A biologically inspired predator-prey study of the effects of emotion and communication on emergent group behavior

Jacqueline Kory

Vassar College

April 2011

Senior thesis submitted in partial fulfillment of
the requirements for the major in Cognitive Science

First Reader: Ken Livingston

Second Reader: Luke Hunsberger

Abstract

Any agent that functions successfully in a constantly changing world must be able to adapt its behavior to its current situation. In biological organisms, emotion is often highlighted as a crucial system for generating adaptive behavior. This paper presents a biologically-inspired predator-prey model to investigate the effectiveness of an emotion-like system in guiding the behavior of artificial agents, implemented in a set of simulated robots. The predator's behavior was governed by a simple subsumption hierarchy; the prey selected actions based on direct sensory perceptions dynamically integrated with information about past motivational/emotional states. Aspects of the prey's emotion system were evolved over time. The first study examined the interactions of a single prey with the predator, indicating that having an emotion system can led to more diverse behavioral patterns, but may not lead to optimal action selection strategies. In the second study, groups of prey agents were evolved. These agents began to utilize alarm signaling and displayed fear contagion, with more group members surviving than in groups of emotionless prey. These results point to the pivotal role emotion plays in social scenarios. The model adds to a critical body of research in which important aspects of biological emotion are incorporated into the action selection mechanisms of artificial agents to achieve more adaptive, context-dependent behavior.

Contents

1. Introduction

In many artificial and robotic systems, inspiration is drawn from the emotion and animal communication literature to get more flexible and adaptive behavior. Examples can be seen in both autonomous robots (e.g., Arkin, 2005; Hasson & Gaussier, 2010; Lee-Johnson & Carnegie, 2007; Miachaud & Audet, 2001; Murphy, Lisetti, Tardif, Irish, & Gage, 2002; and Parisi & Petrosino, 2010) and in many human-robot interaction studies as a way of better enabling humans to interact with the artificial agents (e.g., Arkin, Fujita, Takagi, & Hasegawa, 2003; Bruce, Nourbakhsh & Simmons, 2002; Castellano, Aylett, Paiva, & McOwan, 2008; and Murray & Cañamero, 2009). Emotions, in biological organisms, are often defined as a system that aids in behavior regulation and control, e.g., for choosing between mutually incompatible actions and then preparing the organism's body to perform the selected action (Tooby & Cosmides, 2008). Researchers build models of emotions that are functionally equivalent to portions of biological emotion systems in order to bias reactions and promote the generation of behavior tuned to environmental conditions and appropriate to the agent's goals, whether the goals are survival-oriented as is true for many biological creatures, or whether the goals are specific to some other kind of task, such as the cooperative serving of finger food (see, e.g., Murphy et al., 2002).

Emotions are also involved in signaling and communication. Some researchers suggest that one way of understanding animal signals is to treat them as expressions of affective states that then induce some particular affective state in the receivers of the signal, influencing the future behavior of the receiver in a way that is, on average, beneficial to the signaler (Bradbury & Vehrencamp, 1998; Krebs & Dawkins, 1984). Signals may be exchanged about salient information in the animals' environments, such as the presence of a predator (Seyfarth & Cheney, 2003), a state of distress (Ciani, 2000), or a desire to mate (Rendall, Owren, & Ryan, 2009). Models of animal communication can also be incorporated into artificial agents, again with the goal of generating more adaptive behavior (see, e.g., Jim & Giles, 2000; Loula, Gudwin, El-Hani, and Queiroz, 2010; Weitzenfeld, 2008).

Here, I will explore what emotions are, both in biological organisms as well as what it would mean for an artificial agent to have emotions and how emotions could be useful to an artificial agent (Sections 2.1-2.3). Because of the importance of communication to this study, I will discuss signaling in groups of biological organisms (Section 2.4) and how communication systems are implemented in groups of artificial agents (Section 2.5). I will then present the biologically-inspired predator-prey model I have developed, which is designed to examine the effectiveness of an emotion-like system in guiding behavior and the utility of affective signaling in groups. This model is implemented first in a single prey agent (Section 3) and subsequently in a group of prey agents (Section 4). A general discussion and conclusions follow (Sections 5 & 6).

2. Theoretical background

2.1. What is emotion?

One obtrusive problem with the word “emotion” is the connotation of *feeling*. In accordance with Damasio (1999; 2010), here I take the stance that feelings are not emotions: feelings are *perceptions* of emotions – perceptions of what is happening within one's body and mind during the emoting. One must have emotions to have feelings, but being aware of one's emotions is not a requirement for having them in the first place. Jerome Kagan, in his book *What Is Emotion?* (2007), withholds the title “emotion” from any phenomena in organisms that do not have the symbolic appraisal of sensory states that he thinks is necessary to emotion, but this seems an unneeded distinction, falling too far toward equating emotions with feelings. Restricting one's definition of emotion to those organisms *aware* of their emotions is not particularly useful (except in the case of needing a perfect synonym for “feeling”), since there is no feasible way to determine which organisms *are* aware of their emotions.

As such, emotions are usually defined as a system that aids in behavior regulation and control (e.g., see Arkin, 2005; Bekoff, 2000; Damasio, 2010; and Tooby & Cosmides, 2008). Both Arkin (2005) and Bekoff (2000) point out that “emotion” is a nebulous term – no single theory has yet captured the full extent of the complexity of emotions. Some strict behaviorists reject any discussion of emotions and mental states altogether, since emotions consist at least in part of unobservable phenomena that cannot produce what they think is “proper” scientific data (Bekoff, 2000). More often, however, researchers align themselves approximately with the theories presented by Damasio (2010) and Tooby and Cosmides (2008).

Damasio's (2010) definition is broad, asserting that emotions are complex, largely automated patterns of response, used to regulate the organism's internal state in order to prepare specific reactions to specific environmental cues. They are a high-level life-regulation system. His definition encompasses value systems, reward and punishment devices, drives and motivations, and “proper” emotions such as sadness, fear, disgust, and joy. Arbib and Fellous (2004) and Smith (2004) agree that emotions are evoked in response to particular harms and benefits in the organism's environment, with the purpose of regulating internal states and affecting behavior, without explicitly mentioning whether emotion is a subset of or has a subsets including motivation, value systems, and the rest.

Tooby and Cosmides (2008) present a concurring argument to Damasio's: emotions are “programs” that coordinate action within one's body, directing pieces such as behaviors, reflexes, attention, and effort allocation. Most organisms encounter similar situations time and time again; it is beneficial for the organism to be able to recognize a recurring situation and to direct its body to prepare in specific ways in response. In addition, organisms must have a mechanism for choosing between mutually incompatible actions, such as sleeping and fleeing from a predator, and then coordinating the response of the organism's body for executing the chosen action. Tooby and Cosmides (2008) suggest that emotions are this mechanism, a bet about what internal actions will be most advantageous in the long-term, given the present uncertain circumstances.

Both Damasio (2010) and Tooby and Cosmides (2008) place emotion in a key role in the determination and coordination of behavior, but their definitions leave something to be desired. As Dawkins (2000) explains, many functional definitions of emotion are so general that they encompass *too much* behavior – including the behavior of plants operating without nervous systems. She describes in particular Oatley and Jenkins' (1998) appraisal theory of emotion in relation to a study by Kelly (1992) of the growth of a parasitic plant onto Hawthorn bushes. Oatley and Jenkins' (1998) theory proposes that emotions are composed of appraisal of the situation, action readiness, and physiological changes and the execution of action. In Kelly's (1992) study, pieces of the parasite were seen to grow away from Hawthorne bushes that had been starved of nutrients, and towards bushes that had been fed extra nutrients, thus displaying – over the course of several hours – appraisal of the food source, action readiness, and the execution of action in growing toward or away the bush. Although Damasio (2010) and Tooby and Cosmidies (2008) do take a little more care with their definitions, it could be argued that even a plant has some mechanism for choosing between incompatible actions – such as growing one direction versus another – and readying itself to act.

Arkin (2005) and Parisi and Petrosino (2010) take a different view of emotion: their definitions are still functional, but have finer-grained delineations, naming motivation and separating this phenomenon out from emotion. Motivations can be discussed without reference to emotion. In an organism, motivations are states in which the organism is willing to work to gain a reward or to avoid a punisher (Rolls, 2005). This excludes simple reflexes – motivations are states that may evoke specific behaviors to accomplish particular goals (Arkin, 2005). Emotions, on the other hand, are a specialized sub-mechanism of motivational behavior, states of the organism's body and brain that change the intensity of different motivations, modulating existing behavior rather than leading to the formulation of concrete, goal-achieving behavior.

Both motivation and emotion exist on a strategic level of behavior rather than a merely reactive level; they are often confused. Rolls (1999, 2005) has two well-known books on emotion that are, in fact, primarily about motivation, as Parisi and Petrosino (2010) point out. Rolls (2005) claims that emotions are states elicited by rewards and punishers – side-effects, not mechanisms for influencing motivation at all. In contrast, the definition I am taking here keeps emotion in a key if indirect role in behavior determination – emotions are adaptive, survival-promoting processes (Panksepp, 1998). They influence the appropriateness and effectiveness of the motivations driving behavior, allowing the organism to make motivational decisions more quickly (Parisi & Petrosino, 2010). They are important to determining what information in the environment to attend to given the current circumstances, and may lead to more flexible learning and reasoning (Evans, 2002).

2.2. Emotion in biological organisms

Most of the emotion research controversy comes in when we start talking about which organisms have emotions and whether artificial agents could ever have emotions. Bekoff (2000) suggests that we can make strong interferences about emotional states in many organisms based on features such as posture, facial expressions, and vocalizations, but this presupposes that only a single mechanism – some kind of emotion system – could give rise to those particular displays of behavior. Much human-robot interaction (HRI) research is centered around the very idea that one can achieve emotion-like behavior in an artificial agent without the agent having anything like what we would deem emotions in a biological organism. Having such behavior in the agent helps humans relate to the agent in natural ways, which is particularly important to, e.g., the entertainment robotics sector (Arkin et al., 2003). That said, indirect evidence about the affective states of any creature other than our own self is all we have, and can in many cases provide the kind of strong inferences to which Bekoff (2000) refers.

Dawkins (2000) proposes two methods of studying animal emotions, both based on the idea that we can and do know something about human emotions: 1) functionally, by looking at whether the role of emotions in non-human behavior is similar to that we see in humans, and 2) mechanistically, by examining whether the physiological mechanisms underlying emotions are similar in non-humans and in humans. Dawkins (2000) described several ways we get information about human emotions – verbal reports about perceived emotions; measurements of automatic changes such as heart rate, temperature, and galvanic skin response; and observations of behavior and expressions – and suggests that we can use the latter two for studies of emotion in non-human organisms. She noted that these measures are not, in general, correlated – e.g., sometimes an automatic response is measured, but no subjective experience is reported. Furthermore, there can be multiple causes for any particular observed reaction, such as an increased heart rate as a result of either fear or sexual arousal. This makes it difficult to look at emotion analogs in humans and non-humans: how do we know if what we are looking at is, in fact, an analog?

Dawkins' (2000) solution is unsatisfactory: she claims that certain anticipatory mechanisms may require emotions, and thus, organisms with those mechanisms by necessity have emotions. Specifically, she stipulates that emotions are necessary for an organism to know whether a behavior never performed before ought to be repeated, arguing that organisms capable of reinforcement learning must have a reward-punishment system, which she names emotion. When an organism can learn to perform an arbitrary response to some stimulus, that response cannot be hard-wired in. She notes that receptors can be linked directly to response mechanisms, so mere habituation to some repeated experience is not enough. Although this is a nice story about emotion's adaptiveness in novel situations, Dawkins (2000) fails to acknowledge or explain the utility of emotions in any other circumstance, and furthermore, fails to uncover any other useful ways of investigating animal emotions. The fact that she goes on to equate emotion to cognition does not help her case.

Most other researchers, through studies of neuroanatomy, neurophysiology, and behavior, recognize some level of emotion in many biological creatures – most often, the capacity for fear and pain, since these are most easily seen as adaptive and useful to an organism's daily survival. Panksepp's (1998) book *Affective Neuroscience* is a much-cited source. He accepts that it is a difficult subject – asking, as an exemplar question, whether courting finches in the Galápagos have emotional experiences, or whether their actions mere fixed-action patterns – and reminds us that such methods as fMRIs and behavioral brain research across species and levels now go a great distance toward understanding systems of emotion in both humans and non-humans. He states affirmatively that all mammals have emotions; other researchers, such as Chandroo, Duncan, and Moccia (2004), propose that birds and even fish have some emotion. Chandroo et al. (2004) analyzed whether fish can suffer, investigating functionally comparable regions of the fish brain in relation to tetrapod brains. They noted that the limbic system and some dopaminergic neural systems have been associated with emotion in tetrapods (these systems are consistently implicated – see the discussion in Ledoux and Phelps, 2008), and the telencephalon in fish is comparable – e.g., lesions or stimulation of the amygdalar region in the telencephalon produce aggression changes in fish similar to those in tetrapods, and lesions of the telencephalon in goldfish affected spatial and associative learning in the same way as hippocampal lesions in tetrapods. Some researchers, such as Preston and de Waal (2002) go farther, arguing that animals not only have emotions, but also empathy. They cite studies of rats and pigeons in which subjects showed distress when a shock was administered to an adjacent animal, claiming that this showed both rats and pigeons are emotionally connected to their conspecifics. This conclusion conflates the significance of the animals' reactions – the distress seen could merely be an automatic reaction to seeing a conspecific in distress, since it may be adaptive to one's survival to react fearfully when a nearby rat or pigeon is already acting fearfully. Regardless, the general consensus in the literature is that many animals have emotions, and that emotions are particularly adaptive to have.

2.3. Emotion in artificial agents

Emotion modeling, as a general category, can cover a wide range of models, including some that generate emotion-like expressions or gestures in agents, use the effects of emotions for decision-making and action planning, or use information about a user's emotions, e.g., for games and tutoring programs (Hudlicka, 2008). Here, I am specifically interested in the use of emotion-like systems in the second category – the generation of more flexible and adaptive behavior in autonomous robots. As was discussed in Section 2.1, I am considering emotion to be an adaptive system that, given specific environmental situations, helps coordinate appropriate behaviors by influencing motivation – I am setting aside the matter of feeling for another debate entirely. Since we do not build bodies for robots from living cells, obviously, no emotion model or system built will be mechanistically equivalent to that in a biological organism, but as Fellous (2004) points out, having a strong mechanistic understanding of biological emotions allows us to create functional equivalents, in the same way that having a mechanistic understanding of the human heart allowed artificial hearts to be built. MacLennon asks – if artificial emotion systems fulfill the same functions in robots that emotions do in biological organisms – whether we care what they're called – emotions, artificial emotions, programs?

If an emotion system will be meaningful to a robot, it must be grounded in a relevant internal value system, constituting a “second order” control mechanism that works in conjunction with a motivation system and action selection, monitoring the external and internal environments for significant events (Cañamero, 2005). Such systems are used in a few primary domains in robotics; Fellous (2004) outlines three: 1) communication between individuals, 2) resource mobilization and conservation, and the prioritization of behaviors, and 3) decoupling stimulus and response to allow for context-dependent computations. Numerous other researchers have similar lists (see, e.g., Lee-Johnson & Carnegie, 2007; and Miachaud & Audet, 2001); Rumbell, Barnden, Denham, & Wennekers (2010) add memory control, learning, and self models to the list.

Rumbell et al. (2010) suggest that, when implementing emotions in robots, two primary theories of emotion mechanisms are used. First, the theory that emotions can be classified into a small set of discrete, “basic” emotions, based on identifiable functional roles. Second, that emotions are dimensional, composed of the interacting output of some number of fundamental components. Different emotions are simply particular points in a multi-dimensional emotion space; the problem with this is that it is easy to miss interesting features of emotion.

Miachaud and Audet (2001) use a dimensional model and place emphasis on emotion's role in managing social behavior. Their study placed a pair of robots in an arena with the goal of performing a long-term foraging task. The robots needed to periodically recharge at a charging station. The robots have a two-dimensional bipolar model of four “emotions” as a way of monitoring the accomplishment of their goals over time – being low on energy would, e.g., increase “fear” because the robot was operating in dangerous conditions. The robots could communicate their general affective state to each other; this information was used to coordinate the accomplishment of the task and the long-term survival of the pair via cooperative use of the recharging station.

Lee-Johnson and Carnegie (2007) implemented a model similar to that of Miachaud and Audet's (2001), but instead of a dimensional model, used five basic “emotions” that were linked to specific kinds of events. E.g., “fear” was associated with damage to the robot, and “happiness” was associated with robot achieving its goal. The robot's task was to navigate from location to location in an unknown environment, using subtasks such as learning about its environment and obstacle avoidance, and a hierarchical reactive/deliberation for action selection. Lee-Johnson and Carnegie (2007) evaluated the robot's performance with and without the artificial emotion system, concluding that the inclusion of the system improved the robot's adaptive performance.

Some researchers have moved away from explicitly basic or dimensional emotion models, using emotion systems for learning and memory. Velásquez (1998) built a model with a drive system, representing actions the agent might be motivated to take, and an emotional system, implemented as an associative network that influenced action selection. The network used a modified Hebbian rule to form “emotional” memories related to the actions taken and their results. Yu and Xu (2004) also implemented an associative network, using a fuzzy logic model of emotion and reinforcement learning to get “emotional” memories of actions and relevant situations. Hasson and Gaussier (2010) built an amygdala-inspired neural network that learned the conditioning of a sensor using internal “pain” and “pleasure” states; the network was then used for navigation.

Parisi and Petrosino (2010) have emphasized that only robots that have many competing motivations can ever be said to *have* emotions, because they purport that emotion is a submechanism that exists to help make more appropriate motivational decisions – this is congruous with the definition of emotion I am using here. Parisi and Petrosino (2010) discussed a series of simulated robots that each were given various competing motivations, such as consuming resources to maintain energy levels, avoiding a predator, or resting to heal physical damage. Information from sensory devices and about the robot's internal state were input to a neural network, which directly controlled the robot's motors. They added an emotion circuit to the network; the emotion units received input from relevant sensors (e.g., about the predator's location if the robot was in a scenario with a predator to avoid) and fed into either the internal units of the network or directly to the motor units on the output layer. In both cases, adding the emotion circuit led to better motivational decisions and higher fitness.

Numerous researchers have begun to use emotion systems in groups of robots to aid in action selection and coordination within the entire group. Murphy et al. (2002), e.g., put an emotion system in a pair of robots that performed the asynchronous sequential task of serving hor d'oeuvres, with one serving the role of the waiter (serving food items) and one serving the role of refiller (bringing out trays of refills). The emotion system aided in action selection by modifying behavior via principles similar to homeostatic control, so if a robot was not making sufficient progress on a task, it might attempt to perform faster. It also aided in avoiding deadlock situations by escalating the robot's response, perhaps changing the robot's behavior altogether. Murphy et al. (2002) emphasized that their model only captured part of the many functions of biological emotions. The robots did not need to interpret or understand one another's “emotions” for the emotion systems to be beneficial to the accomplishment of their interdependent tasks.

Another multi-robot system was presented by Nozawa, Sato, and Ide (2005), which used an emotion-inspired model of affinity/antipathy to build behavior rules for robot interaction. A robot with no obstacles in its path entered a “pleasure” state; if there were obstacles, it would enter a “non-pleasure” state. Four simulated robots were placed in an empty arena; they would approach one another when in “pleasure” states and avoid one another when in “non-pleasure” states. When an attractor point – the robots' “destination” – was introduced to the arena, the approach-avoid behavior between robots led to general obstacle avoidance of walls when approaching the destination point. In a more complex environment, the “pleasure” fields resulting from the robots' movement could be used to draw environment maps. Nozawa et al. (2005) suggested that the emergent behavior seen was more flexible than the behavior that would have been generated by a central control mechanism. When implementing the same model in real robots in a later study, Kusano, Nozawa, and Ide (2008) suggested that it led to the emergence of burden sharing. The use of an emotion system can be beneficial both to agents acting individually and to agents that act as members of a group – even when the agents are not strictly aware of other group members' “emotions.”

2.4. Signaling

When talking about groups of agents, a key question is what kind of communication exists between those agents. Animal signals are well-studied – in biological organisms, signals are used by senders to affect and exploit the behavior of organisms receiving the signals in ways beneficial to the sender (Endler, 2000; Krebs & Dawkins, 1984). For both senders and receivers, signals serve to reduce uncertainty about other organisms' future behavior (Krebs & Dawkins, 1984). They provide reliable pieces of information about relations in the organism's environment that can be used to the organism's advantage (Bradbury & Vehrencamp, 1998; Endler, 2000). Signals are used to determine aspects of the sender's identity, such as social group, sex, or age, and to determine the sender's location; they are used in mate attraction, courtship, parent-offspring interactions, in aggressive encounters to signal intentions and commitment, and also to share information about conditions external to either the sender or the receiver, such as alarm signals to indicate the presence of a predator or other thread, or food signals to indicate the presence of a shareable food source (Bradbury & Vehrencamp, 1998). Rendall et al. (2009) are careful to point out that signals are not just about an exchange of information, however – it is the *influence* of a signal on another organism that matters.

Seyfarth and Cheney (2003) agree: senders rarely signal except in the presence of other organisms. The meaning and function of a signal may be very different from the point of view of a sender versus a receiver; they have potentially divergent interests and asymmetrical mechanisms for signal production and reception.

Signals consist of two parts: *content*, the "what" that is being communicated, the piece of information being used by the receiver to make a decision about what action to make next, and some detectable and processable *means* that carry the content, such as light, sound, or chemicals (Endler, 2000). The best signals are tuned to the receiver's senses and are inexpensive to produce so they can be given many times – using, e.g., resonant structures or shorter signals with less structure. In addition, the environment can modify signals by adding noise or masking. E.g., the amplitude of an auditory signal is particularly sensitive to noise, since the amplitude of the received signal is the sum of the amplitudes of the signal and the noise, and a high noise level adds a high amplitude to the total sum. The best signals are in environments with the least amount of distortion and interference.

There is a careful balance in adaptive signaling between efficiency in communicating the message to the intended receiver and inefficiency at communicating to unintended receivers (Endler, 2000). Signals best avoid eavesdroppers when they match the background environment under the receiver's conditions, are only perceivable at close range, are non-repetitive, or have properties that allow certain environmental influences. E.g., guppies have strong red and orange courtship coloration in areas of the sea where they have a red-insensitive prawn predator, and blue coloration where one of their predators is the blue-insensitive cichlid; they also have a short courtship distance (Bradbury & Vehrencamp, 1998; Endler, 2000). Unfortunately, making a signal harder for eavesdroppers to detect also makes it harder for receivers to detect. Senders must balance the the conspicuousness of the signal, as well as the time and energy required for production, and physical constraints – e.g., a small body size constrains the intensity and transmission distance of visual signals (Bradbury & Vehrencamp, 1998).

The signal must also balance efficiency and content on the receiving end: a receiver's costs are reduced signals are reduced in complexity, or when complex signals are sent more slowly (Endler, 2000). The receiver must also balance the time to assess the signal, potential vulnerability if a close distance is required to perceive the signal, and numerous physical constraints – e.g., body size constrains the effectiveness of determining the directionality of auditory signals (Bradbury & Vehrencamp, 1998). For all organisms involved, it is a matter of trade-offs.

As noted previously, an integral part of signaling is emotion. Emotions are crucial to communication and social coordination (Arbib & Fellous, 2004), acting to promote particular signal-producing behavior. E.g., Ciani (2000) studied how rage could signal to an opponent an affective state of distress, a warning that the situation could escalate into a physical conflict, influencing other organism's future actions. Rage, Ciani (2000) argued, is useful both in communicating the organism's state to others, and in preparing an organism for antagonistic actions. In dogs, e.g., curling the lips to expose teeth may aid in avoiding biting the lips when attacking; the accelerated heartbeat, similarly, may be useful in an oxygen-demanding attack action. Rage can give way to aggression, if the organism perceives that the opponent can be defeated, or fear, in preparation for submission or flight.

An emotion common to most organisms is fear, as discussed in Section 2.2. Fear reactions come into play when a threat to the organism is present – most prominently, the threat of being eaten by a predator. In many animal populations, this threat triggers an alarm signal. Alarm signal production tends to be reflexive, an affective reaction elicited by the detection of a predator or by a sound, such as the call of a conspecific (Seyfarth & Cheney, 2003). Bradbury and Vehrencamp (1998) distinguish three types of alarm signals, based on the kind of response elicited in conspecifics: 1) flee alarms, with a shorter range and generally harder to localize, are given to cause all receivers to avoid an attacking predator; 2) assembly alarms, with a longer range, are used to induce conspecifics to converge with the sender; and 3) alerting signals, often with a longer range as well, alert receivers to the presence of danger without inducing immediate action, and can also alert predators that they have been spotted. Sherman (1977) add that alarm signals may also function to divert a predator's attention to another prey and reduce the likelihood of later attacks by the same predator if that predator often returns to the same area. Some animals use more than one type of alarm signal – the type is generally dependent on the particular predator and the urgency of the situation (Seyfarth & Cheney, 2003).

Alarm signals are only found in relatively social species – they transmit information about presence of predator or rival, and their production depends heavily on the presence of potential recipients (Bradbury & Vehrencamp, 1998; Seyfarth & Cheney, 2003). The perception of alarm calls may involve very different mechanisms than sending because of the subtle information that may be available regarding different predators or the urgency of the situation (Seyfarth & Cheney, 2003). Seyfarth and Cheney also note that most animals cannot distinguish the mental states of other animals – calls are not made intentionally to inform receivers and and receivers do not recognize that a signal reflects the sender's knowledge of a situation.

Sherman (1977) emphasized that producing alarm signals can be costly to an organism – it may reveal the signaler's location or the signaler's group's location to a predator. That said, alarm signals exist in many animal populations – vervet monkeys, ground squirrels, ring-tailed lemurs, and domestic roosters, to name a few – and so the benefits must outweigh the costs. Sherman (1977), after observing ground squirrels for thousands of hours and tabulating the circumstances of their alarm calls, argued that alerting kin to the presence of a threat was the primary function of ground squirrels' alarm calls. Females gave calls more frequently when relatives might hear; when they were surrounded by unrelated females, they refrained from calling. Females also called more frequently than males, and male squirrels with several mates with offspring called more frequently than males without. Bradbury and Vehrencamp (1998) also propose that most alarm signals are given when kin are among the receivers. Having kin survive may help in gene transmission to future generations.

2.5. Alarm signals in artificial agents

Numerous researchers have begun to develop biologically-inspired strategies of communication in groups of artificial agents. Loula et al. (2010) investigated emergent symbolic communication systems in groups of artificial agents, drawing heavily on the vervet monkey alarm call research of Seyfarth and Cheney. They built a simulated world populated by prey agents and predator agents. The behavior of these agents was determined by a hierarchical control architecture in which the action with the highest motivation was selected. Motivation was influenced by an emotion-like system of “drives” – such as hunger and tiredness – and actions included wandering, resting, and either chasing or fleeing. The prey also had drives for fear, boredom, solitude, and curiosity, as well as actions for vocalizing, learning, following, and scanning, which in conjunction led to the prey to stay together so that communicative interactions would occur more frequently, and learn to associate vocalizations with predatory events. Four prey were put in an environment with three predators; the prey developed a self-organized system of alarm calls that were functionally referent to predatory events. The fact that a system of alarm signals independently developed in the prey suggests that signaling about threats in the environment is beneficial to the prey's survival.

Other artificial predator-prey models exist, but their focus has tended not to be on alarm signaling and emergent prey group behavior. Jim and Giles (2000) developed a model of four predator agents that communicated with one another to capture a single prey. In their simulation, each predator could “speak” a string of symbols, and could read all strings of symbols “spoken” by other predators before determining its next move. The predators' “language” was evolved over time using a genetic algorithm; predator performance improved when using the “language,” and more so when the number of possible messages that could be “spoken” was increased. Jim and Giles (2000) suggest that communication among agents is especially helpful in partially observable environments – that is, when agents have access to potentially different pieces of information about parts of the entire situation.

Weitzenfeld (2008) also investigated predator behavior, developing a prey catching and predator avoidance architecture based on a frog and praying mantis stimulus-response behaviors, both in simulation and with real robots. The behavior of each agent was controlled by a multi-level neural network and schema architecture. He extended this model using strategies from wolf pack group behaviors, creating a scenario in which several “beta” predators maintained a formation behind the “alpha” predator to pursue a prey.

Nolfi and Floreano's (1998) study had a different focus: Their goal was to examine the co-evolution of predator and prey agents over time, modeling the “arms race” of competing adaptations and levels of complexity. They placed one simulated predator and one simulated prey in an otherwise empty environment, each controlled by an artificial neural network, and used a genetic algorithm to evolve aspects of the network and strategies for survival. Floreano and Keller (2010) discuss in detail the benefits of taking an evolutionary approach to robotics, suggesting that evolving strategies of navigation, cooperation, and survival can lead to unique an powerful solutions, particularly when there may be no single “right” combination of traits.

These predator-prey models, in conjunction with the research described in previous sections on emotion systems in groups of robots, the functions of emotion in biological organisms, and the uses of alarm signaling in biological organisms, have inspired the experiment described in Section 3. I have developed a novel action selection architecture – including an emotion system and capabilities for alarm signaling – for prey robots with a single predator.

3. Study One

3.1. Overview

The purpose of this study was twofold: First, because a crucial assumption is that an emotion system *is* beneficial, to demonstrate that a prey agent with an emotion system is *more fit* than a prey agent without an emotion system – i.e., that having emotions is beneficial and increases the prey's lifespan. Second, to establish a baseline with the interactions of a single prey agent with a predator agent for later studies with multiple prey agents. A genetic algorithm was used to determine good values for the variables in the emotion and motivation equations for a prey agent with emotions. This agent's behavior and fitness was compared to an emotionless prey agent.

3.2. Methods

3.2.1. Initial simulations

The simulated robots were implemented in C# using Microsoft Robotics Studio and Microsoft Visual Studio .NET. A genetic algorithm, also implemented in C#, was run with just one prey agent and one predator agent in order to evolve appropriate thresholds for variables in the motivation and action-selection equations. Twenty trials were run per generation for sixty generations, the results of which were used to inform a second run of the genetic algorithm, with twenty trials per generation for twenty-five generations. Each trial consisted of 720 timesteps of 250ms each for an overall trial time of three minutes; at each timestep, the prey and the predator calculated which action they would each perform next, as detailed below. Each predator-prey pair ran in its own simulation. Details of the genetic algorithm are provided below.

The four best prey agents from the second run of the genetic algorithm were determined based on fitness. These agents were compared to 1) prey agents acting in a predator-less environment, and 2) an emotionless prey agent.

3.2.2. Comparison with predator-less prey agent

As a control and to establish a baseline lifespan, each of the best prey agents (as determined above) was tested without a predator in the environment. The hypothesis was that a prey without a predator would, on average, survive longer. Eighty trials, each consisting of 720 timesteps of 250ms each, were run for each prey. In the first forty, both the predator and the prey were placed in an obstacle-free environment with their starting locations randomly determined. In the second forty, only the prey was placed in the environment. The predator could seek out the prey and attempt to attack; the prey could “eat” to maintain energy levels and move to avoid the prey, as is described in detail below.

3.2.3. Comparison with emotionless prey agent

To determine whether the addition of an emotion/motivation system was beneficial to the prey agent, forty trials of 720 timesteps of 250ms each were run for each of the best prey, and for an emotionless prey agent. Again, both the predator and prey were placed in an obstacle-free environment with their starting locations randomly determined. The predator could seek out the prey and attempt to attack; the prey could “eat” to maintain energy levels and move to avoid the prey, as is described in detail below. The only difference between the emotion and the emotionless prey was that the emotionless prey did not have “fear” or “vigilance” modifying its action-selection function.

3.2.4. Genetic Algorithm

The genetic algorithm served the same function as evolution does for biological organisms: to select the individuals whose reactions to conditions and situations were relevant and beneficial (Floreano & Keller, 2010; Eiban & Smith, 2003). Individuals should react to recurring situations in a consistent, survival-oriented way, even if their behavior is often not so explicit. Each prey agent in the initial population was given a genome with randomly generated values in a normal distribution with mean zero and a standard deviation of one, within a specified range for each gene (see Figure 3.1).

Each successive generation of prey agents was produced from the previous generation. After all trials in a generation had finished, the overall fitness of each prey agent was calculated: prey that survived the most number of timesteps were considered most fit. If two or more prey survived an equal number of timesteps, they were secondarily ranked based on their motivation, with prey that had lower hunger and lower fear during their lifetime ranked higher – this is a "quality of life" measure, in the vein of O'Bryne et al.'s (2009) fitness functions. A set of "parents," used to produce eighty percent of the next generation, were selected based on these fitness values, using rank-based selection mapped to selection probability and stochastic universal sampling (for a detailed account of these algorithms, see, e.g., Eiban & Smith, 2003).

Pairs of parents were randomly selected from the parents pool to undergo recombination and mutation, thus allowing for the transfer of information between successful individuals. Whole arithmetic recombination occurred with eighty percent probability to create two offspring; otherwise, each parent's genome was copied asexually to produce the two offspring. Mutation of the offspring's genome occurred with probability one, at a rate in the range [0.1, 0.35]. This mutation rate was also evolved, allowing for some self-adaptation of the genetic algorithm. It is often the case that different values of some evolutionary parameters are optimal at different points during the genetic algorithm's lifetime (Eiben & Smith, 2003).

The remaining twenty percent of the new generation were selected based on fitness from the previous generation. These "survivors" were copied whole and not altered or mutated.

The initial population of the second run had genomes randomly generated in a normal distribution with means informed by the best agents of the first run, and a standard deviation of one (see Figure 3.1). The first run was also used to inform the structure of the algorithm itself (e.g., the maximum rate of mutation). Results from the second genetic algorithm are discussed further in Section 3.3.

3.2.5. Environment

The environment was a flat, textured plane 100 meters by 100 meters, bounded by walls that represented “rivers” and “mountains.” It contained one safe zone, denoted by a yellow patch on the plane, and twenty randomly distributed food patches, denoted by green patches. Because the transmission of visual signals can be improved by choosing different micro-environments with different water absorbency or different light levels in order to maximize the radiance of the color patterns or maximize the difference between the color pattern and the background environment (Endler, 2000), the background environment was gray, and the lighting was relatively constant, thus simplifying the scenario (see Figure 3.5).

3.2.6. The agents

The predator and the prey each consisted of a robot with differential drive, bump sensors, infrared sensors, and a webcam. As a visual signal of which agents are of what “species,” both the predator and prey were color-coded: the predator was red; the prey was blue (as can be seen in Figure 3.5). This allowed for each agent to use color blob tracking to detect and track the other.

3.2.7. Predator behavior

The predator's behavior was controlled by a simple hierarchy: If a bump sensor was hit, avoid the obstacle by backing up and turning a random direction. Otherwise, search for the prey agent using both color blob tracking with the webcam and monitoring of any alarm messages the prey emitted, which would provide data on the prey's location (see Figure 3.6 for a schematic). The predator drove at a moderately slow speed during both these behaviors. If the prey was detected within a threshold distance, the predator would attack the prey at a faster speed. The predator's top speed was slower than the prey's top speed, but the predator accelerated faster. This is consistent with biological data, as well as previous artificial predator-prey scenarios, such as Nolfi and Floreano's (1998), in which they gave their prey a top speed twice that of the predator.

3.2.8. Prey behavior – with emotions

The prey's behavior is based on a winner-take-all action selection function, with an obstacle avoidance behavior subsuming all other behaviors (if a bump sensor was hit, the prey would back up and turn away from the obstacle). The prey was provided with a “safe zone,” based off the idea that many animals (e.g., rabbits, voles) have a burrow or safe hole to hide in. The possible actions are *search-for-predator, forage, avoid-predator,* and *alarm*; these actions were as follows (see Figure 3.7):

Search-for-predator: Spin in place to scan the environment, searching for predator-colored blobs. If the predator is detected (i.e., a red blob above the threshold size), stay fixated on the predator and track the blob. Otherwise, continue scanning.

Forage: Spin in place to scan the environment, searching for food-colored blobs. If food is detected after one revolution (i.e., a green blob above threshold size), move toward the food. Otherwise, randomly move in a direction and scan again. When within a threshold distance of the food, stay there and “eat.”

Avoid-predator: If the predator is visible (i.e., a red blob above the threshold size has been detected), turn away so that the predator is not visible and move forward in that direction. Search for safe zones (i.e., yellow-orange blobs above threshold size) and approach these.

Alarm: Send an alarm message to all agents. Agents must be within 2.0 meters of the sender to be able to “hear” a message. In accordance with Seyfarth and Cheney's (2003) observations of biological organisms, alarm calls have the potential to be elicited by detecting a predator, by fear level alone, or by the call of a conspecific.

Motivation for each action was calculated based on “fear” and “hunger” levels. Fear was a measure of endangerment to the prey's well-being and was a function of nearness to the predator and alarm calls from other prey. Hunger was a measure of the prey's energy deficit and was a function current energy levels, which in turn were incremented and decremented based on movement and amount “eaten.” Although the words “fear,” “hunger,” etc. are laden with emotional connotations, recall that I have defined emotion such that it does not require the *feeling* of the emotion – emotion is just a complex pattern of response that prepares the agent to act in particular ways (see Section 2.1). These words best denote the kind of motivations interacting within the prey.

In all equations below, the sensitivity factors and exponents were evolved in the genetic algorithm.

Fear is an function of distance to predator and alarms, where k_f is a constant that was also evolved, D is the distance to the predator, A_N is the number of alarms detected, A_I is the intensity of the alarms, and d and a are the evolved exponents.

$$\text{FEAR} = k_f + D^{-d} + (A_n * A_i)^a$$

Hunger is a function of energy levels, where k_h is a constant that was also evolved, E is the current energy level, and n is an evolved exponent.

$$\text{HUNGER} = k_h + E^{-n}$$

Vigilance is a leaky integrator function, where k is a constant that was evolved, L is the rate of the leak, also evolved, t is time, and C is the input to the function.

$$\text{VIGILANCE} = k_v e^{(-Lt)} + C/L$$

Search-for-pred is a linear function of fear level and vigilance, where F is the current fear level, V is the vigilance factor, and s_{sf} and s_{sv} are sensitivity factors.

SEARCH = $s_{sf}F + s_{sv}V$

Forage is a linear function of hunger, where *H* is the current hunger level and s_{fo} is a sensitivity factor.

FORAGE = $s_{fo}H$

Avoid-predator is a linear function of fear, where *F* is the current fear level, *V* is the vigilance factor, s_{av} and s_{av} are sensitivity factors, and k is a constant.

AVOID = $s_{av}F + s_{av}V + k_{av}$

Alarm is a linear function of fear, where *F* is the current fear level, s_{al} is a sensitivity factor, and k is a constant.

ALARM = $s_{al}F + k_{al}$

3.2.9. Prey behavior – without emotions

The motivation for selecting any particular action was not modified by *fear* or *vigilance* values in the emotionless agent. As such, the agent had a reduced set of actions, containing only *forage* and *avoid-predator* because *alarm* served as an expression of affective state, and *search-for-predator* was in response to vigilance and fear alone. If the predator came within a threshold distance (1.0m) of the prey, the prey selected *avoid-predator.* Otherwise, the prey selected *forage* and sought out food.

3.2.10. Food

Twenty food patches were randomly distributed throughout the plane, color-coded green and each with a randomly picked starting number of energy units between 60 and 200. Food patches replenish at a constant rate of 4 energy units per timestep. When eaten by the prey, a food patch's energy potential decreases at a rate of 10 energy units per timestep. The prey gets 10 energy units per timestep from eating.

The prey's linear energy decay and restoration rates were based on graphs of energy expenditure in lions and gazelles (Elliott, Cowan, & Holling, 1977), which showed that different tissues are more expensive to maintain and some gaits require more energy expenditure. Each prey started with an energy of 800, which meant that unless it successfully foraged, it would die before the end of the trial.

3.2.11. Data Analysis

The lifespan in timesteps of each prey agent in each simulation was recorded, as well as the agent's motivations, energy, and action selected each timestep. Two-tailed t-tests were run in Excel to compare the fitness of prey with predators to prey without a predator, and to compare the fitness of prey with emotion to prey without emotion.

3.3. Results

Results from the second genetic algorithm are shown in Figures 3.2 through 3.4. Figure 3.2 displays the average lifespan of prey agents in each generation, with a trend of increasing lifespan over time and a noticeable inflection point after ten generations. The evolution of each gene over time is shown in Figures 3.3a-3.3c; the same inflection point can be seen. The genes continued changing across all generations with the exception of the mutation rate, which maxed out, and the exponent modifying *Energy* in the *Hunger* equation, which leveled out as well. The genomes of best prey, calculated based on fitness and used for the emotion system vs. no emotion system comparisons, are shown in Figures 3.4a-3.4d. These demonstrate that despite average trends in genome change, the genomes of individual prey varied significantly and the genomes of successful prey did not necessarily follow the same trends.

The lifespans of the four best prey agents (Prey1, Prey2, Prey3, Prey4) with and without a predator are plotted in Figure 3.8. Two-tailed t-tests revealed that Prey3 and Prey4 had significantly longer lifespans in timesteps when no predator was present (Prey3 M = 510.225, SD = 176.3065; Prey4 M = 522.12, SD = 195.683) than when there was a predator (Prey3 M = 160.475, SD = 147.312; Prey4 M = 395.55, SD = 200.935), Prey3 t(78) = 9.6279, $p < 0.0001$; Prey4 t(78) = 2.8541, p = 0.0055. Prey1 had a longer lifespan when no predator was present (M = 542.175, SD = 209.657) than when it was (M = 457.875, SD = 218.5438), but this difference was not quite significant, t(78) = 1.6706, p = 0.0988. Prey2 was anomalous: its lifespan was longer on average when the predator was present (M = 630.05, SD = 173.378) than when it was not (M = 517.8, SD = 208.704), t(78) = 2.6165, p = 0.0107.

Figure 3.8 also shows the lifespans of the emotion prey and the emotionless prey. A two-tailed t-test revealed that the lifespans of Prey1 (M = 457.875, SD = 218.544) did not differ significantly from the prey without emotions (M = 526.025, SD = 240.617), t(78) = 1.3260, p = 0.1887. Prey2's lifespan (M = 630.05, SD = 173.358) was significantly longer than that of the emotionless prey, t(78) = 2.2185, p = 0.0294. Prey3's lifespan (M = 160.475, SD = 147.312) was significantly shorter than that of the emotionless prey, t(78) = 8.1946, $p < 0.0001$, as was Prey4's (M = 395.55, SD = 200.935), t(78) = 2.6323, p = 0.0102. The average lifespan of the emotionless prey did not differ significantly in the presence of the predator versus not (M = 552.20, SD = 208.48), t(78) = 0.5200, p = 0.6046, though the trend was in the expected direction.

3.4. Discussion

In this study, the primary goal was to determine whether having an emotion system was beneficial to a prey agent, extending its lifespan beyond what could be attained via a less adaptive action selection system. The results did not conclusively show that this was the case: in only one of four cases was the average lifespan of a prey with an emotion system significantly greater than that of a prey agent without an emotion system. Two of the others' average lifespans did not differ significantly from the emotionless prey, and the fourth had a lifespan significantly shorter. The control, in which prey agents with and without a predator were compared, was marginally more successful at verifying that the presence of a predator significantly decreases a prey's lifespan, regardless of whether it had an emotion system or not.

The emotionless prey had predictable behavior: it would forage for food until the predator came within a threshold distance (1m), and then it would avoid. It was frequently successful at escaping the predator and then continuing to forage, surviving entire trials by switching between these two actions. When it died, it was more often a failure to locate a food patch than a failure to escape the predator. During the predator-less trials, the emotionless prey only foraged, and any deaths were solely the result of "starvation." The fact that many prey failed to locate food patches suggests that their *forage* algorithm could be improved. Similarly, because the predator in many cases failed to ever find the prey, the predator's *seek-prey* algorithm could also be improved. It would likely prove interesting to instantiate multiple different search strategies in the agents and to compare which strategies work better in relation to the agent's body, motivations and survival strategy, and environment.

Prey1's strategy was most similar to the emotionless prey's, and thus it makes sense that its average lifespan did not differ significantly from the emotionless prey's. It foraged near-continuously, never selecting *search-for-predator,* and only selecting *avoid* when the predator came very close. It tended to be unsuccessful at avoiding, switching to the action when the predator was already too close (5-10 timesteps away), and dying as a result. When there was no predator, Prey1 had a longer lifespan on average, but not significantly so. This makes sense in light of Prey1's patterns of behavior: without the predator, Prey1 foraged and occasionally "starved" to death. With the predator, it occasionally got "eaten," thus slightly increasing its average number of deaths.

Like Prey1, Prey2 only selected *avoid* when the predator came very close, and tended to be unsuccessful at avoiding, though there were a few isolated instances of successful escapes. Its primary strategy was to *search* more than *forage* when there was a predator; with no predator, Prey2 foraged near-continuously and lived for a shorter span on average. Searching uses slightly less energy than foraging on average because it involves less movement, so switching between searching and foraging instead of just foraging could lead to the conservation of more energy overall. In addition, *search* involves spinning in place to scan the environment for the predator, which could, in addition, lead to better detection of nearby food patches and thus allow the prey to forage more successfully.

Prey3 also selected *search* more than *forage*, and its motivation to avoid was better tuned, allowing it to more frequently escape the predator. Indeed, in several instances, the predator chased Prey3 continuously for more than three hundred timesteps. Prey3 did not live as long as the emotionless prey on average, which was likely because it failed to start foraging soon enough to locate and approach a food patch. When there was no predator, Prey3 was much more likely to be able to forage successfully.

Similar to Prey3, Prey4 did not live as long as the emotionless prey on average. This was the result of both failing to start foraging soon enough and avoiding soon enough. When there was no predator, it foraged successfully more frequently.

Overall, more variation in average lifespan was seen when the predator was in the environment than when it was not, which was as expected. Without a predator, the task was very simple and there was just one optimum survival strategy: find food patches. With a predator, the prey evolved with the genetic algorithm developed very different strategies for survival compared both to each other and to the emotionless prey. There was no clear "best" strategy; having an emotion system was not demonstrably better than not. This could be in part due to the continued simplicity of the task and the environment. In other models of artificial agents with emotions, the tasks required of the agents have tended to be more complex: e.g., long-term foraging with a pair of robots (Miachaud & Audet, 2001), coordination of the serving of hor d'oeuvres (Murphy et al., 2002), navigation in an unknown environment (Hassaon & Gaussier, 2010; Lee-Johnson & Carnegie, 2007). Both O'Bryne et al. (2009) and Parisi and Petrosino (2010), whose simulated robots performed tasks fairly similar to this one, included more competing motivations: not only avoiding a predator and foraging to obtain energy, but also resting to heal physical damage, seeking "water" in addition to "food," or "mating" with other robots to generate "offspring." Indeed, simple biological organisms such as bacteria (Berg, 2000) and ants (Gordon, 2000), get along just fine with a simpler, emotionless architectures. Only more complex creatures that must perform comparably more complex tasks to survive have emotions – mammals, birds, fish (Chandroo et al., 2004; Panksepp, 1998; and Preston and de Waal, 2002).

Study Two investigates whether emotions are beneficial in more complex scenarios, specifically, when a prey agent acts as a member of a group.

4. Study Two

4.1. Overview

The goal of this study was to examine what it is that makes emotions beneficial to a prey agent. A genetic algorithm was used to evolve a group of prey with emotions to determine what kind of action selection and behavior would be beneficial in a group – specifically, whether alarm signaling would be utilized. This group of evolved prey was compared to a group of emotionless prey.

4.2. Methods

4.2.1. Simulations

The simulation environment, prey motivations and behavior, and predator behavior were the same as in Study One, with the following addendum: Instead of placing only a single prey in the simulation environment, a group of prey were used. The genomes of the four best emotion prey agents, as determined through the genetic algorithm in Study One, were each used to initialize five genomes in the first population for the multiple-prey genetic algorithm: two "clones" plus three copies randomly mutated a small amount. The best individuals from this population were selected and mutated to produce the subsequent generation, using the same genetic algorithm structure as was described in Section 3.2.4. Twenty prey were run per generation; each trial consisted of 720 timesteps of 250ms each for an overall trial time of three minutes. As detailed in Sections 3.2.7 and 3.2.8, at each timestep, the prey and the predator calculated which action they would each perform next, and each predator-prey group ran in its own simulation. This same genetic algorithm was run three times for fifteen, twenty-five, and twenty-two generations, respectively. Each time, evolution was halted because the genomes present had converged on a solution. Slightly different solutions were obtained each time. These results are discussed further in Section 4.3.

The five best individuals were determined based on longest average lifespan of the individual in respect to its population and with respect to the population as a whole. From the first run of the genetic algorithm, the best agent used the final solution found; for the second two runs, the best agents had genomes very slightly different from the final average solution. To account for this, two agents from each of these runs were selected. Figures 4.3 and 4.4 display these agents' genomes both in a table and visually. Each individual was cloned to create its own group of twenty agents; each of these groups was compared to a group of emotionless prey agents. In addition, a composite group formed by taking four copies of each of the best individuals was also compared.

4.2.2. Comparison with emotionless prey

To determine whether the emotion/motivation system was beneficial to prey acting in a group, twenty trials of 720 timesteps of 250ms each were run for each prey group – the five single genome groups, the one composite group, and the group of emotionless prey agents, using the emotionless prey described in Section 3.2.9. As before, the predator and prey agents were placed in an obstacle-free environment with their starting locations randomly determined. The predator could seek out the prey and attempt to attack; the prey could “eat” to maintain energy levels, move to observe the predator, send alarm signals, and move to avoid the prey, as described previously.

4.2.3. Data Analysis

The lifespan in timesteps of each prey agent in each simulation was recorded, as well as the agent's motivations, energy, and action selected each timestep. Two-tailed t-tests were run in Excel to compare the fitness of each emotion prey group to the emotionless prey group, and to compare the numbers of prey in each group that survived an entire trial.

4.3. Results

Results from the three genetic algorithms are displayed in Figures 4.1 through 4.5. Figure 4.1 shows the average lifespan over time for each run, while Figure 4.2 displays genetic change over time. It is not obvious from the change in average lifespan that the solutions improved, but from the inflection points and the genetic changes, it is clear that the solutions found changed significantly from the starting genomes from Study One. In each case, the genomes present in the population converged on a solution, with slightly different solutions found each time (Figures 4.3 and 4.4). As before, the genomes of successful prey varied widely and did not necessarily follow the same trends. The genomes represent complex interactions among many variables and it can be difficult to tell simply by looking at gene values what the resultant phenotype may be.

The average lifespans of the prey groups with emotion (PreyG1, PreyG2, PreyG3, PreyG5), the composite group (PreyGC) containing copies of each of these five, and the prey group without emotion are plotted in Figure 4.5. Two-tailed t-tests revealed that PreyG1, PreyG2, and PreyG4 all had significantly longer average lifespans in timesteps (PreyG1: M = 432.18, SD = 60.96; PreyG2: M = 350.59, SD = 41.08; PreyG4: M = 373.45, SD = 75.87) than the emotionless prey group (M = 319.26, SD = 36.15), PreyG1: $t(38) = 7.1262$, $p < 0.0001$; PreyG2: $t(38) = 2.5606$, $p = 0.0145$; PreyG4: $t(38) = 2.8834$, $p = 0.0064$. PreyG3 had significantly shorter mean lifespans in timesteps (M = 218.76, SD = 44.76) than the emotionless prey group, $t(38) = 7.8124$, $p < 0.0001$. PreyG5's average lifespan in timesteps (M = 316.32, SD = 51.81) did not significantly differ from of the emotionless prey group, $t(38) = 0.2081$, $p = 0.8362$; PreyGC also did not differ significantly (M = 337.54, SD = 46.08), $t(38) = 1.3953$, $p = 0.1710$.

In support of these results, in all but one case significantly more prey per group survived for the entire trial when they had an emotion system (PreyG1: M = 2.80, SD = 1.61; PreyG2: M = 3.10, SD = 1.25; PreyG4: M = 3.55, SD = 1.96; PreyG5: M = 2.15, SD = 1.73; PreyGC: M = 4.10, SD = 1.25) than when they did not (M = 1.9, SD = 0.85), PreyG1: $t(38) = 2.2118$, $p = 0.0331$; PreyG2: $t(38) = 3.5463$, $p = 0.0011$; PreyG4: $t(38) = 3.4555$, $p = 0.0014$; PreyG5: $t(38) = 2.9137$, $p = 0.0060$; PreyGC: $t(38) = 6.5087$, $p < 0.0001$. PreyG3 (M = 1.55, SD = 1.53) did not differ significantly from the emotionless prey group, t(38) = 0.9065, p = 0.3704. These means are displayed in Figure 4.6.

In light of the results of PreyG3, a two-tailed t-test was run to compare the lifespan of all prey from PreyG3 to all the prey with the same genome in PreyGC. The test revealed that the PreyG3 prey had shorter lifespans on average (M = 218.76, SD = 161.92). than the same genome in PreyGC (M = 260.75, SD = 190.77), but that this difference was not quite significant, t(458) = 1.8151, p = 0.0702.

4.4. Discussion

The purpose of this study was to determine whether having an emotion system would be beneficial when a prey agent acted as a member of group. The results suggest this is indeed the case: in all but one emotion prey group, more prey per trial survived on average than in the emotionless prey group. Additionally, three of the emotion prey groups survived significantly longer on average than the emotionless group. The fact that in a two cases, despite having a lower average lifespan than the emotionless prey, more prey survived per trial overall can be understood when one sees that in those cases, the tendency was for a few prey were picked off early, bringing the average down, with the remaining prey succeeding at living longer. These results stand in contrast to the previous study and suggest that in more complex scenarios – perhaps specifically when acting as a member of a group in a “social” environment – an emotion system becomes useful.

In terms of behavior, putting the prey agents in a group situation selected for prey that utilized the behavior of the other individuals in their environment – exactly as predicted. The emotionless prey group followed the same pattern of behavior as any single emotionless prey: they primarily foraged and avoided the predator when it came within a threshold distance. When they did not survive an entire trial, it was either because of 1) a failure to locate a food patch, 2) being caught by the predator, or 3) successfully avoided the predator, but being continuously chased, not allowing for a switch in behavior to foraging.

The prey with emotions displayed much more interesting behavior. The PreyG1 agents primarily searched until their energy dropped below approximately 50 (depending the predator's proximity). When the predator was near, they would alarm, but rarely avoided. Alarms, however, tended to be contagious – after one alarm signal was sent, nearby prey that received it were motivated to alarm as well, creating a chain reaction. In the PreyG2 agents, the contagion was even more pronounced. PreyG2 agents defaulted to foraging rather than searching – in fact, none ever utilized the *search* action – but when the predator approached, they would first alarm and then avoid. Nearby agents would similarly both alarm and avoid. Figure 4.7 shows the pattern of action selection across all individuals in one trial, as an example of this kind of behavior.

The one anomaly here was PreyG3, which performed worse than every other group, particularly when compared to the emotionless group. Although it is plausible that the initial performance of the PreyG3 genome during the genetic algorithm runs was a fluke, allowing to be selected as one of the best prey, this is unlikely. PreyG3 agents were highly responsive to the fear of other agents, swiftly switching to *avoid* as alarm signals spread throughout the group. As a result, they tended to alarm and avoid continuously until their energy fell below about 170 (depending on proximity to the predator). Because avoiding drains energy faster, these prey ran out of energy faster, and because they responded to alarms received so readily, they often failed to start foraging soon enough to find food patches. Figure 4.8 displays the pattern of action selection for one of the PreyG3 trials, as an example. When placed in a group with prey that did not alarm as readily (the composite PreyGC), prey with the PreyG3 genome performed better.

Alarm contagion was also seen in PreyG4 and PreyG5, but minimally – more often, receiving an alarm signal provoked an *avoid* response rather than an alarm. PreyG4 agents tended to search initially, switching to *avoid* or *alarm* in response to fear levels, and switching to *forage* after their energy levels fell below about 460. PreyG5 agents similarly searched initially, but switched to *forage* later, when their energy fell below about 290. In both groups, if there were no alarms and many prey were searching, when one (or more) prey alarmed, most if not all of the searching prey switched to *avoid*. Figure 4.9 shows an example.

The behavior of PreyGC was composite of all of the above. The presence of prey with the Prey3G genome led to the whole group being high-fear, with a great deal of *alarm* and *avoid* behavior. Alarms led to more avoiding than alarming overall. As a result, more prey successfully avoided the predator and survived the entire trial; very few prey were caught after the first fifty timesteps.

These results unequivocally show that when acting in a group, the prey agents evolved to take advantage of the behavior of other prey. Specifically, emotional contagion via alarm signals provoked an increase in fear responses that, on average, increased the likelihood that prey agents would survive longer.

5. General Discussion

The prey model developed here is both a starting point for designing architectures for intelligent robot behavior as well as a hypothesis about what drives the adaptability and flexibility of biological organisms – namely, emotion. The results reported support this hypothesis: artificial agents that have an emotion-like system behave more adaptively and tune their behavior to the particular environment in which they find themselves.

Study One suggested that in relatively simple scenarios, having an emotion system can lead to remarkably different solutions than what one might find in an emotionless architecture – some of which perform better than an emotionless agent, and others worse. In Study Two, I investigated whether an emotion system would be more beneficial when an agent was placed in a more complex group scenario. This study revealed that prey agents evolved from those in Study One began to utilize alarm signaling; emotional contagion of fear emerged as a result. More prey in these groups survived longer per trial than prey in the emotionless group, which confirmed my predictions that when an agent must select between mutually incompatible actions, such as foraging and fleeing from a predator, having an emotion system will aid in selecting advantageous actions, as was suggested by Tooby and Cosmides (2008). Additionally, it was only in the group scenario that alarm signaling became a useful action. This is in line with Bradbury and Vehrencamp (1998) and Seyfarth and Cheney (2003), who both noted that alarm signals are only found in social species, most likely because their production depends on the presence of potential recipients.

This is certainly not the first emotion model incorporated into a artificial agent, as was detailed in Section 2.3, but its contributions to this body of research are three-fold: First, it reaches farther than many of the models already developed. Instead of looking at predator-prey relations, emotion architectures, or communication between multiple agents in isolation, all three were merged to create a more realistic, challenging scenario. This demanded digging to the core of each domain to determine how best to capture enough relevant features about the phenomena being modeled, without making the model unnecessarily complex.

Second, this model is a new take on biologically-inspired artificial emotion systems, drawing on key mechanisms that appear to be constitutive of emotional phenomena in biological organisms. As Ziemke and Lowe (2009) argue, this kind of approach accounts much more for the whole organism than strongly representationalist approaches that amount to "shallow" modeling of affective states, and generally results in greater flexibility and scope in robot performance. The model validates the claim that an emotion-like system can be beneficial, particularly in uncertain scenarios when other agents are present and many variables need to be accounted for. The fact that the contagion of fear seen in Study Two emerged as a beneficial behavior suggests this model is on the right track. Emotional contagion of this kind is seen in many biological organisms, such as Sherman's (1977) observations of alarm signaling as a way of spreading fear and alertness in voles. De Waal (2008) notes that even birds display contagion responses: If a single bird in a flock is startled into flight, this fear can spread to the rest of flock, causing them all to take off at once. Emotional contagion tends to be an automatic, reflex-like response not requiring any understanding or conscious intention – a bet that reacting in a particular way will be beneficial.

Third, following the lead of O'Bryne and Cañamero (2010), this model supports a broader look at action selection. As the results of PreyG3 in Study Two demonstrated, one cannot look at the action selection architecture of an agent in isolation. The context in which the agent finds itself – both in terms of the environment itself as well as in terms of the behavior of other agents present – can significantly change what kind of behavior the agent displays. Nor can one ignore that action selection is, at its core, a decision about trade-offs. Performing one action necessarily means the agent cannot be performing a different action. Emotion can help assess risk and opportunity cost in a fast, heuristic way. Contagion of fear responses, mentioned above, is an example of this.

It is worthwhile to recall that in Study Two, although more prey in emotion groups survived for the entire trial time, in only three of the six cases did the entire group have a longer average lifespan than the emotionless group. Perhaps, as discussed earlier, this is because the scenario was still relatively simple. The prey only needed to determine whether to forage or to deal with the predator – Parisi and Petrosino (2010) suggest adding a third competing motivation could make the benefits of emotion clearer, as in the hazardous three-resource problems described by O'Bryne et al. (2009) and O'Bryne and Cañamero (2010). It is also possible that the prey evolved in Study Two were not, in fact, optimal solutions, instead the result of premature convergence (see Eiben & Smith, 2003). Adjusting the genetic algorithm further or improving both the predator's and the prey's *seek-food* strategies such that fewer prey would die of “starvation” and more prey that fall into local optima would be “eaten” by the predator could allow better solutions to evolve.

It also seems plausible that if the prey were given a simple “herd” behavior, in which the agents would attempt to stay near other prey agents as in a flock of birds or school of fish (for a classic simulation of this kind of behavior, see Reynolds, 1987), alarm signaling may become even more relevant to the prey and would in fact correlate with an increased average lifespan. In addition, the range over which an alarm could be perceived in this simulation was large in relation to the speed at which they traversed the space. It is likely that any given prey agent sometimes received alarm signals from distant prey that did not predict an immediate threat to the agent, thus selecting for prey that did not necessarily place as great weight on alarm signals. Reducing the range at which an alarm could be perceived could make alarm signals more immediately useful and could correlate with increased average lifespan. Similarly, Jim and Giles (2000) suggested that communication among agents helped more in partially observable environments – the environment in Study Two was a flat plain. It would prove interesting to see how alarm behavior among the prey changes in a different environment. Jim and Giles (2000) examined predator communication rather than prey communication, but is plausible that the same principles apply to any agents attempting to coordinate or inform one another's actions.

5.1. Future directions

There are many directions that this work could be taken. As mentioned earlier, several competing motivations may be necessary for an emotion system to be useful in a robot (Parisi & Petrosino, 2010; O'Bryne et al., 2009). It could be worthwhile to complexify the prey agents' tasks by add competing motivations beyond hunger and fear, such as a desire to mate with other prey agents and produce offspring. For that particular case, a new factor could be added to the genetic algorithm during the evolution of prey agents in which the offspring for subsequent generations would based on which individuals succeeded at “mating” with each other. Relatedly, because emotions are inherently embedded in the flow of time, it could be promising to follow the lead of Morgado and Gasper (2008) and build a more dynamic emotion model that takes change of emotion over time into account. The “receptive field” of emotion abilities could be modified as well – O'Bryne et al. (2009) found that the hormone-like mechanism biasing their prey agent's action selection was only adaptive when linked to distal detection of the predator, not when just linked to proximal detection.

There is great potential for investigating the effects of the environment on prey group behavior. As previously mentioned, obstacles could be added to create a partially-observable environment. The food available to the prey could be varied – either in availability/scarcity, or in the kinds of food patches in the environment. E.g., multiple kinds of food, each providing a different amount of energy to the prey and perhaps also differing in ease of accessibility, as well as evolving the prey's foraging strategy over time, could produce interesting comparisons of different search strategies, as was previously mentioned. One could study resource density in relation to population size, or could examine the co-evolution of search and avoid strategies by evolving both the predator and the prey, building on past work by Nolfi and Floreano (1998).

Modifying the kind of alarm signal, such as using a visual alarm, in conjunction with evolving the predator or changing the environment could have interesting results. Endler (2000) noted that even small changes in lighting can influence how effectively coloring patterns can be detected. Relatedly, in animal societies alarm signals are produced primarily to warn kin – female voles stay quiet around non-kin (Sherman, 1977). It would be relatively simple to extend this model to include “kin” and “non-kin” tagging of prey in the prey group, as well as a relevant motivation for a “desire” to inform kin of danger. This would allow for the study of kin selection effects on behavior and signaling and may prompt a different pattern of alarm utilization to arise.

The complexity of signals could also be increased: prey could be required to discriminate information from multiple signals simultaneously. E.g., a second predator could be added with a corresponding alarm signal, or mating signals from other prey. This would lead to an interesting interplay with eavesdropping, the observation of signaling between others; the predator may use observations of mating signaling between prey to locate them, or prey may use mating signaling between other prey to determine whether or not to court a particular prey. The degree of resolution on the receiving end of a signal could be varied as well, since sensory systems are not perfect and cannot extract or code all aspects of signals, and are dependent on the sensitivity range of organs and receptors (Endler, 2000). Or, following the lead of Floreano, Mitri, Magnenat, and Keller (2007), signaling and communication themselves can be evolved.

Because the goal is to develop more adaptive behavior in artificial agents, primarily robots, it could be worthwhile to base the simulated robots' energy depletion rates on battery decay rates and their energy restoration rates at food patches on battery charge rates. This would allow for an easy transition to real robots and would make the scenario more realistic. Moving into the real world would also emphasize the role of the agent's body, which many researchers argue is particularly important for adaptive behavior and learning (e.g., Johnson, 2007; Pfeifer & Scheier, 1999). One might explore different kinds of bodies in relation to emotion and action selection systems, or evolve the agent's body over time.

6. Conclusions

As artificial agents are used to solve an increasing number of problems, the world in which they operate grows steadily more complex and requires them to perform even more diverse tasks. To function successfully, these agents need to be able to adapt their behavior to many divergent scenarios in real-time. The inclusion of an emotion system in their behavior architectures is a critical step in the this direction.

The model I developed here incorporates several functionally important aspects of biological emotion into a prey agent's action selection mechanisms. Information about the prey's immediate environment via its direct sensory perceptions and information about past motivational states and encounters with the predator are dynamically integrated to inform the prey's future actions, allowing it to continuously adapt to an existence in an ever-changing world.

References

Arbib, M., & Fellous, JM. (2004). Emotions: from brain to robot. *TRENDS in Cognitive Science,* 8 (12): 554-561.

Arkin, R. (2005). *Moving Up the Food Chain: Motivation and Emotion in Behavior-based Robots.* In J. Fellous & M. Arbib, (Eds.), *Who Needs Emotions: The Brain Meets the Robot.* London: Oxford University Press.

Arkin, R., Fujita, M., Takagi, T., & Hasegawa, R. (2003). An ethological and emotional basis for human-robot interaction. *Robotics and Autonomous Systems*, 42*(3-4)*.

Bekoff, M. (2000). Animal Emotions: Exploring Passionate Natures. *BioScience,* 50 (10): 861-870.

Berg, H. (2000). Motile behavior of bacteria. *Physics Today,* 53: 24-29.

Bradbury, J., & Vehrencamp, S. (1998). *Principles of Animal Communication.* Sunderland, MA: Sinauer Associates, Inc.

Bruce, A., Nourbakhsh, I., & Simmons, R. (2002). The role of expressiveness and attention in human-robot interaction. In *Proceedings of the IEEE International Conference on Robotics and Automation*, 4138-4142.

Cañamero, L. (2005). Emotion understanding from the perspective of autonomous robots research. *Neural Networks,* 18, 445-455.

Castellano, G., Aylett, R., Paiva, A., & McOwan, P. (2008). Affect recognition for interactive companions. In *Workshop on Affective Interaction in Natural Environments (AFFINE), ACM International Conference on Multimodal Interfaces (ICMI'08)*. ACM.

Damasio, A. (1999). *The Feeling of What Happens: Body and Emotions in the Making of Consciousness*. New York, NY: Harcourt Brace & Company.

Damasio, A. (2010). *The Self Comes to Mind: Constructing the conscious brain.* New York, NY: Pantheon Books.

Dawkins, M. (2000). Animal minds and animal emotions. *American Zoology,* 40: 883-888.

de Waal, F. (2008). Putting the altruism back into altruism: the evolution of empathy. *Annual Review of Psychology*, 59: 279-300.

Eiben, A. & Smith, J. (2003). *Introduction to Evolutionary Computing.* Springer, Berlin, Heidelberg, New York.

Elliott, P., Cowan, I., & Holling, C. (1977). Prey capture by the African lion. *Canadian Journal of Zoology,* 55(11): 1811-1828.

Endler, J. (2000). Evolutionary implications of the interaction between animal signals and the environment. In Y. Espmark, T. Amundsen, & G. Rosenqvist, *Animal Signals: Signaling and signal design in animal communication* (pp. 11-46). Trondheim, Norway: Tapir Academic Press.

Evans, D. (2002). The search hypothesis of emotion. *British Journal of Philosophical Science,* 53: 497-509.

Fellous, JM. (2004). From human emotions to robot emotions. In *Proceedings of the AAAI Spring Symposium-Architectures for Modeling Emotion*, TR SS-04-02. Menlo Park, CA: AAAI Press.

Floreano, D., & Keller, L. (2010). Evolution of adaptive behaviour in robots by means of Darwinian selection. *PloS Biology,* 8*(1),* e1000292.

Floreano, D., Mitri, S., Magnenat, S., & Keller, L. (2007). Evolutionary conditions for the emergence of communication. *Current Biology,* 17, 514-519.

Gordon, D. (2000). *Ants at Work: How an Insect Society is Organized*. New York, NY: The Free Press.

Hasson, C., & Gaussier, P. (2010). From conditioning of a non-specific sensor to emotional regulation of behavior. In K. Diamantaras, W. Duch, & L. Iliadis, (Eds.), *ICANN Part II, LNCS 6353,* pp. 317 – 326.

Hudlicka, E. (2008). *What are we modeling when we model emotions?* In *Proceedings of the AAAI Spring Symposium "Emotion, Personality, and Social Behavior,"* Menlo Park, CA: AAAI Press.

Jim, K., & Giles, C. (2000). Talking helps: Evolving communicating agents for the predator-prey pursuit problem. *Artificial Life,* 6*(3)*, 237-254.

Johnson, M. (2007). *The Meaning of the Body.* Chicago, IL: The University of Chicago Press.

Kagan, J. (2007). *What Is Emotion?* Binghamton, NY: Vail-Ballou Press.

Kelly, C. (1992). Resource Choice in *Cuscuta europaea.* Proceedings of National Academy of Science, USA, 89: 12194-12197.

Krebs, J., & Dawkins, R. (1984). Animal Signals: Mind-Reading and Manipulation. In J.R.

Krebs, N.B. Davies, (Eds.), *Behavioural ecology: an evolutionary approach, 2nd ed.* (pp. 380–402). Oxford: Blackwell Scientific.

Ledoux, J., & Phelps, E. (2008). Emotional networks in the brain. In M. Lewis, J. Haviland-Jones, & L. Barrett, *Handbook of Emotions* (pp. 159 - 179). New York, NY: Guilford Press.

Lee-Johnson, C. & Carnegie, D. (2007). Emotion-based parameter modulation for a hierarchical mobile robot planning and control architecture. *Proceedings of the 2007 IEEE/RSJ International Conference on Intelligent Robots and Systems,* 2839 – 2844.

Loula, A., Gudwin, R., El-Hani, C., & Queiroz, J. (2010). Emergence of self-organized symbol-based communication in artificial creatures. *Cognitive Systems Research,* 11, 131 – 147.

MacLennan, B. (2008). A Protophenomenological Analysis of Synthetic Emotion in Robots. Unedited draft of “Robot React, but Can They Feel?” (*Handbook of Research on Synthetic Emotions and Sociable Robotics: New Applications in Affective Computing and Artificial Intelligence*, edited by Jordi Vallverdú and David Casacuberta, Hershey, NJ: IGI Global, 2009, pp. 133–53).

Miachaud, F., & Audet, J. (2001). Using motives and artificial emotions for long-term activity of an autonomous robot. *AGENTS '01,* 188 – 189.

Murphy, R., Lisetti, C., Tardif, R., Irish, L., & Gage, A. (2002). Emotion-based control of cooperating heterogeneous mobile robots. *IEEE Transactions on Robotics and Automation,* 18*(5),* 744 – 757.

Murray, J. & Cañamero, L. (2009). Developing Preferential Attention to a Speaker: A Robot Learning to Recognize its Carer. In *Proceedings of the IEEE Symposium Series on Computational Intelligence: ALife 2009*, Nashville, TN, USA.

Nolfi, S., & Floreano, D. (1998). Co-evolving predator and prey robots: do “arms races” arise in artificial evolution. *Artificial Life,* 4, 311-335.

Oatley, K., & Jenkins, J. (1998). *Understanding Emotions.* Oxford, UK: Blackwell.

O'Bryne, C., & Cañamero, L. (2010). Emotions in decisions of life and death – its role in brain-body-environment interactions for predator and prey. In *Proceedings of the Alife XII Conference,* Odense, Denmark, 812-819.

O'Bryne, C., Cañamero, L., & Murray, J. (2009). The importance of the body in affect-modulated action selection: A case study comparing proximal versus distal perception in a prey-predator scenario. *Affective Computing and Intelligent Interaction (ACII 2009): Proceedings 3rd International Conference,* IEEE Computer Society Press.

Panksepp, J. (1998). *Affective Neuroscience: The Foundations of Human and Animal Emotion.* New York, NY: Oxford University Press.

Parisi, D., & Petrosino, G. (2010). Robots that *have* emotions. *Adaptive Behavior,* 0*(00)*, 1-17.

Pfeifer, R., & Scheier, C. (1999). *Understanding Intelligence.* Cambridge, MA: MIT Press.

Rendall, D., Owren, M., & Ryan, M. (2009). What do animal signals mean? *Animal Behaviour,* 78, 233-240.

Reynolds, C. (1987). Flocks, herds, schools: A distributed behavior model. *Computer Graphics,* 21 *(4):* 25-34.

Rolls, E. (1999). *The brain and emotion.* New York, NY: Oxford University Press.

Rolls, E. (2005). *Emotion Explained.* New York, NY: Oxford University Press.

Rumbell, T., Barnden, J., Denham, S., & Wennekers, T. (2010). A generic framework for the analysis of emotion mechanisms in autonomous agents. In *Proceedings of BICS 2010 - Brain Inspired Cognitive Systems 14-16 July 2010*, Madrid, Spain.

Schmickl, T., Möslinger, C., Thenius, R., Crailsheim, K. (2007). Bio-inspired Navigation of Autonomous Robots in Heterogenous Environments. *International Journal of Factory Automation, Robotics and Soft Computing,* 3, 164-170

Smith, C. (2004). A Functional Perspective on Emotion Elicitation: Some Considerations for the Development of Emotional Architectures. In E. Hudlicka, & L. Cañamero (Eds.), *Architectures for modeling emotion: Cross-disciplinary foundations. Papers from the 2004 AAAI spring symposium*, (pp. 135-143). Menlo Park, CA: AAAI Press.

Tooby, J., & Cosmides, L. (2008). The evolutionary psychology of the emotions and their relationship to internal regulatory variables. In M. Lewis, J. Haviland-Jones, & L. Barrett, *Handbook of Emotions* (pp. 114 - 137). New York, NY: Guilford Press.

Weitzenfeld, A. (2008). A prey catching and predator avoidance neural-schema architecture for single and multiple robots. *Journal of Intelligent and Robotic Systems, Springer,* 51*(2)*, 203-233.

Ziemke, T., & Lowe, R. (2009). On the role of emotion in embodied cognitive architectures: From organisms to robots. *Cognitive Computation:* 1, 104-117.

Figures

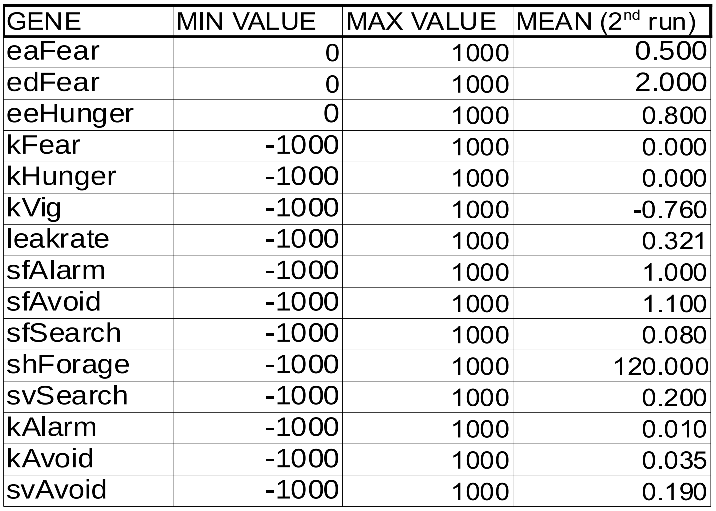

GENE	MIN VALUE	MAX VALUE	MEAN (2nd run)
eaFear	0	1000	0.500
edFear	0	1000	2.000
eeHunger	0	1000	0.800
kFear	-1000	1000	0.000
kHunger	-1000	1000	0.000
kVig	-1000	1000	-0.760
leakrate	-1000	1000	0.321
sfAlarm	-1000	1000	1.000
sfAvoid	-1000	1000	1.100
sfSearch	-1000	1000	0.080
shForage	-1000	1000	120.000
svSearch	-1000	1000	0.200
kAlarm	-1000	1000	0.010
kAvoid	-1000	1000	0.035
svAvoid	-1000	1000	0.190

Figure 3.1: The valid ranges for each gene in the prey's genome in Study 1. The initial genomes for the second run of the genetic algorithm was randomly generated in a normal distribution with the mean in the fourth column and a standard deviation of one.

The genes *eaFear* and *edFear* are the exponents in the *fear* equation in Section 3.2: *a* and *d*, respectively; kFear is k_f in the same equation. *EeHunger* is the exponent *n* in the *hunger* equation; *kHunger* is k_h. From the *vigilance* equation, *kVig* corresponds to k_v and *leakrate* to *L*. The genes *sfAlarm*, *sfAvoid*, and *sfSearch* are the sensitivity factors modifying *fear* in the *alarm*, *avoid*, and *search* equations; *svSearch* and *svAvoid* are the sensitivity factors modifying *vigilance* in the *search* and *avoid* equations. The genes *kAlarm*, *kAvoid* are the constants k_{al} and k_{av}, respectively, and *shForage* is the sensitivity factor modifying *hunger* in the *forage* equation.

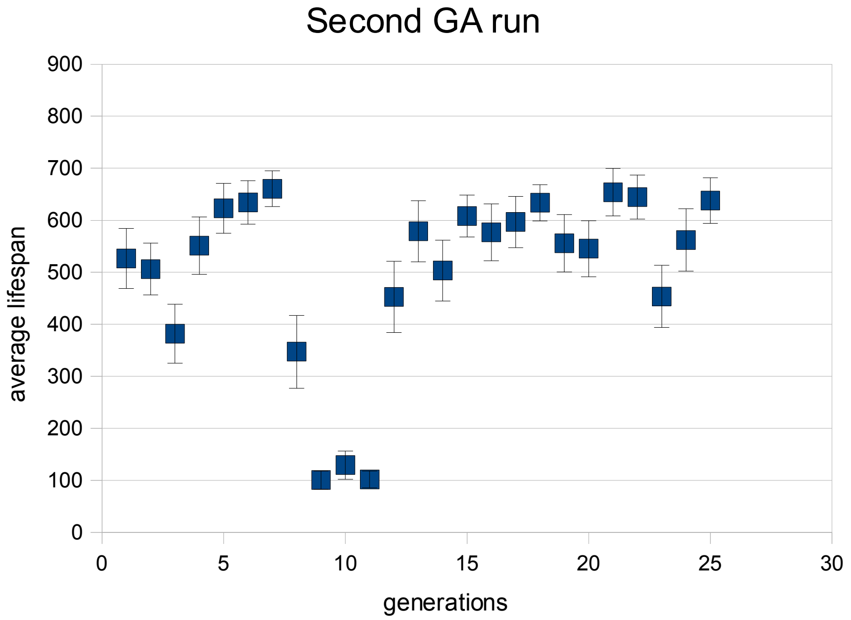

Figure 3.2: The average lifespan of the prey agents for each of the twenty-five generations of the second genetic algorithm in Study One.

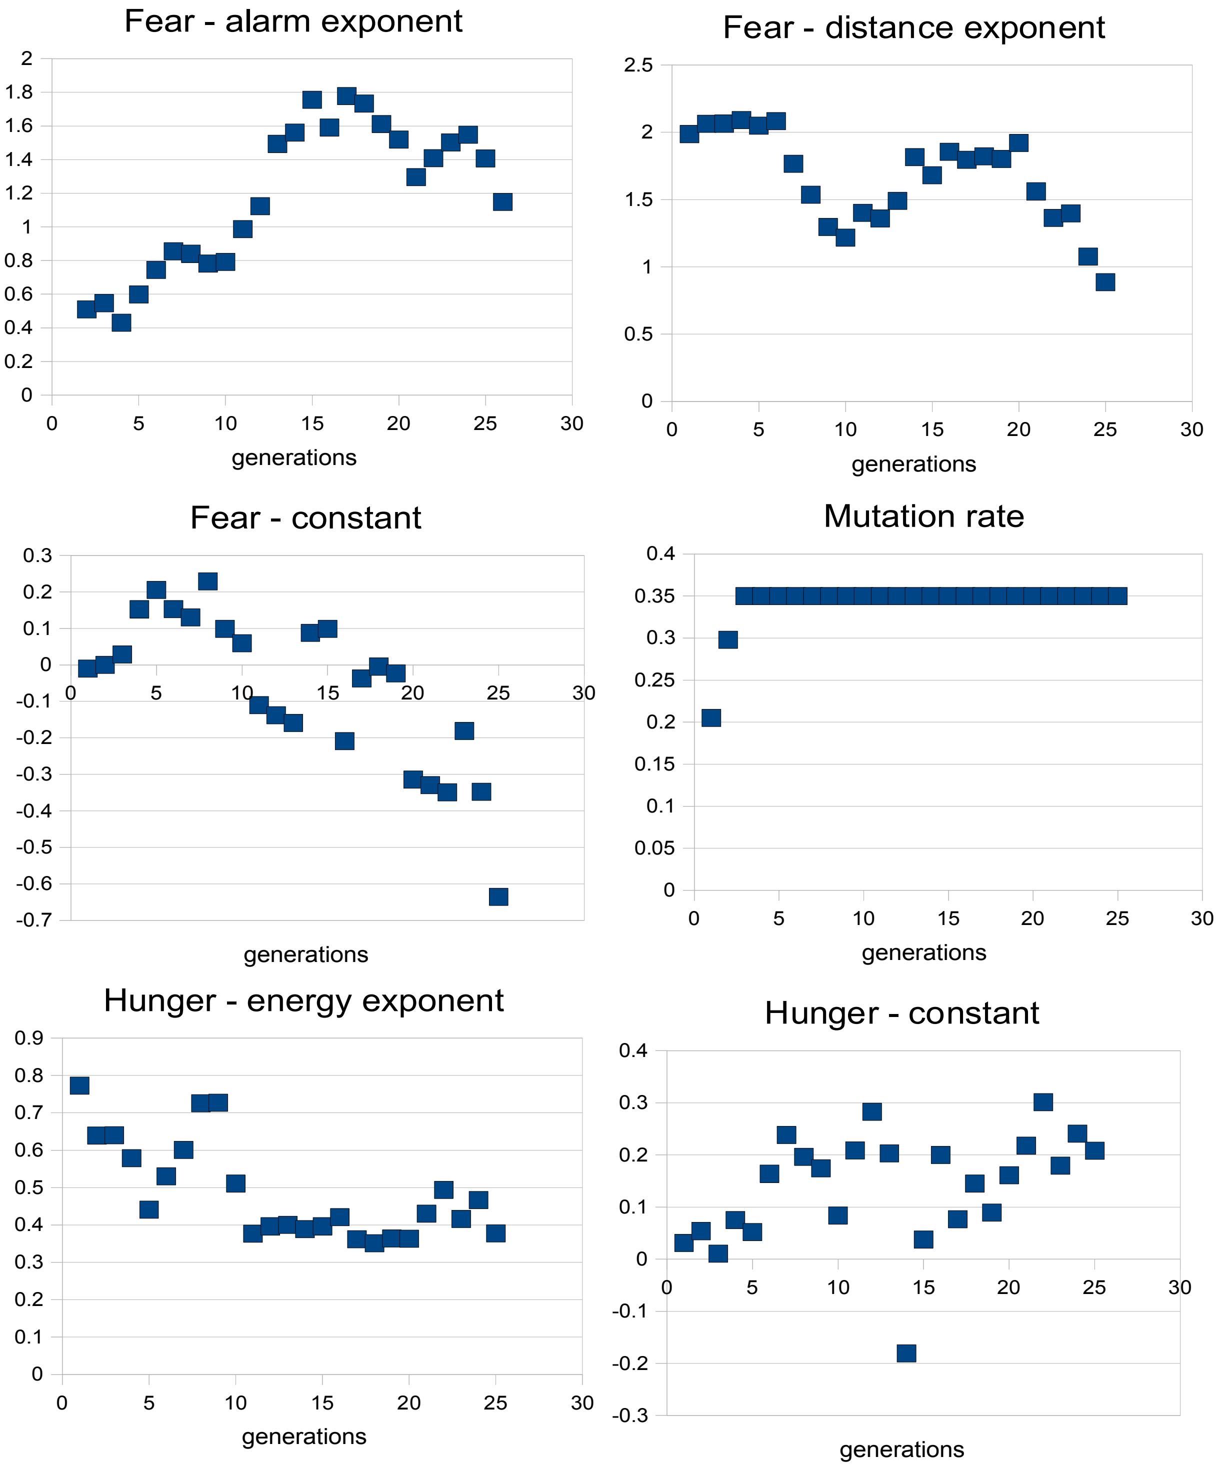

Figure 3.3a: The change in each gene across the twenty-five generations of the second genetic algorithm in Study 1. Each dot represents the average of that gene for that generation.

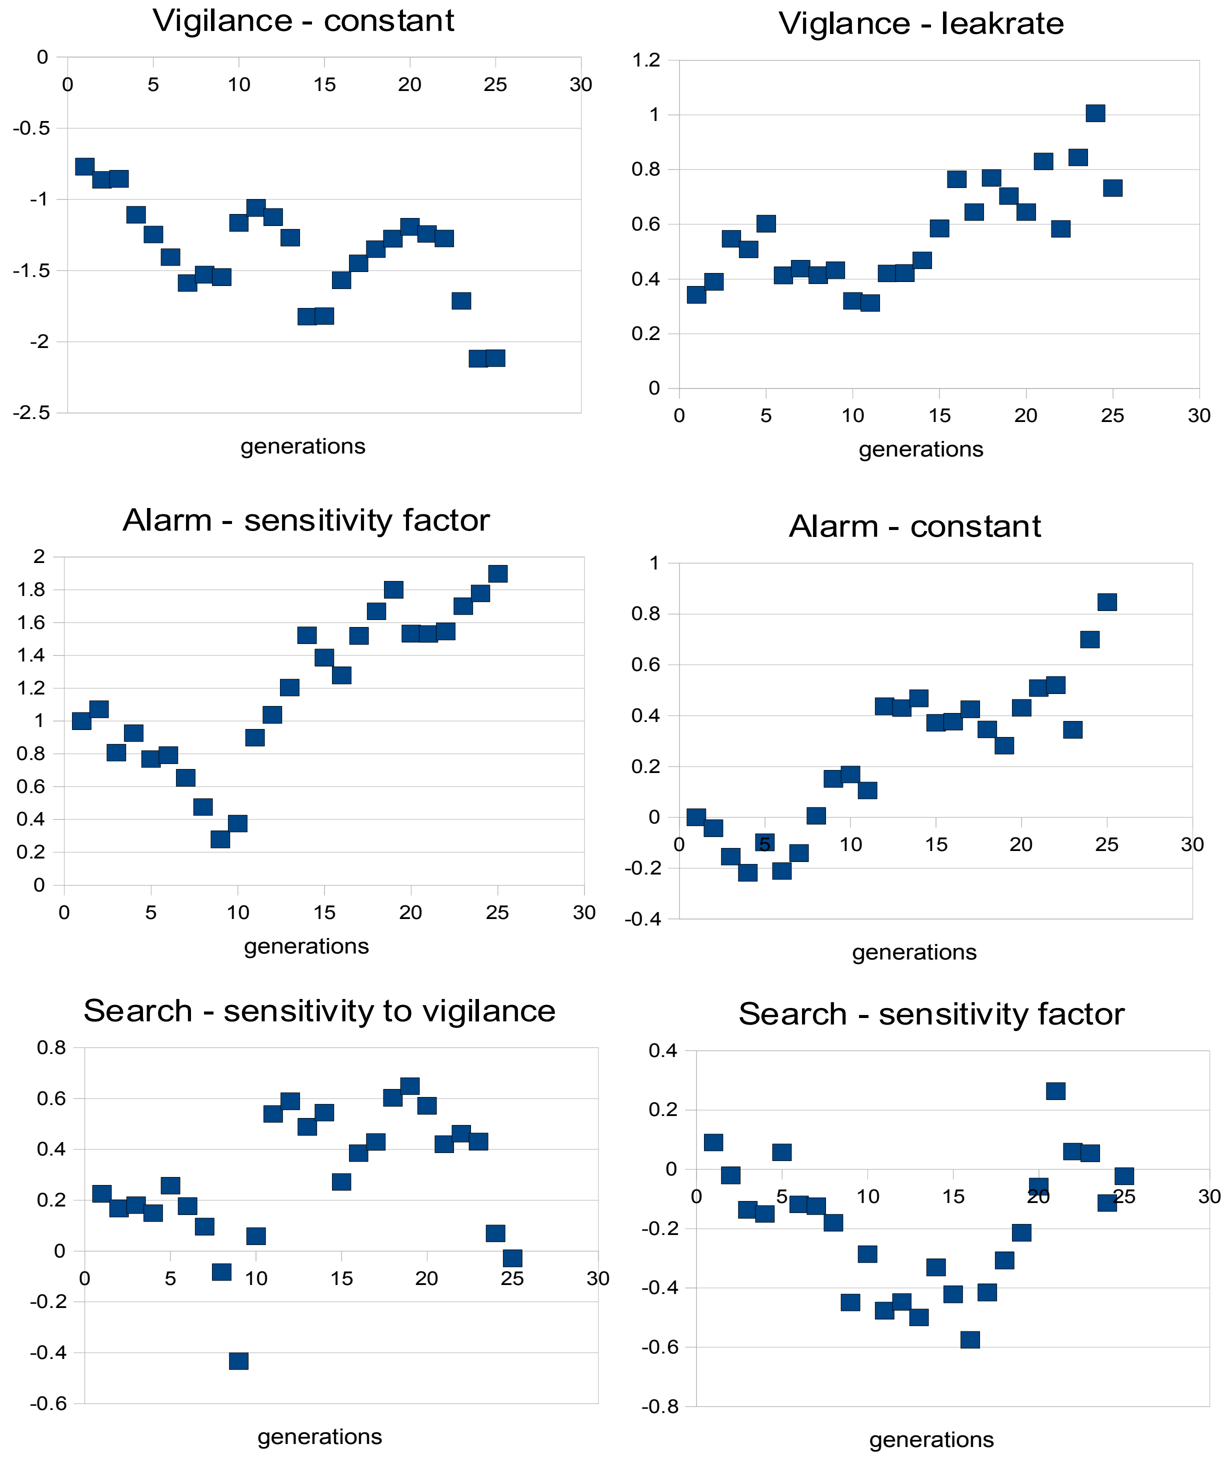

Figure 3.3b: The change in each gene across the twenty-five generations of the second genetic algorithm in Study 1, continued.

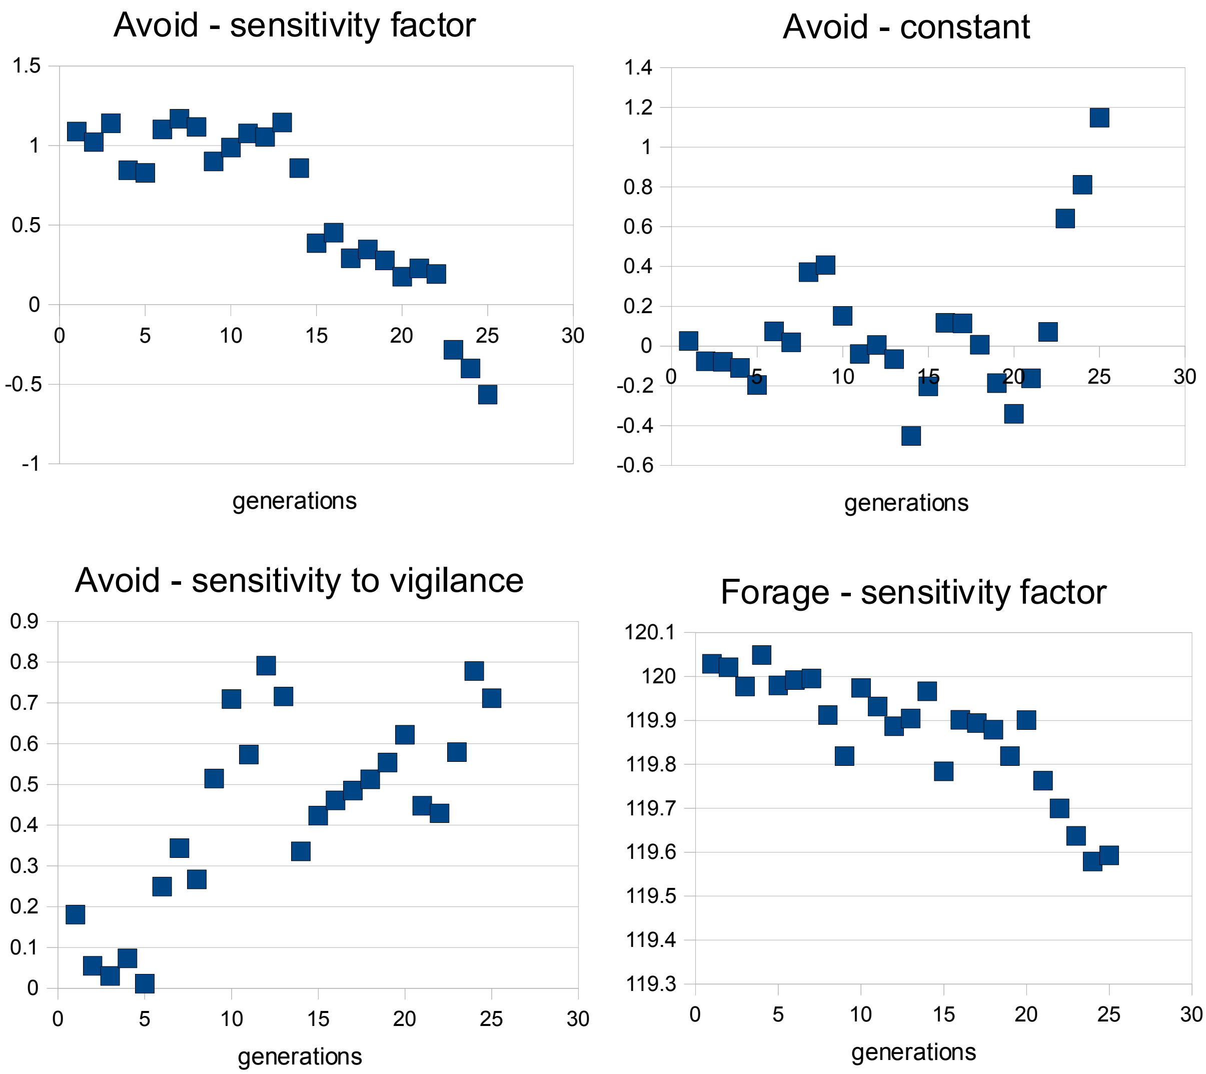

Figure 3.3c: The change in each gene across the twenty-five generations of the second genetic algorithm in Study 1, continued.

GENE	Prey1	Prey2	Prey3	Prey4
eaFear	0.1983	1.5816	1.3656	0.0000
edFear	1.5858	0.7817	1.1889	1.3102
eeHunger	0.3282	0.7347	0.3034	0.3872
kFear	0.4252	-0.0392	0.1458	-0.0103
kHunger	-0.0620	-0.0061	-0.1613	-0.0833
kVig	-1.8761	-0.4422	-1.1776	-2.3033
leakrate	0.6967	1.4661	0.2281	0.0100
sfAlarm	-0.2523	0.5231	1.5866	-0.4356
sfAvoid	1.1685	0.9146	0.0539	0.2802
sfSearch	0.1402	0.4176	0.5860	0.9522
shForage	119.9738	120.0733	119.5470	120.4045
svSearch	-0.1717	0.0197	0.1074	0.1594
kAlarm	0.4873	0.8831	1.7890	-0.2043
kAvoid	-0.0678	-0.7629	-0.3622	-0.3489
svAvoid	0.2868	0.0227	1.7370	0.4563
mut_rate	0.3500	0.3500	0.3500	0.3500

Figure 3.4a: The genomes of the four best prey agents evolved in Study 1.

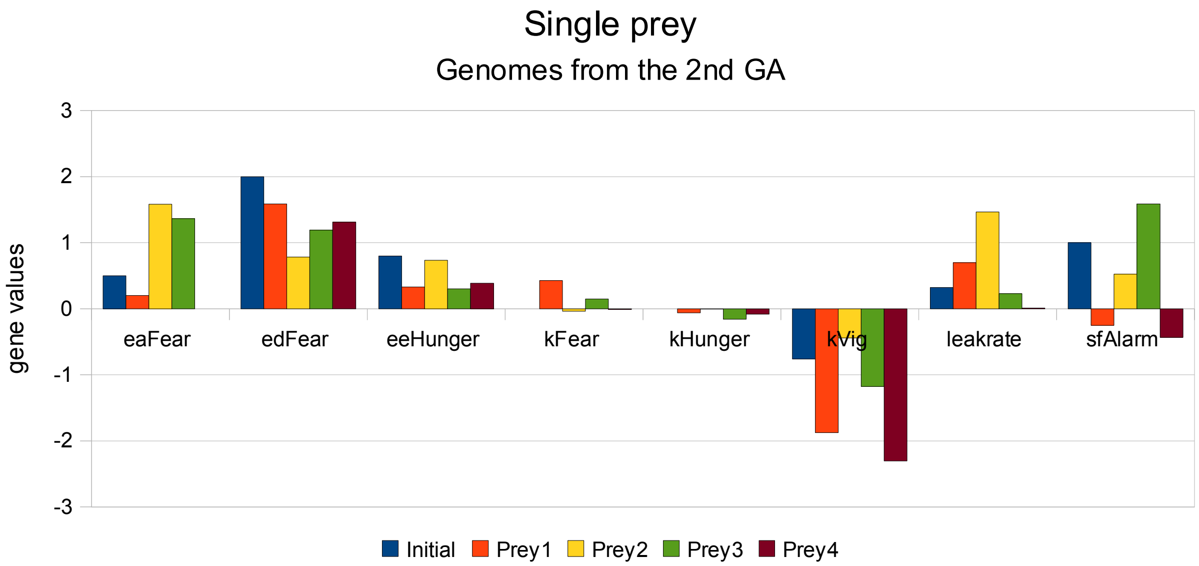

Figure 3.4b: The first eight genes in the genomes of the four best prey agents evolved in Study 1, displayed visually and compared to the “seed” genome for this run of the genetic algorithm (see Figure 3.1).

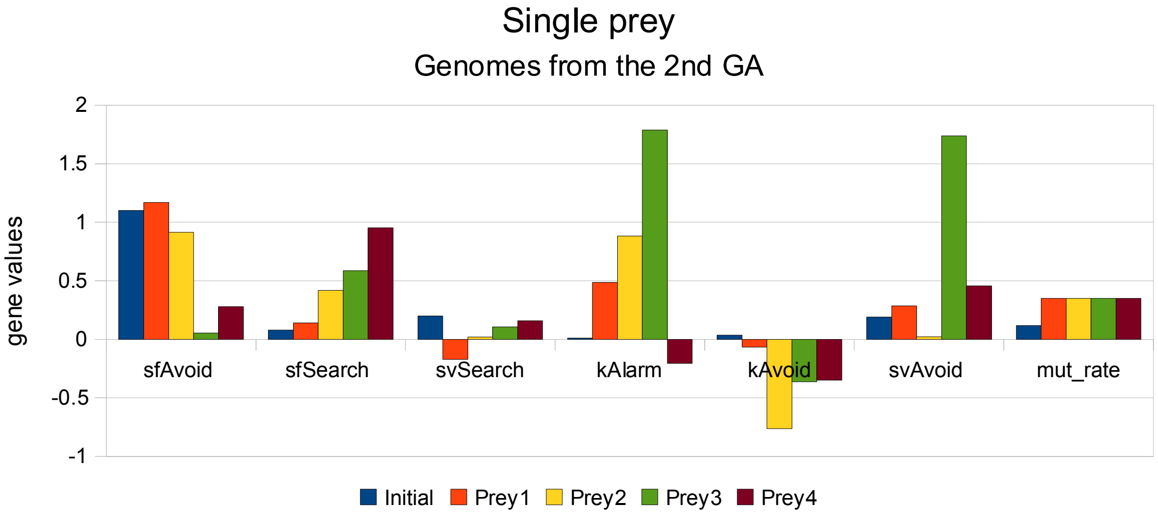

Figure 3.4c: A continuation of the genomes of the four best prey agents evolved in Study 1, displayed visually and compared to the "seed" genome for this run of the genetic algorithm (see Figure 3.1).

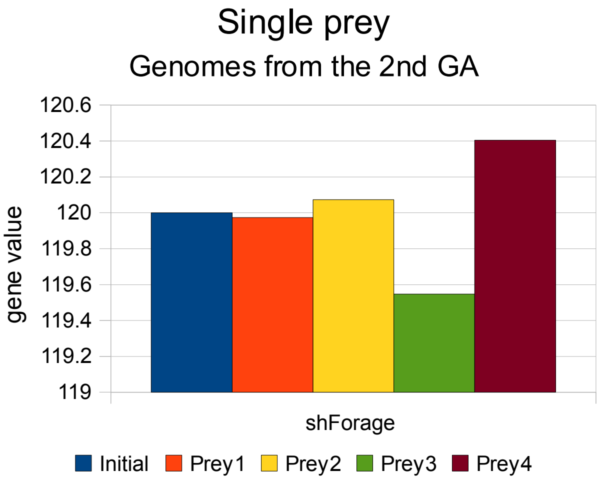

Figure 3.4d: The last gene in the genomes of the four best prey agents evolved in Study 1, displayed visually, compared to the "seed" genome for this run of the genetic algorithm.

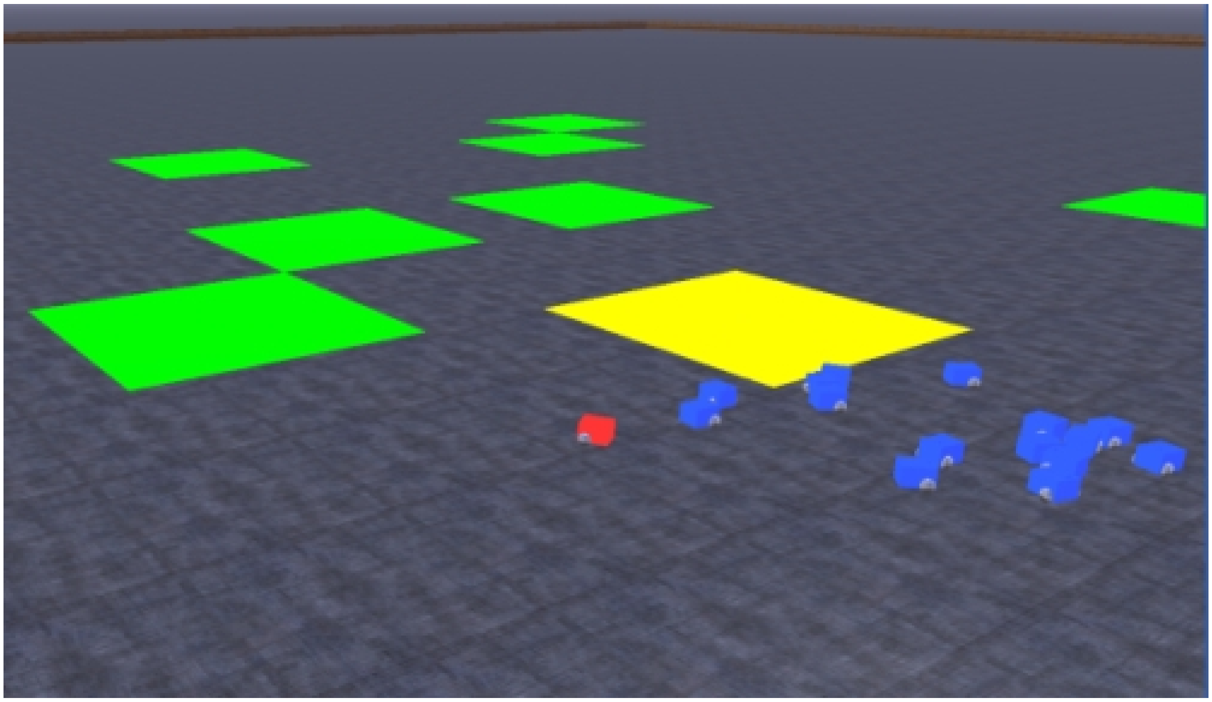

Figure 3.5: A screenshot of the simulation environment. The safe zone is the yellow square; the green squares are food patches. The predator was color-coded red, while each prey agent was blue.

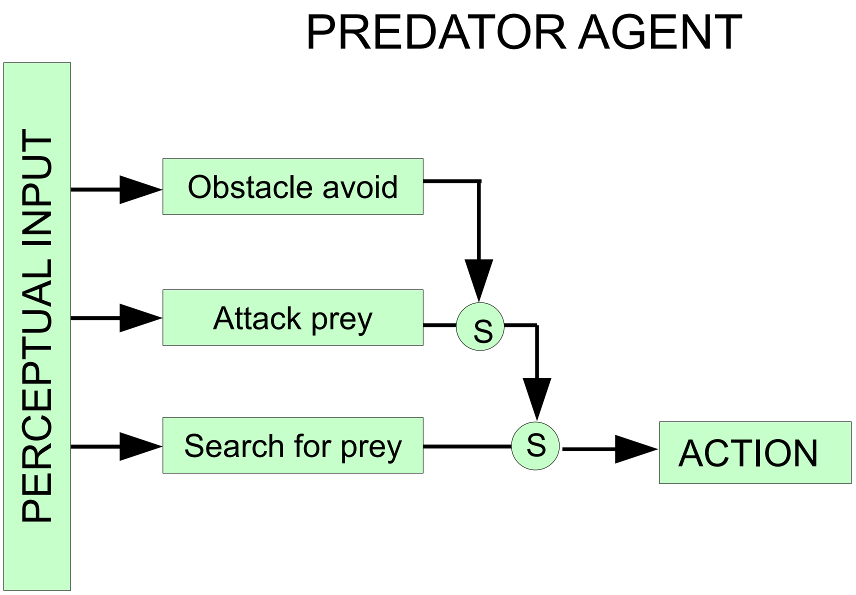

Figure 3.6: Diagram of the predator's action-selection and behavior hierarchy.

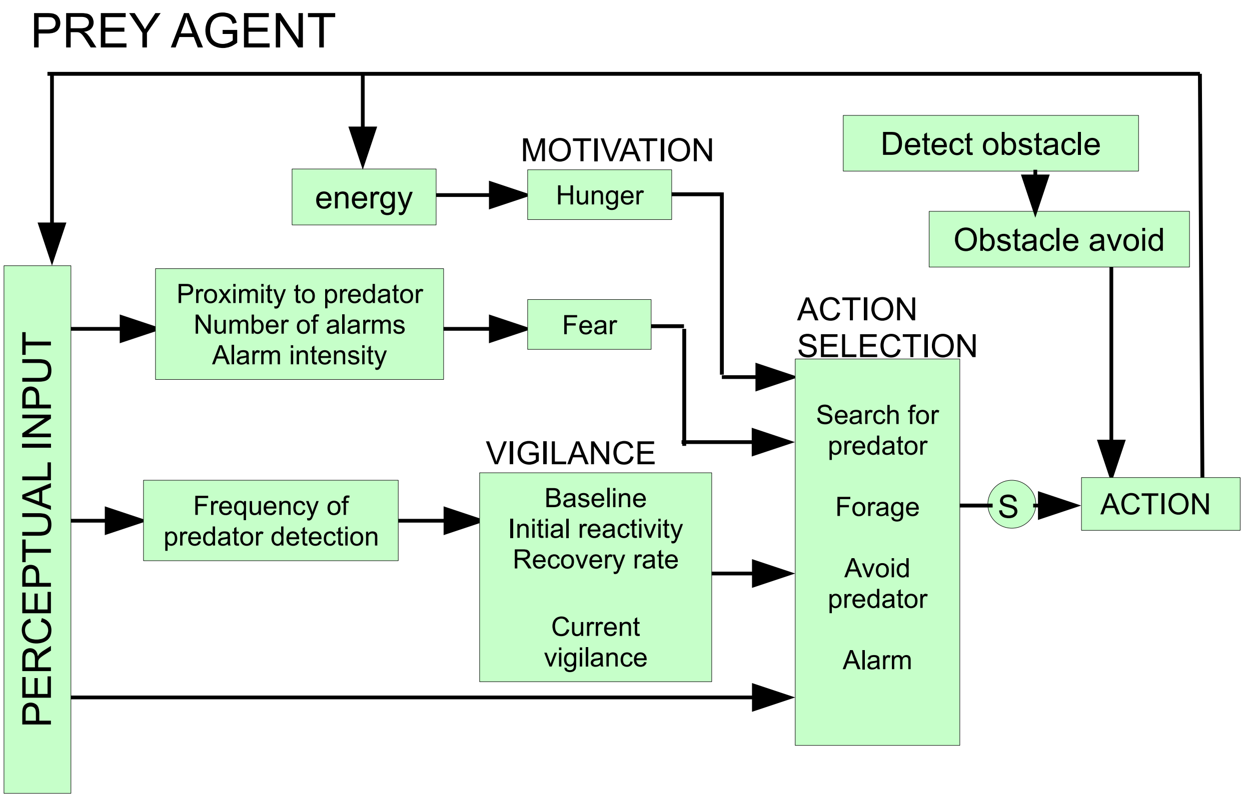

Figure 3.7: Diagram of the prey's action-selection and behavior functions.

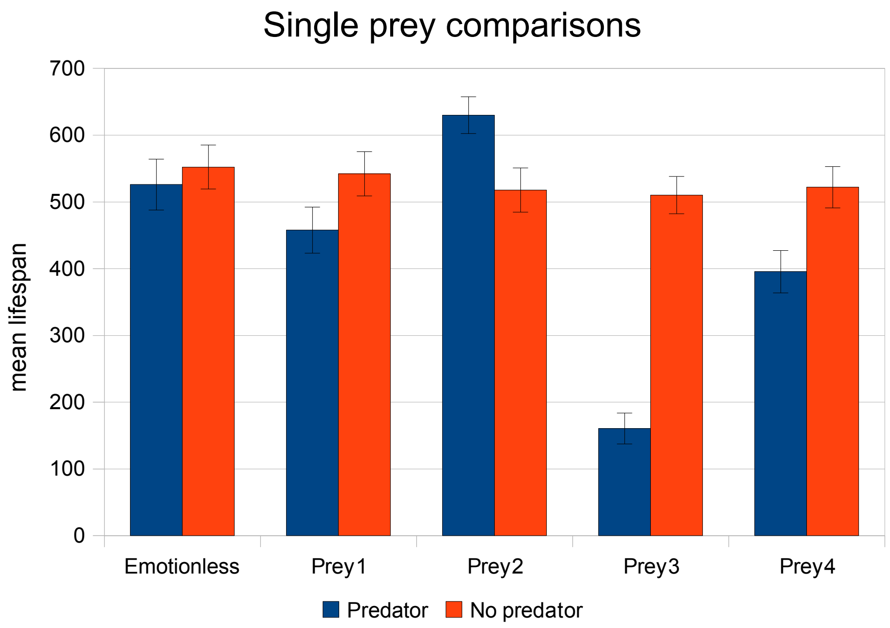

Figure 3.8: The mean lifespans of the four best prey agents tested in Study 1 when a predator was present in the environment (blue) and when no predator was present (orange).

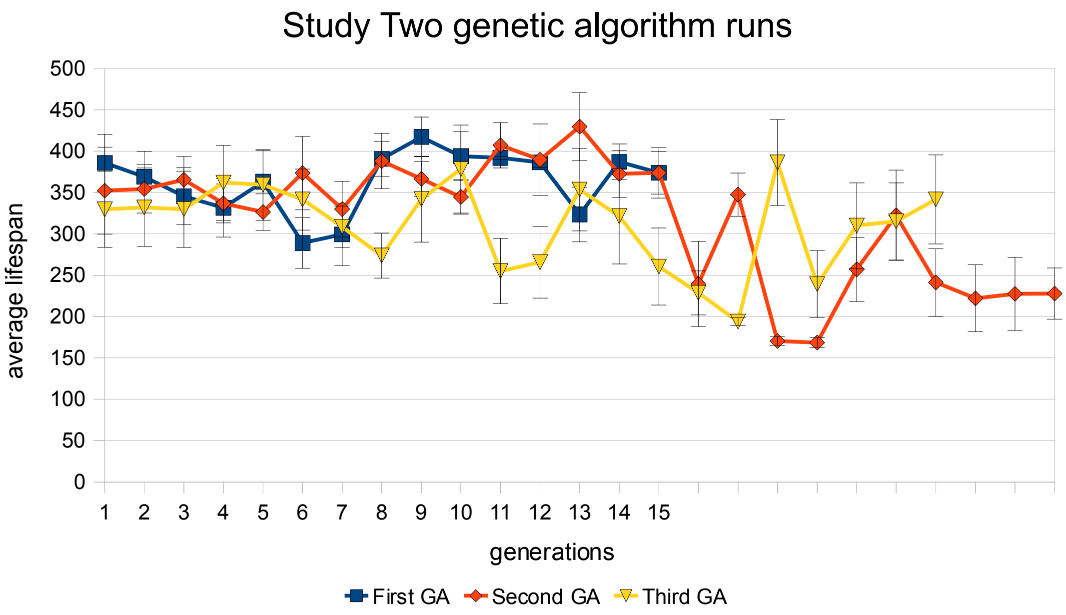

Figure 4.1: The mean lifespan over time for each of the genetic algorithm runs in Study Two.

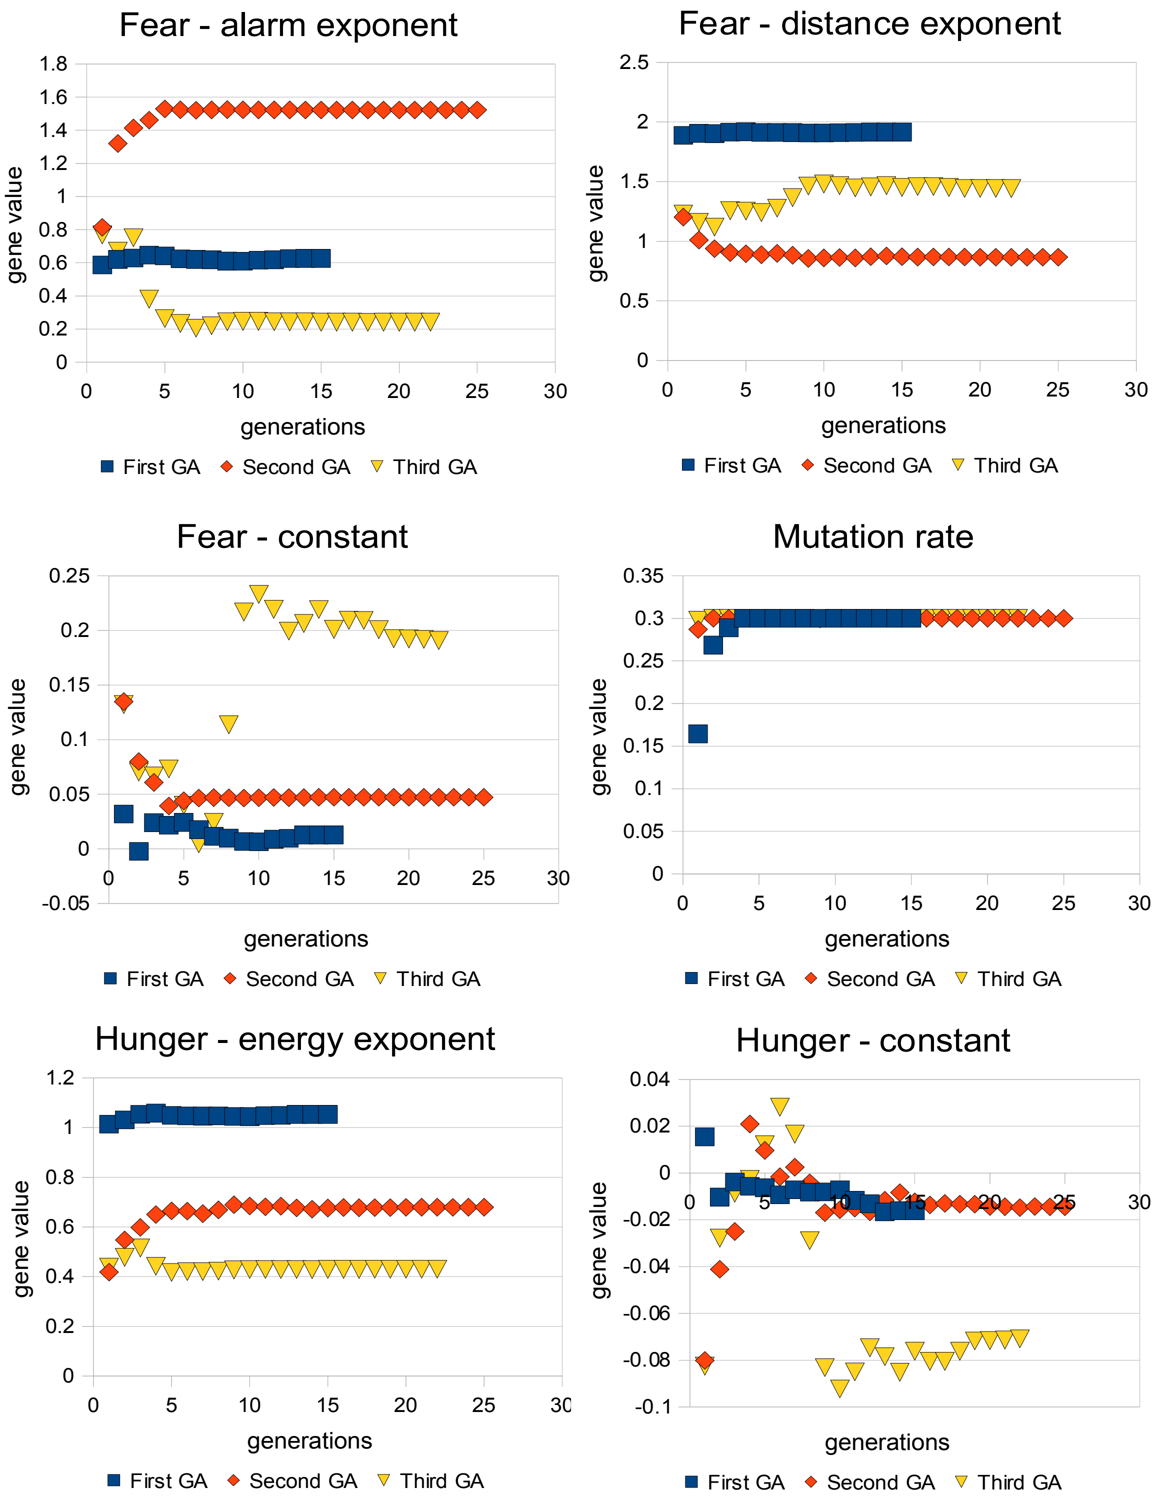

Figure 4.2a: The genetic change over time for each of the genetic algorithm runs in Study Two.

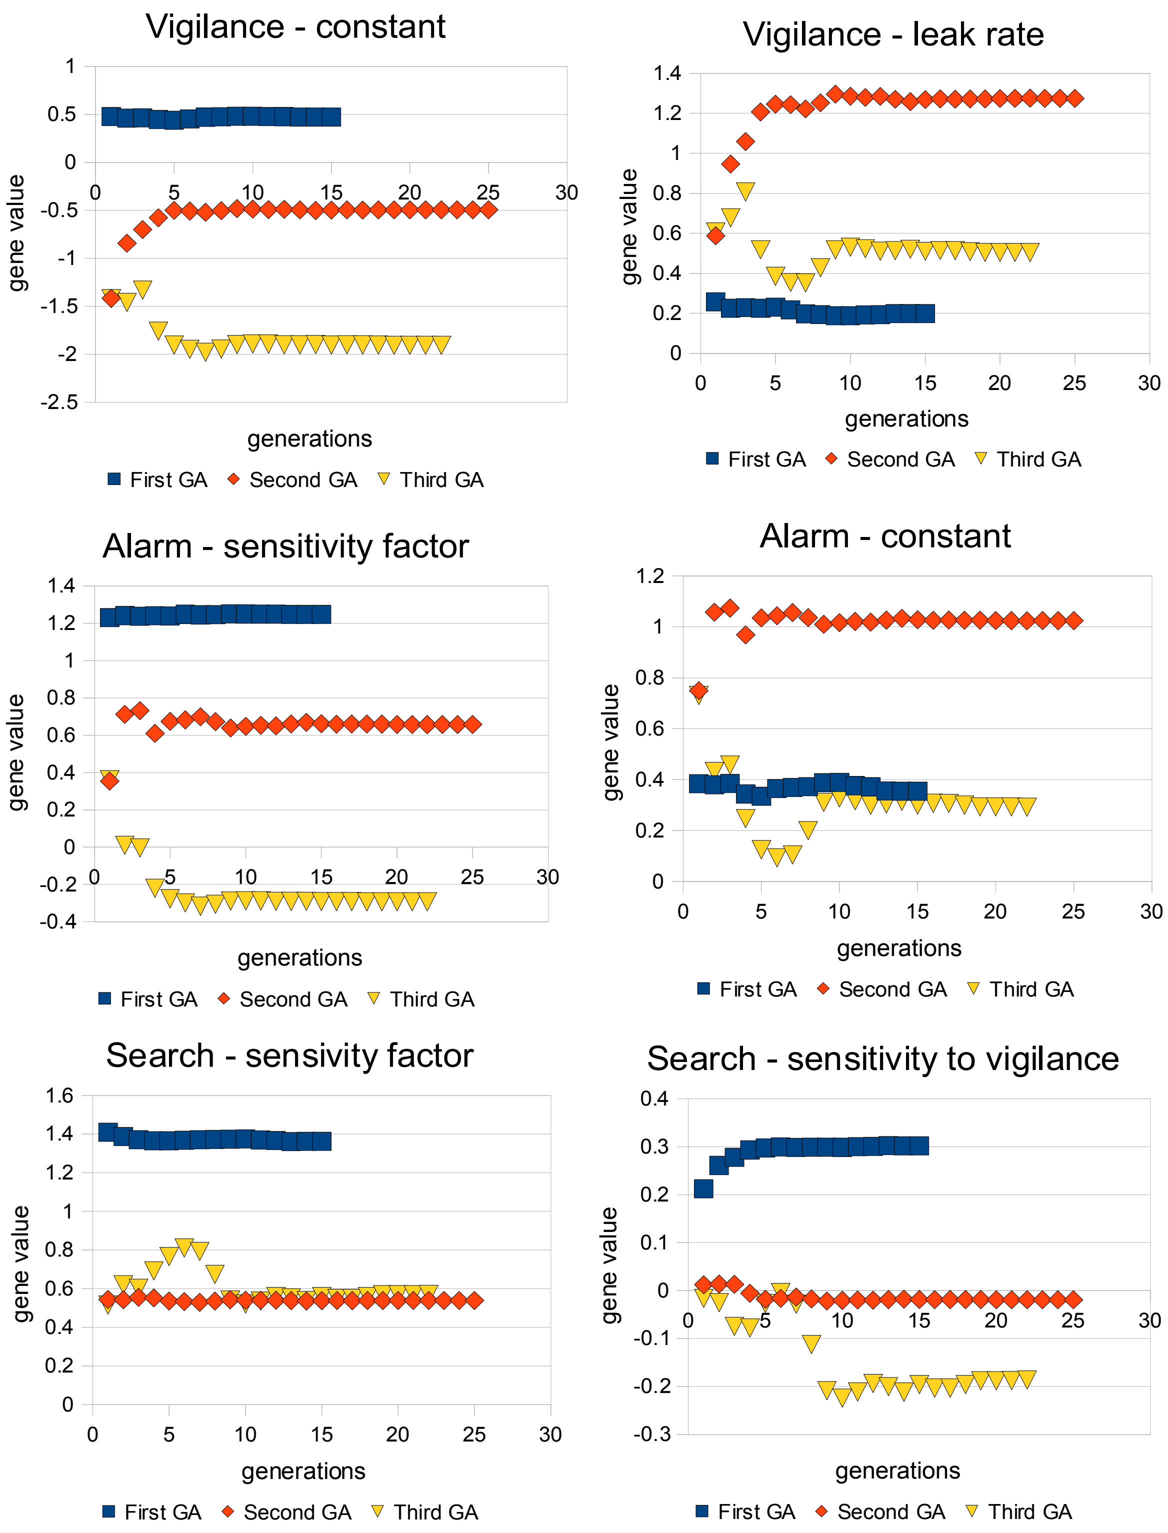

Figure 4.2b: The genetic change over time for each of the genetic algorithm runs in Study Two.

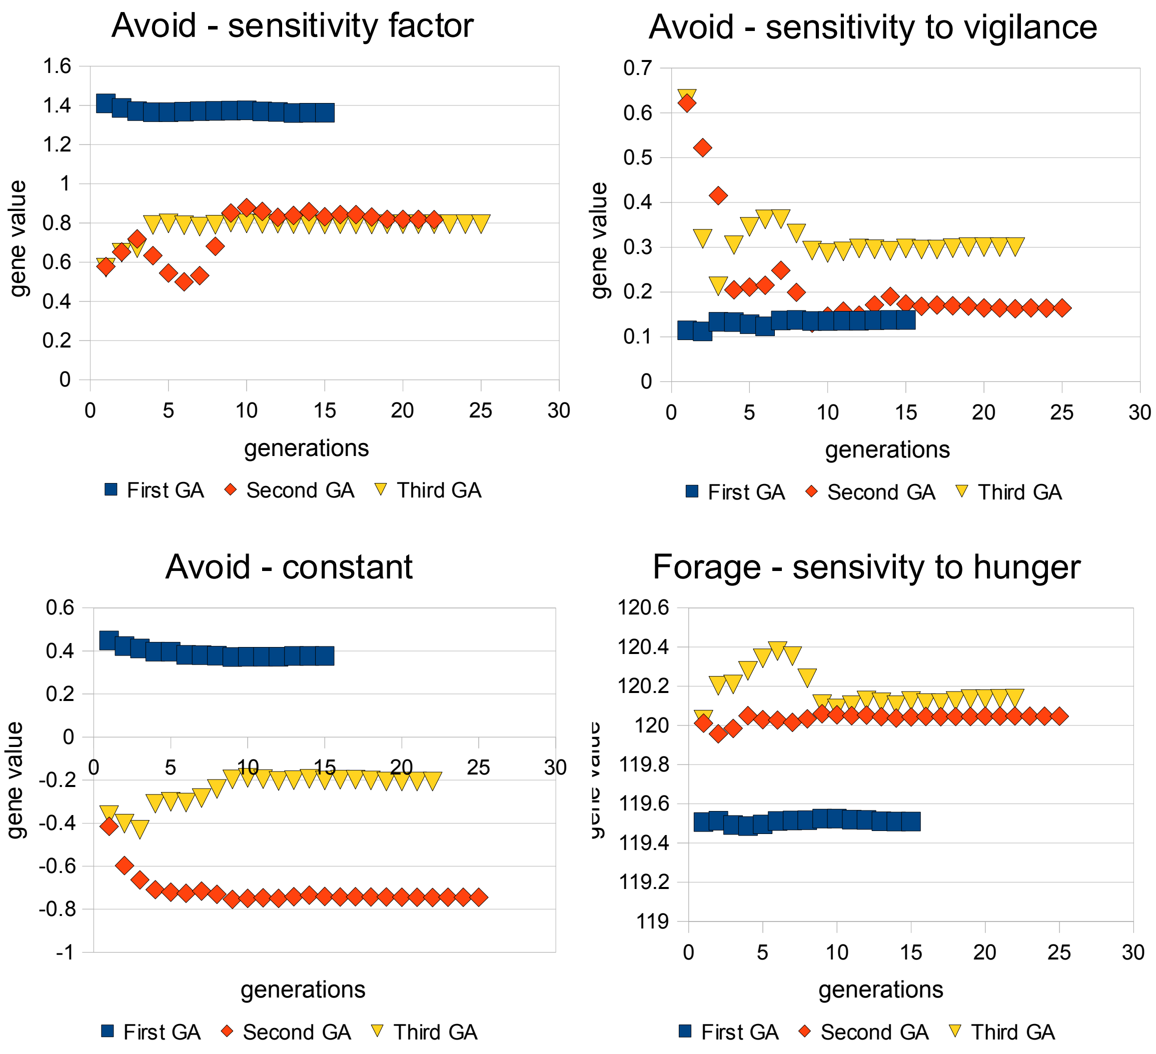

Figure 4.2c: The genetic change over time for each of the genetic algorithm runs in Study Two.

GENE	PreyG1	PreyG2	PreyG3	PreyG4	PreyG5
eaFear	0.6260	1.5224	1.5214	0.2420	0.2397
edFear	1.9154	0.8955	0.8611	1.4687	1.4307
eeHunger	1.0527	0.6593	0.6820	0.4292	0.4298
kFear	0.0128	0.0471	0.0476	0.2201	0.1790
kHunger	-0.0162	0.0021	-0.0183	-0.0862	-0.0648
kVig	0.4729	-0.5121	-0.4922	-1.8994	-1.9097
leakrate	0.1981	1.2331	1.2803	0.5202	0.4961
sfAlarm	1.2475	0.6875	0.6536	-0.2907	-0.2948
sfAvoid	1.0749	0.7867	0.7935	0.8590	0.7988
sfSearch	1.3620	0.5335	0.5394	0.5378	0.5845
shForage	119.5102	120.0225	120.0506	120.1009	120.1508
svSearch	0.3019	-0.0159	-0.0192	-0.2130	-0.1759
kAlarm	0.3551	1.0476	1.0221	0.3164	0.2810
kAvoid	0.3771	-0.7186	-0.7488	-0.1936	-0.2120
svAvoid	0.1378	0.2297	0.1548	0.2916	0.3042
mut_rate	0.3000	0.3000	0.3000	0.3000	0.3000

Figure 4.3: A table of the genomes of the five best prey agents in Study Two, used for the emotion prey group versus emotionless prey group comparisons.

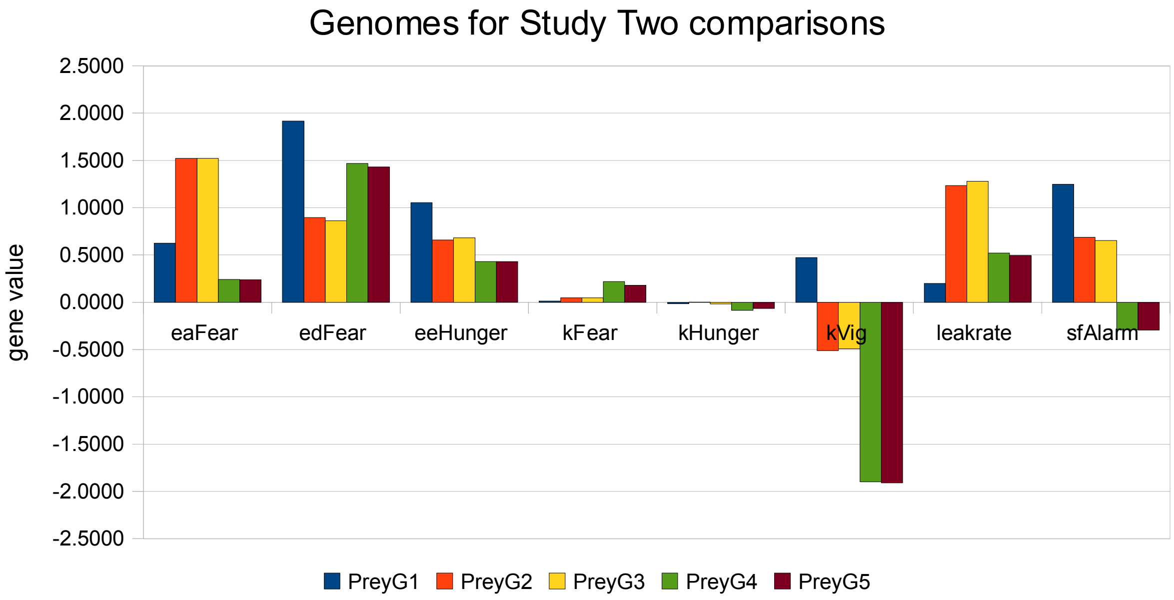

Figure 4.4a: The first eight genes in the genomes of the five best prey agents in Study Two, displayed visually.

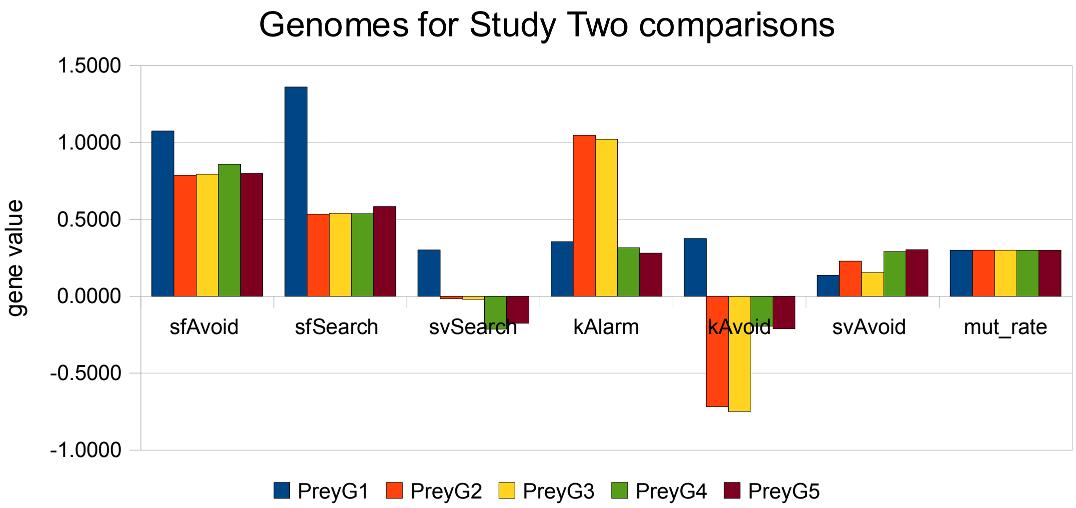

Figure 4.4b: A continuation of the genomes of the five best prey agents in Study Two, displayed visually.

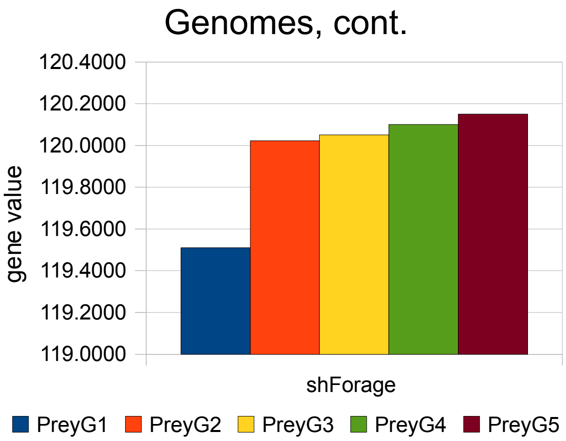

Figure 4.4c: The last gene in the genomes of the five best pre agents in Study Two, displayed visually.

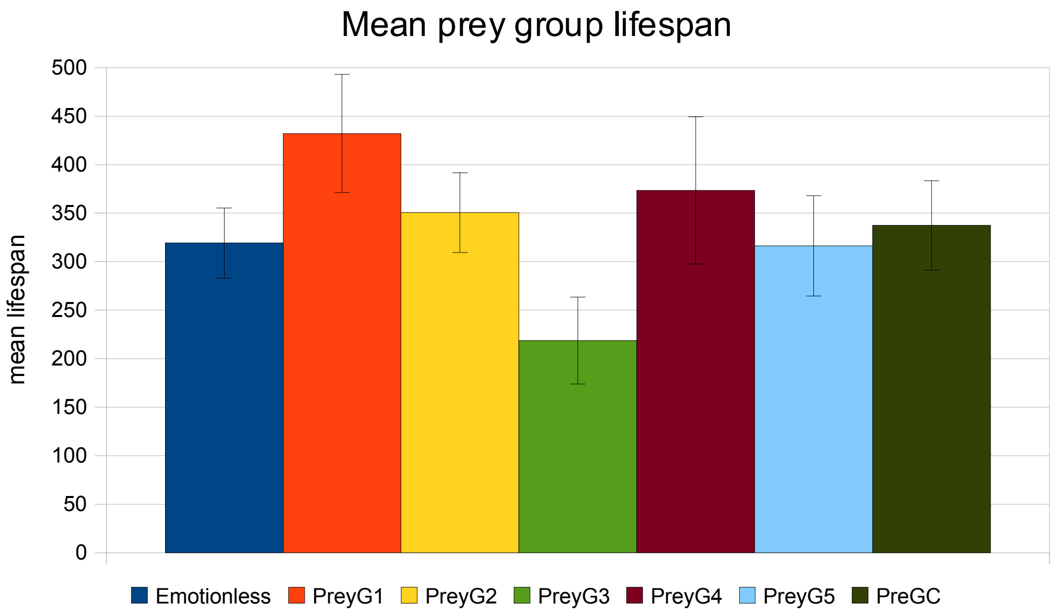

Figure 4.5: The mean lifespans of the five emotion prey tested, of the composite group containing copies of each of these five, and of the emotionless group.

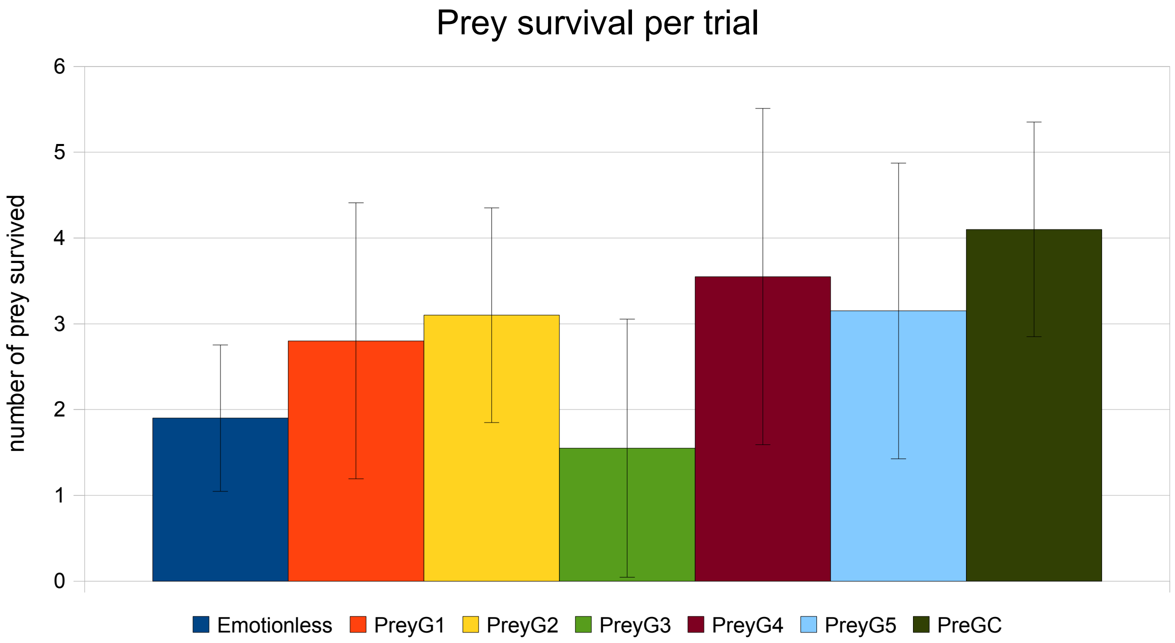

Figure 4.6: The mean number of prey per group that survived for the entire trial for each prey group tested.

Figure 4.7: This graph shows the pattern of action selection over time across all the individuals in one of the PreyG2 trials. Alarms were contagious for the PreyG2 agents and provoked fear responses (alarming and avoiding) throughout the group. Search was not utilized by these agents.

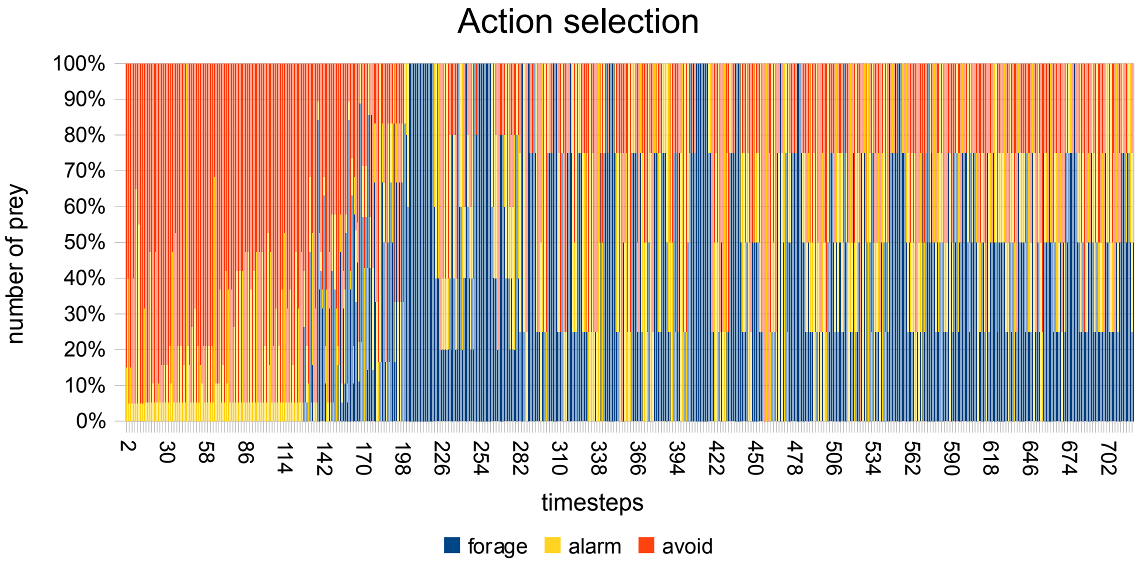

Figure 4.8: This graph shows the pattern of action selection over time across all individuals in one of the PreyG3 trials. Fear was highly contagious and again, these agents did not use *search*.

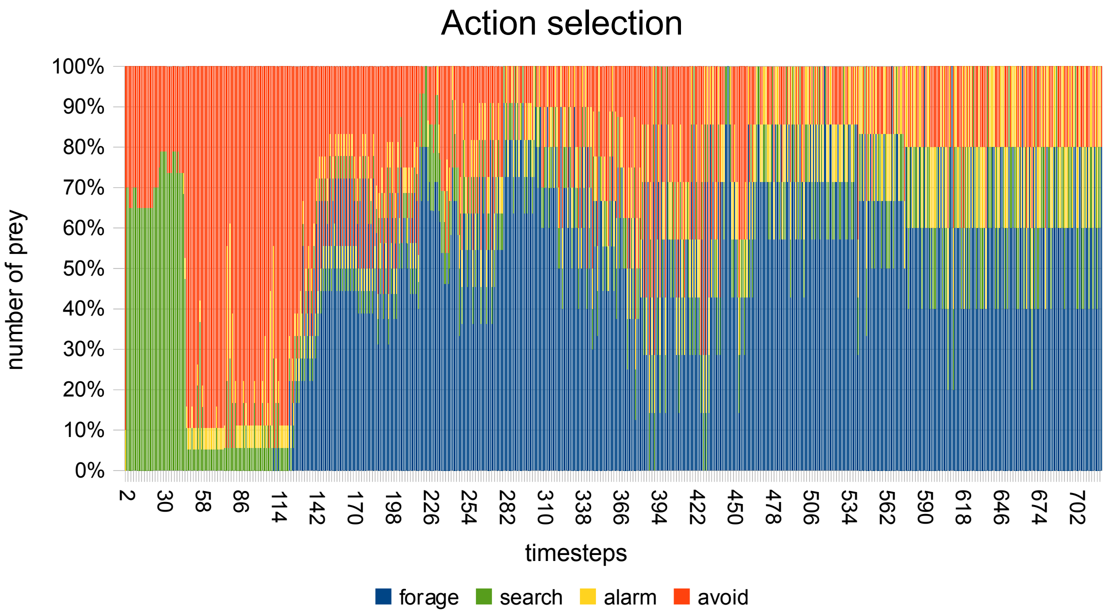

Figure 4.9: This graph shows the pattern of action selection over time across all individuals in one of the PreyG5 trials. Alarms tended to provoke avoid behavior rather than further alarms.