behaviour from a limited set of recognisable external events and of internal interaction scripts. This can be useful for situations such as games, FAQ-bots and personal digital assistants, where for technical reasons the set of available programmed behaviours maybe limited, but some variation in the expression of these behaviours can be in-

roduced unwittingly by the users themselves, by generating varied sequences of interactions and thus varied emotional states in the avatars. However, the complexity of the system makes it impractical for some situations e.g. rapid prototyping and results analysis. Consequently we have not so far introduced its full form into our scientific simulations. However, the practice of creating limited forms of mutual inhibition between some emotions (and not others) lead us to a proposed solution for a problem with our latched drive mechanism, which we describe next.

Beyond Latches: Consciousness?

A problem with the latching system described earlier emerged as we began making more complex primate social simulations. As mentioned earlier, once a drive has passed the threshold required for it to begin influencing behaviour, it continues to do so until either a higher-priority activity interrupts, or another threshold indicating satiation has been reached. However, sometimes a dynamic external environment³ intervenes in a way that causes an *external* interruption. Return to our boiled egg example. If an agent is eating boiled eggs until satiated, and then runs out of eggs when it is not quite fully satiated, would it really start boiling more eggs? In our primate simulation, a monkey that had not quite satisfied itself with grooming might pursue another grooming partner for five minutes, only to satiate after just five seconds of further grooming.

Our suggestion is that when an agent is reliably engaged in an activity consummatory to its goal (or, indeed, is failing to engage in a consummatory action for some length of time) it may essentially reset its lock status and “reconsider” its current top level drive. Such a system still escapes dithering, but also provides a more efficient pursuit of goals — a heuristic mechanism for sensing when expectations have been violated and it is sufficiently likely that a goal may no longer be worth pursuing to justify reassessment. The cost is only adding mechanisms for recognising external interruption, reassessing priorities after interruption is the same as initial assessment. Our early results indicate a massive increase in over-all efficiency when the agents abandon sufficiently satisfied goals in the face of interruption (see Figure 4).

Interestingly, at least two established models of consciousness are similar to our new model of flexibly-latched drives. Norman & Shallice (1986) describe consciousness as a higher-cost attentional system which is brought on line whenever the more basic, reliable, low-cost action-sequencing mechanism is unable to proceed. More recently, Shanahan (2005) proposes a model of mutually-inhibiting motives in a global workspace. We do not agree with Shanahan that such models can account for all of action selection (see e.g. the Tyrrell 1994 critique of Maes 1991). However, his model is similar to what we propose here for arbitration between certain types of high-level tasks. Of course, a simple system of eliciting drive levels and (possibly) weighing them against expected costs does not ex-

³“The real world” isn’t quite the right phrase for virtual agents, but the concept is similar.

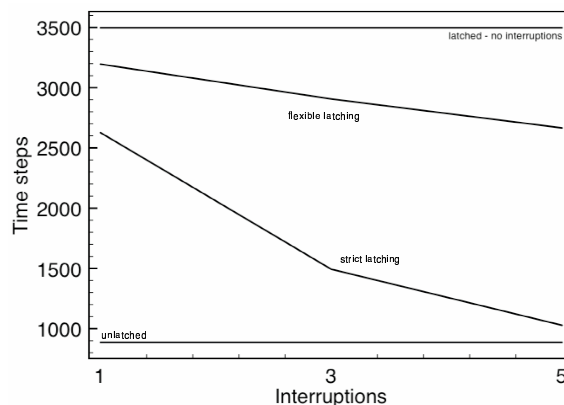


Figure 4: Performance of three action-selection systems measured in terms of the number of time steps allocated to low-level priorities. This measure indicates the overall efficiency with which higher level goals are met while the agents’ basic needs are satisfied sufficiently for it to stay “alive”. Latching is always more efficient than dithering, but in the case of external interruptions strict latching is significantly less so. Flexibility in reassigning latched drives after an interruption ameliorates this problem (Rohlfshagen & Bryson 2008).

plain all the phenomena ordinarily associated with the term *consciousness*. That term is a repository of aggregated folk-psychological theories for aspects of behaviour ranging from perceptual experience through self-concept and on to the soul (Hobson, Pace-Schott, & Stickgold 2000; Dennett 2001). We note only in passing a possible similarity in functional outcome between these systems.

Returning to POSH, flexible latching requires that some sets of elements in the top-level of the drive hierarchy should share the same priority level, and determine control more or less “between themselves” if action selection passes to that priority level. This determination is made on the bases of need as modelled by internal variables, analogous to the hormonal signalling provided by NI drive and emotion systems. Notice that drives should not be mistaken for chemical signals of true starvation or somesuch. Rather, they represent periodic shifts in motivation intended to guarantee that the system as a whole stays well within a safe operating conditions and is able to pursue ancillary, non-safety-critical goals such as procreation (Rohlfshagen & Bryson 2008). We have successfully implemented and evaluated a basic version of this system (see Figure 4). We are now evaluating its usability in the context of modelling individual primate social behaviours such as coalition formation, behaviour associated with emotions as well as drives.

Conclusions

In this paper we have argued that emotions and other durative states such as drives, mood and possibly conscious attention represent a fundamental component of biological action selection — an intervening representation controlling behaviour for durations shorter than memory but longer than

the current instant's reaction. We have described an ontology of action selection which we have previously shown to be wide-spread in use. Here we focussed mostly on the ontology's expression in our research tools for supporting AI systems development⁴. We then described a series of extensions to that ontology, all motivated by a need to improve arbitration between high-level goals while maintaining real-time performance of the system in ordinary action selection.

References

- Bonasso, R. P.; Firby, R. J.; Gat, E.; Kortenkamp, D.; Miller, D. P.; and Slack, M. G. 1997. Experiences with an architecture for intelligent, reactive agents. *Journal of Experimental and Theoretical Artificial Intelligence* 9(2/3):237–256.
- Breazeal, C., and Scassellati, B. 1999. How to build robots that make friends and influence people. In *International Conference on Intelligent Robots (IROS-99)*, 858–863.
- Breazeal, C. 2003. Emotion and sociable humanoid robots. *International Journal of Human-Computer Studies* 59(1–2):119–155.
- Broekens, J.; Kusters, W. A.; and Verbeek, F. J. 2007. Affect, anticipation, and adaptation: Affect-controlled selection of anticipatory simulation in artificial adaptive agents. *Adaptive Behavior* 15(4):397–422.
- Brom, C.; Gemrot, J.; Bida, M.; Burkert, O.; Partington, S. J.; and Bryson, J. J. 2006. Posh tools for game agent development by students and non-programmers. In Mehdi, Q.; Mtenzi, F.; Duggan, B.; and McAtamney, H., eds., *The Ninth International Computer Games Conference: AI, Mobile, Educational and Serious Games*, 126–133. University of Wolverhampton.
- Brooks, R. A. 1986. A robust layered control system for a mobile robot. *IEEE Journal of Robotics and Automation* RA-2:14–23.
- Brooks, R. A. 1991. Intelligence without representation. *Artificial Intelligence* 47:139–159.
- Bryson, J. J., and Hauser, M. D. 2002. What monkeys see and don't do: Agent models of safe learning in primates. In Barley, M., and Guesgen, H. W., eds., *AAAI Spring Symposium on Safe Learning Agents*.
- Bryson, J. J., and Leong, J. C. S. 2006. Primate errors in transitive 'inference': A two-tier learning model. *Animal Cognition* 10(1):1–15.
- Bryson, J. J., and McGonigle, B. 1998. Agent architecture as object oriented design. In Singh, M. P.; Rao, A. S.; and Wooldridge, M. J., eds., *The Fourth International Workshop on Agent Theories, Architectures, and Languages (ATAL97)*, 15–30. Providence, RI: Springer.
- Bryson, J. J., and Stein, L. A. 2001a. Architectures and idioms: Making progress in agent design. In Castelfranchi, C., and Lespérance, Y., eds., *The Seventh International Workshop on Agent Theories, Architectures, and Languages (ATAL2000)*, 73–88. Springer.
- Bryson, J. J., and Stein, L. A. 2001b. Modularity and design in reactive intelligence. In *Proceedings of the 17th International Joint Conference on Artificial Intelligence*, 1115–1120. Seattle: Morgan Kaufmann.
- Bryson, J. J., and Thórisson, K. R. 2000. Dragons, bats & evil knights: A three-layer design approach to character based creative play. *Virtual Reality* 5(2):57–71.
- Bryson, J. J.; Ando, Y.; and Lehmann, H. 2007. Agent-based models as scientific methodology: A case study analysing primate social behaviour. *Philosophical Transactions of the Royal Society, B — Biology* 362(1485):1685–1698.
- Bryson, J. J.; Caulfield, T. J.; and Drugowitsch, J. 2005. Integrating life-like action selection into cycle-based agent simulation environments. In North, M.; Sallach, D. L.; and Macal, C., eds., *Proceedings of Agent 2005: Generative Social Processes, Models, and Mechanisms*, 67–81. Chicago: Argonne National Laboratory.
- Bryson, J. J. 2000a. Cross-paradigm analysis of autonomous agent architecture. *Journal of Experimental and Theoretical Artificial Intelligence* 12(2):165–190.
- Bryson, J. J. 2000b. Hierarchy and sequence vs. full parallelism in reactive action selection architectures. In *From Animals to Animats 6 (SAB00)*, 147–156. Cambridge, MA: MIT Press.
- Bryson, J. J. 2001. *Intelligence by Design: Principles of Modularity and Coordination for Engineering Complex Adaptive Agents*. Ph.D. Dissertation, MIT, Department of EECS, Cambridge, MA. AI Technical Report 2001-003.
- Bryson, J. J. 2003. Where should complexity go? Cooperation in complex agents with minimal communication. In Truszkowski, W.; Rouff, C.; and Hinchey, M., eds., *Innovative Concepts for Agent-Based Systems*, 298–313. Springer.
- Carlson, N. R. 2000. *Physiology of Behavior*. Boston: Allyn and Bacon.
- Carruthers, P. 2003. The cognitive functions of language. *Brain and Behavioral Sciences* 25(6).
- Davelaar, E. J. 2007. Sequential retrieval and inhibition of parallel (re)activated representations: A neurocomputational comparison of competitive queuing and resampling models. *Adaptive Behavior* 15(1). special issue on Mechanisms of Action Selection, in press.
- De Rosis, F.; Pelachaud, C.; Poggi, I.; Carofiglio, V.; and De Carolis, B. 2003. From greta's mind to her face: Modelling the dynamics of affective states in a conversational agent. *International Journal of Human Computer Studies* 59:81–118.
- Dennett, D. C. 2001. Are we explaining consciousness yet? *Cognition* 79:221–237.
- Dunbar, R. 1997. *Grooming, Gossip, and the Evolution of Language*. Harvard University Press.
- Gadanhó, S. C. 1999. *Reinforcement Learning in Autonomous Robots: An Empirical Investigation of the Role of Emotions*. Ph.D. Dissertation, University of Edinburgh.
- Gallistel, C.; Brown, A. L.; Carey, S.; Gelman, R.; and Keil, F. C. 1991. Lessons from animal learning for the

⁴All software described here is available for download, linked from the AmonI Software page.

- study of cognitive development. In Carey, S., and Gelman, R., eds., *The Epigenesis of Mind*. Hillsdale, NJ: Lawrence Erlbaum. 3–36.
- Georgeff, M. P., and Lansky, A. L. 1987. Reactive reasoning and planning. In *Proceedings of the Sixth National Conference on Artificial Intelligence (AAAI-87)*, 677–682.
- Glasspool, D. W. 2005. The integration and control of behaviour: Insights from neuroscience and AI. In Davis, D. N., ed., *Visions of Mind: Architectures for Cognition and Affect*. Idea Group. 66–89.
- Hiller, M. J. 1995. The role of chemical mechanisms in neural computation and learning. Technical Report AITR-1455, MIT AI Laboratory.
- Hobson, J. A.; Pace-Schott, E. F.; and Stickgold, R. 2000. Dreaming and the brain: Toward a cognitive neuroscience of conscious states. *Behavioral and Brain Sciences* 23(6):793–842.
- Laird, J. E., and Rosenbloom, P. S. 1996. The evolution of the Soar cognitive architecture. In Steier, D. M., and Mitchell, T. M., eds., *Mind Matters*. Erlbaum.
- Lashley, K. S. 1951. The problem of serial order in behavior. In Jeffress, L. A., ed., *Cerebral mechanisms in behavior*. New York: John Wiley & Sons.
- LeDoux, J. 1996. *The Emotional Brain: The Mysterious Underpinnings of Emotional Life*. New York: Simon and Schuster.
- Maes, P. 1991. The agent network architecture (ANA). *SIGART Bulletin* 2(4):115–120.
- Marcella, S., and Gratch, J. 2002. A step toward irrationality: Using emotion to change belief. In *Proceedings of the First international joint conference on autonomous agents and multiagent systems*, 334–341. Bologna, Italy: ACM Press.
- Minsky, M.; Singh, P.; and Sloman, A. 2004. The St. Thomas common sense symposium: Designing architectures for human-level intelligence. *AI Magazine* 25(2):113–124.
- Morgado, L., and Gaspar, G. 2005. Emotion based adaptive reasoning for resource bounded agents. In *The 4th International Joint Conference on Autonomous Agents and Multi Agent Systems (AAMAS '05)*, 921–928. Utrecht: ACM.
- Newell, A. 1990. *Unified Theories of Cognition*. Cambridge, Massachusetts: Harvard University Press.
- Nilsson, N. J. 1994. Teleo-reactive programs for agent control. *Journal of Artificial Intelligence Research* 1:139–158.
- Norman, D. A., and Shallice, T. 1986. Attention to action: Willed and automatic control of behavior. In Davidson, R.; Schwartz, G.; and Shapiro, D., eds., *Consciousness and Self Regulation: Advances in Research and Theory*, volume 4. New York: Plenum. 1–18.
- Norman, D. A.; Ortony, A.; and Russell, D. M. 2003. Affect and machine design: Lessons for the development of autonomous machines. *IBM Systems Journal* 38–44.
- Partington, S. J., and Bryson, J. J. 2005. The behavior oriented design of an unreal tournament character. In Panayiotopoulos, T.; Gratch, J.; Aylett, R.; Ballin, D.; Olivier, P.; and Rist, T., eds., *The Fifth International Working Conference on Intelligent Virtual Agents*, 466–477. Kos, Greece: Springer.
- Prescott, T. J.; Bryson, J. J.; and Seth, A. K. 2007. Modelling natural action selection: An introduction to the theme issue. *Philosophical Transactions of the Royal Society, B — Biology* 362(1485):1521–1529.
- Prescott, T. J. 2007. Forced moves or good tricks in design space? landmarks in the evolution of neural mechanisms for action selection. *Adaptive Behavior* 15(1):9–31.
- Rohlfshagen, P., and Bryson, J. J. 2008. *in prep.*
- Schaal, S., and Atkeson, C. G. 1994. Robot juggling: An implementation of memory-based learning. *Control Systems Magazine* 14(1):57–71.
- Shanahan, M. P. 2005. Global access, embodiment, and the conscious subject. *Journal of Consciousness Studies* 12(12):46–66.
- Sloman, A., and Croucher, M. 1981. Why robots will have emotions. In *Proceedings of the 7th International Joint Conference on Artificial Intelligence (IJCAI '81)*, 1537–1542. Vancouver: William Kaufmann.
- Tanguy, E. A. R.; Willis, P. J.; and Bryson, J. J. 2006. A dynamic emotion representation model within a facial animation system. *The International Journal of Humanoid Robotics* 3(3):293–300.
- Tanguy, E.; Willis, P.; and Bryson, J. J. 2007. Emotions as durative dynamic state for action selection. In *Proceedings of the 20th International Joint Conference on Artificial Intelligence*, 1537–1542. Hyderabad: Morgan Kaufmann.
- Tanguy, E. A. R. 2006. *Emotions: The Art of Communication Applied to Virtual Actors*. Ph.D. Dissertation, Department of Computer Science, University of Bath, United Kingdom. Bath CS Technical Report CSBU-2006-06.
- Tomb, I.; Hauser, M. D.; Deldin, P.; and Caramazza, A. 2002. Do somatic markers mediate decisions on the gambling task? *Nature Neuroscience* 5(11):1103–1104.
- Tyrrell, T. 1994. An evaluation of Maes's bottom-up mechanism for behavior selection. *Adaptive Behavior* 2(4):307–348.
- Velásquez, J. D., and Maes, P. 1997. Cathexis: A computational model of emotions. In *Proceedings of the First International Conference on Autonomous Agents (Agents '97)*, 518–519. New York, NY, USA: ACM.
- Whiteson, S.; Taylor, M. E.; and Stone, P. 2007. Empirical studies in action selection for reinforcement learning. *Adaptive Behavior* 15(1). special issue on Mechanisms of Action Selection, *in press*.
- Zadeh, S. H.; Shouraki, S. B.; and Halavati, R. 2006. Emotional behaviour: A resource management approach. *Adaptive Behaviour*.