Emotions and Action Selection in an Artificial Life Model of Social Behavior in Non-Human Primates

Joanna Bryson and Jessica Flack

Artificial Intelligence Laboratory, MIT
545 Technology Square, Cambridge MA 02139, USA
joanna@ai.mit.edu

The Yerkes Regional Primate Research Center
202 Psychology Building, Emory University, Atlanta GA 30322, USA
jflack@emory.edu

Abstract. This is an extended abstract describing work in progress. We are developing an artificial life (ALife) model the various sorts of social behaviors displayed by colonies of non-human primates. We hope to use this ALife model to support work on an ethological model also under development which examines the relationship between conflict management behaviors displayed by a species and the dominance relationships between individuals of that species. In this paper, we describe the relationship between emotions and action selection in our ALife model.

1 Introduction

The topic of emotions is losing its taboo both in artificial intelligence and in the animal sciences. Nevertheless, emotions seem necessarily an emotional subject, often raising unusually passionate responses, both in support and in criticism for systems and theories. This paper describes work on modeling complex social activity in non-human primates. As a byproduct of this work, we are integrating emotional responses into complex agent control.

By all indications, the various phenomena we know as emotions characterize a set of behaviors that evolved at significantly different points in our ancestral history [LeDoux, 1996]. Emotions are effectively an intervening variable used to explain categories of species-typical behaviors that are related not only by the behaviors and the environmental contexts in which they tend to be displayed, but by expressive body postures in the behaving animal. These emotion “variables” have a real, biological correlate: relatively slow and diffuse chemical transmissions within the brain (and the rest of the organism) which create a significant internal context for the brain’s operation, affecting both perception and action selection.

This paper demonstrates agents exhibiting complex social behavior which is typically associated with emotions. The work demonstrates the implementation of such behaviors into a general-purpose agent action selection mechanism. Emotional state is represented by drive-like variables, in an approach at least somewhat analogous to that of Cañamero [1997] and Frankel and Ray [2000]. Emotional responses are expressed by modular *behaviors* and coordinated by reactive plans. This action selection system is an incremental advance on standard hybrid software agent control but is also biologically plausible [Bryson and Stein, 2001a,b].

Our agents are designed to model animal behaviors that humans readily describe as emotional — the interactions between individuals in a primate colony. Our current model shows the animals oscillating between two "drives", the desire to groom, and the desire for privacy. Grooming is an interesting behavior, associated with bonding between animals and a calming effect on the recipient. Although most primates seem to derive pleasure from grooming, they normally engage in this behavior relatively infrequently. Frequency of grooming tends to increase in times of certain social stresses.

The "desire for privacy" in our model stands in for a number of other ways primates spend their time, such as foraging and napping. We model only seeking isolation from other agents for simplicity's sake. For monkeys living in a community, one monkey's desire to groom can interfere with another monkey's desire for privacy. There are a number of possible solutions to such conflict [de Waal, 2000], but for our pilot study we are only modeling two: *tolerance* and *aggression*. Aggression is of course associated with two other emotions, anger and fear.

In the remainder of this abstract, we describe in more detail the ethological model we are examining, and the ALife model we are developing. We conclude a short section about the current status of the ALife work, and the format of the experiments and results we expect to report in the final paper.

2 Basic Primate Conflict Management

One of the most interesting questions in the study of animal societies is how individuals negotiate their social relationships. This question of how conflict among lower level units (individual group members) is regulated in the formation of higher level units (societies) has been described as the fundamental problem in ethology, [Egbert Giles Leigh, 1999]. Although research on non-human primate societies indicates that there are a variety of mechanisms — such as aggression, social tolerance, and avoidance — by which conflict is managed or resolved [de Waal, 2000], it is not well understood how and why the expression of these mechanisms varies across and even within social systems. For example, there is tremendous variation across the macaque genus in terms of how conflict is managed despite similar patterns of social organization. Aggression in some species is common and severe while in others, it is extremely frequent but rarely escalates to levels that produce injuries [de Waal and Luttrell, 1989]. Corresponding to this variation in the degree to which aggression is employed to settle conflicts of interest is variation in the degree of social tolerance by dominant individuals of subordinate ones, particularly in the context of resource acquisition, and variation in the degree to which relationships damaged by aggression are repaired via reconciliation [de Waal and Luttrell, 1989]. Although it appears that this co-variation in conflict management mechanisms varies in predictable ways across species, it does not appear that the co-variation is species-specific. Rather, the variation seems to be emergent from patterns of social interaction among individuals, and self-reinforced through social learning.

The importance of social learning on styles of interaction was made clear by the results of a cross-fostering study of two macaque species the individuals of which have drastically different proclivities for aggression and reconciliation [de Waal and Johanowicz, 1993]. In this study, juvenile rhesus macaques, which typically live in social systems characterized by high levels of severe aggression and low levels of reconciliation, were cross-fostered with slightly older "tutor" stumptailed macaques, which live typically in social systems characterized by high levels of mild aggression and high levels of reconciliation. Over the course of the study, the rhesus monkeys learned to reconcile more frequently and adopted the stumptailed style of interaction, and even retained this pattern after all stumptail tutors were removed. Although this type of study indicates that social learning plays a role in what types of conflict management strategies individuals adopt, it remains unclear what role emotion plays as a mediating factor in this process and it remains unclear how differences in interaction patterns at the relationship level translate generate conflict management patterns at the social system level.

3 The Model as it Currently Stands

3.1 Action Selection under Behavior-Oriented Design

Our ALife model is implemented under Behavior-Oriented Design (BOD) [Bryson and Stein, 2001a]. This methodology divides the architecture of an agent's intelligence into two different representations. First, there are modular, semi-autonomous *behaviors*. These are inspired by the behaviors of the behavior-based approach to artificial intelligence (BBAI) [Brooks, 1991, Steels, 1994, Matarić, 1997], which are responsible for encapsulating the perception and action necessary for a particular aspect of an agents expressed behavior. BOD extends this notion to include whatever variable state is necessary for that perception and control to operate, whether this represents short-term memory or lifetime learning. As such, BOD moves the notion of a behavior closer to that of an active object [van Eijk et al., 2001].

Different behaviors may run processes independently and in parallel so long as they do not interfere with other behaviors. For example, a behavior underlying the social aspects of grooming might continuously make note of any other colony-mates that have been spotted recently as possible grooming partners, so that if the agent starts a search it will recall a good area to start. On the other hand, to the extent that a behavior's actions might interfere with those of another behavior, they are arbitrated through a dedicated action selection module. In a BOD agent, such action selection relies on hierarchical reactive plans.

Relying on such plans is a dominant strategy in the control of complex, reactive agents [Hexmoor et al., 1997, Kortenkamp et al., 1998, Bryson and Stein, 2001b]. What distinguishes BOD from other such ("hybrid") architectures is a particularly flexible form of hierarchy referred to as POSH (Parallel-rooted, Ordered Slip-stack Hierarchical), and the fact that the behaviors are not simple primitives manipulated by plans, but are complex, persistent semi-autonomous modules. The plan primitives are *method calls* on the behavior objects, which serve to release the expression of a particular action, which is then governed by the behavior itself.

3.2 Representing Emotions under BOD

The emotional responses of the agents in our simulation are represented exactly as any other behavior — through a combination of reactive plans representing particular orderings of actions (action patterns) and behaviors that determine how and in which way these actions are expressed. The fact that emotional responses are so continuous with normal action selection makes sense in light of understandings of emotions such as proposed by Damasio [1999] which suggests that essentially any species-typical behavior *is* an emotional response, because emotions are central to motivation.

The emotions themselves are represented by simple behaviors which have at their heart a level variable, with time-dependent methods for increasing and decreasing that level. The exact nature of the methods depends on the particular emotion. For example, the desire for isolation increases slowly in the presence of other animals, and decreases slowly when the animal is alone. Fear on the other hand increases radically in the context of a direct threat, and more slowly in the context of a fight between other nearby agents. It decreases slowly in isolation, or more quickly when being groomed out of the context of danger.

Emotional body postures are very much abstracted in our ALife simulation: they are simply colors for the agents. Their expression is also controlled by the emotion behaviors.

3.3 Extensions for a Multi-Agent Simulation

A more significant change to the BOD system was actually caused by this being the first multi-agent application of BOD. Still, this was also a relatively straight-forward improvement. The version of the BOD architecture that we are currently using for our pilot model is implemented in a multi-process implementation of common lisp [Xan, 1999]. We therefore simply dedicated a process to each agent, and relied on the platform to perform the time sharing between agents. Since, under BOD, behaviors and POSH reactive plans are already represented in an object-oriented framework, making the system multi-agent simply required replacing the individual instances of the POSH plan hierarchy and the behaviors with hash tables of such instances, each indexed on the process ID of the agent.

3.4 Model Showing Grooming Oscillating with Privacy

The current system controlling our primates is as follows. First, there are currently four behaviors. The first three, **grooming**, **novelty**, and **explore**, are fairly simple behaviors as described above, controlling latent variables that might be identified with emotions or drives. The fourth, **primate**, has the relatively large responsibility of handling the primates' "bodies" — it controls navigation of the agents around their enclosure.

If the simulations of the primates were particularly complex, **primate** would probably be decomposed into more behaviors. However, the most important output of our system is a simple list of events in the colony such as is produced by primatologists, since this is the information that is being compared to existing models and data. For the purpose of debugging, we also have a GUI representation of the agents, but they are represented by simple buttons, with color indicating their expression, and ASCII characters indicating their identity and some gestures. The observing function that produces the log of activity actually has more information than can be gleaned by observing the GUI, but not more than can generally be gathered by field observers. Although there is of course less noise from ambiguities arising in the field in determining the intended object of a gesture in a crowded location, in general the level of reporting is plausible because it records only expressed behaviors.

```
   CH    Charlie (all names should be unique in first 2 letters)
   CH-   Charlie gesturing right
  =CH    Charlie touching left
  ^CH    Charlie gesturing above
  vCHv   Charlie touching below

Grey           - neutral: normal motion or sitting
Pink           - displaying
Red            - angry (fighting)
Orange         - frightened (screeching)
Lavender       - aroused
Purple         - mounting
Blue           - playing
Green          - grooming
```

Fig. 1. Labels and colors indicating the state and identity of a particular agent in the simulated colony. The only colors in current use are grey and green.

Figure 2 shows the current reactive plan for coordinating potential conflicting actions between these behaviors. In this figure, each vertical line represents a *basic reactive plan* (BRP): a planning element similar to the teleo-reactive plans of Nilsson [1994]. BRPs are hierarchical, and their elements are prioritized. The leftmost one (labeled with a *D*) is a specialized form of a BRP called a *drive collection*. On every cycle of the action-selection module, the drive-collection is checked to set action attention to its highest priority element that can fire given its trigger list. Triggers are in parentheses with desired values. Drive collection elements can also be scheduled so that drives run in course-grained parallel, see further [Bryson, 2000, Bryson and Stein, 2001a].

The content of a drive-collection element is the root of a further plan hierarchy. In this case, these plan hierarchies have only two levels: another BRP, and a primitive action at each leaf of the BRP. The plain-text label at the far left of every plan element is for readability, it does not affect action selection. The bold-face labels are interfaces to the behaviors listed above. Here is the assignment of reactive-plan primitives to behaviors:

– **grooming**: want-to-groom, partner-chosen, groom, choose-grooming-partner
– **explore**: want-new-loc, leave

life *D*	grooming *C* (want-to-groom)	groom-gp (partner-chosen aligned-w-target)	**groom**
		no-vert-gp (partner-overlap)	**go-to-partner-edge**
		align-w-gp (partner-chosen very-near-target)	**engage**
		get-near-gp (partner-chosen)	**approach**
		choose-gp ()	**choose-grooming-partner**
	wandering *C* ()	get-moving (want-new-loc target-chosen touching-target)	**forget-target**
		move-away (want-new-loc target-chosen)	**leave**
		choose-wander (want-new-loc)	**choose-isolated-place**
		sit ()	**wait**

Fig. 2. Reactive plan supporting the coordination of the ALife agents. See text for explanation.

– **primate**: aligned-w-target, go-to-partner-edge, engage, touching-target, approach, wait

The behavior **novelty** is not monitored or controlled by this plan, and in fact has no impact in the primate behavior below.

3.5 Current Results

We are still working on building the behaviors of our primates, and have thus not yet begun quantitative analysis of our results. However the transcript in Figure 3 shows a brief episode in the experience of a colony. When the animals want to be alone, they move towards a location (in 2-D space), when they are attempting to groom they move towards each other. Mutual grooming as seen between Murray and Jean is at this stage coincidental — Roger, who is not particularly in the mood to groom, but not particularly concerned with being isolated, ignores Alice. George, on the other hand, wants to move to an isolated spot, and is being chased rather pathetically by Monica who repeatedly tries to sit down by him as he moves. The agents here do not have any simulation of a “theory of mind” — they cannot yet notice when their company is not wanted.

4 Expected Progress

We hope by August to have at least rudimentary anger and aggression behaviors modeled, and to have results showing the tradeoffs for individuals in the colony between

```
George   APPROACH (0 54)                 485591
 Ringo   APPROACH (0 67)                 487575
Murray       WAIT (374 9)                487881
Monica       WAIT (45 236)               491908
  Jean   APPROACH Murray                 497864
  Jean     ALLIGN Murray                 500125
  Jean      GROOM Murray                 500254
 Alice   APPROACH Roger                  500275
 Alice      GROOM Roger                  503282
Murray   APPROACH Jean                   505554
  Jean   APPROACH Murray                 505772
Murray      GROOM Jean                   506143
  Jean      GROOM Murray                 506237
Monica   APPROACH George                 509439
Monica     ALLIGN George                 510684
Monica   APPROACH George                 510972
Monica     ALLIGN George                 513842
Monica   APPROACH George                 513958
```

Fig. 3. Part of a sample observation log for the simulated primate colony. "Approach" indicates walking, "wait" sitting or foraging, "align" engaging (sitting closely) for interactions, and "groom" is obvious. See comments in text.

meeting individual behavioral goals (e.g. in amount of time spent grooming and exploring) for various levels of conflict and tolerance. We next expect to add a few more complex social behaviors, such as knowing one's place in a dominance hierarchy, and third-party conflict management. However, our programmer is busy with a dissertation so these more exciting behaviors might not be done by August. Nevertheless, the system as it stands now already demonstrates rudimentally an interaction between emotions and action selection.

5 Conclusion

We are in the process of developing what could be an exciting model, both from the perspective of examining primate social interactions, and from the perspective of improving the state-of-the-art for modeling complex agent behaviors. Our model already shows one way to express species-typical emotional responses which incorporates an explicit modeling of individual emotion levels. In this system, "emotions" as a whole are not modeled as special subsystem that interacts with action selection, but rather each emotion is modeled individually. An emotion is one of many internal states used to control the expression of behavior.

Acknowledgements

References

Rodney A. Brooks. Intelligence without representation. *Artificial Intelligence*, 47:139–159, 1991.

Joanna Bryson. Hierarchy and sequence vs. full parallelism in reactive action selection architectures. In *From Animals to Animats 6 (SAB00)*, pages 147–156, Cambridge, MA, 2000. MIT Press.

Joanna Bryson and Lynn Andrea Stein. Modularity and design in reactive intelligence. In *Proceedings of the 17th International Joint Conference on Artificial Intelligence*, Seattle, August 2001a. Morgan Kaufmann.

Joanna Bryson and Lynn Andrea Stein. Modularity and specialized learning: Mapping between agent architectures and brain organization. In Stefan Wermter, Jim Austion, and David Willshaw, editors, *Emergent Neural Computational Architectures Based on Neuroscience.* Springer, 2001b. *in press*.

Dolores Cañamero. A hormonal model of emotions for behavior control. Technical Report 97-06, Vrije Universiteit Brussel, Artificial Intelligence Laboratory, 1997.

Antonio R. Damasio. *The Feeling of What Happens: Body and Emotion in the Making of Consciousness*. Harcourt, 1999.

Frans B. M. de Waal. Primates—a natural heritage of conflict resolution. *Science*, 289: 586–590, 2000.

Frans B. M. de Waal and Denise L. Johanowicz. Modification of reconciliation behavior through social experience: An experiment with two macaque species. *Child Development*, 64:897–908, 1993.

Frans B. M. de Waal and Lesleigh Luttrell. Toward a comparative socioecology of the genus *macaca*: Different dominance styles in rhesus and stumptailed macaques. *American Journal of Primatology*, 19:83–109, 1989.

Jr. Egbert Giles Leigh. Levels of selection, potential conflicts, and their resolution: The role of the ”common good”. In Laurent Keller, editor, *Levels of Selection in Evolution*, pages 15–30. Princeton University Press, Princeton, NJ, 1999.

Carl Frankel and Rebecca Ray. Emotion, intention and the control architecture of adaptively competent information processing. In Aaron Sloman, editor, *AISB'00 Symposium on Designing a Functioning Mind*, pages 63–72, 2000.

Henry Hexmoor, Ian Horswill, and David Kortenkamp. Special issue: Software architectures for hardware agents. *Journal of Experimental & Theoretical Artificial Intelligence*, 9(2/3):147–156, 1997.

David Kortenkamp, R. Peter Bonasso, and Robin Murphy, editors. *Artificial Intelligence and Mobile Robots: Case Studies of Successful Robot Systems*. MIT Press, Cambridge, MA, 1998.

Joseph LeDoux. *The Emotional Brain : The Mysterious Underpinnings of Emotional Life*. Simon and Schuster, New York, 1996.

Maja J. Matarić. Behavior-based control: Examples from navigation, learning, and group behavior. *Journal of Experimental & Theoretical Artificial Intelligence*, 9(2/3): 323–336, 1997.

Nils Nilsson. Teleo-reactive programs for agent control. *Journal of Artificial Intelligence Research*, 1:139–158, 1994.

Luc Steels. A case study in the behavior-oriented design of autonomous agents. In Dave Cliff, Philip Husbands, Jean-Arcady Meyer, and Stewart W. Wilson, editors, *From Animals to Animats (SAB94)*, Cambridge, MA, 1994. MIT Press. ISBN 0-262-53122-4.

Rogier M. van Eijk, Frank S. de Boer, Wiebe van der Hoek, and John-Jules Ch. Meyer. Generalised object-oriented concepts for inter-agent communication. In C. Castelfranchi and Y. Lespérance, editors, *Intelligent Agents VII (ATAL2000)*. Springer, 2001.

Lispworks Professional Edition. Xanalys, Waltham, MA, 4.1.18 edition, 1999. (formerly Harlequin).