

# Modeling the Effects of Mass and Age Variation in Wolves to Explore the Effects of Heterogeneity in Robot Team Composition

John D. Madden and Ronald C. Arkin, *Fellow, IEEE*

**Abstract**— No two wolves are the same. Some of the variations between individuals such as variation in mass and age have major implications on their hunting abilities. The predatory success of a wolf pack has been found to depend on its composition of individuals along these dimensions. Building from a system for simulating wolf hunting behavior in earlier work [1], this paper explores the effects of heterogeneity among individuals in a team and the utility of these variations for the team as a whole. The implications for robot team control and organization are presented.

## I. INTRODUCTION

The Office of Naval Research has taken interest in studying heterogeneous unmanned network teams (HUNT) of robots, which has led to a project that attempts to draw ideas and models through the study of biological systems. This project has explored many species as diverse as dolphins [2] and prairie chickens [3] and this is the second paper from our group that explores the hunting behavior of the gray wolf as a model for robotic teams [1]. The previous work demonstrated that wolf pack hunting behavior could be simulated with relatively high fidelity to what is observed in nature by using a system of releasers and a weighted roulette wheel of probabilities using a finite state acceptor (FSA) plan representation. This work diverged from other work using wolf packs as models for robot behaviors, (e.g., Weitzenfeld et al. [4]), as it did not assume tight structures or roles for individuals (e.g., alpha male) but rather followed a specific and contemporary biological model [5,6]. Specifically, the model utilized in our earlier work [1] assumed that wolves employed no obvious pattern of coordinated hunting behavior, drawn directly from ethological observations of wolves hunting in Yellowstone National Park (YNP) [5].

Our previous work on wolf hunting behavior assumed complete homogeneity among the simulated wolves, both physically and behaviorally. In nature, no two wolves are the same. Individuals vary on dimensions such as age, mass, sex, experience, and personality. This variation is outwardly apparent from the different coloring of wolves even from the same litter (Figure 1). Any given wolf pack may vary in its structure on any of these dimensions; namely, the age structure could vary from one extreme of a pack of primarily older wolves to the other of a pack of wolves yet to reach

their prime, and any combination in between. In two separate projects, MacNulty et al. explored the affects of mass and aging on predatory performance and the success of wolf packs [7]. Here, these findings are used to explore the effects of simulated aging and difference in mass of robots on the overall success of the robot team.



Figure 1. Dark and light wolves approaching American bison [16]. Differences in color provide visible evidence of variability in wolves.

## II. WOLF BEHAVIOR AND HETEROGENEITY

Details about a generic wolf's physical properties and particulars about the hunt were presented in [1] as were the mechanics of the hunt itself. In this paper, the focus is on the effects of heterogeneity among individual wolves and thus only the basics needed to understand the discussion are reprised from [1]. For the discussion below, it may be assumed that the location is YNP and the prey animal or animals are all elk or herds of elk. Data used for this project were specifically gathered by biologists situated in YNP, as evidenced from the interactions of wolves and elk.

### A. Wolf Hunting Behavior

Wolves were chosen as a species of interest for the HUNT project due to their success as a large predator. They have found considerable success across the northern hemisphere in a variety of dissimilar ecosystems. The robustness of this species appears to be due to their having a generalized physical structure and hunting behavior rather than being tailored to specific purposes. Notably, the hunting behavior of wolves does not appear to employ a well-structured set of

This work was supported in part by the Office of Naval Research under MURI Grant # N00014-08-1-0696.

J.D. Madden and R.C. Arkin are with the Mobile Robot Laboratory, College of Computing Georgia Institute of Technology, 85 5<sup>th</sup> ST NW, Atlanta, GA, 30332. email {jmadden,arkin}@gatech.edu

strategies but rather generalized ‘rules of thumb’ that are used to react to the prey’s escape behavior [5].

The wolf hunt can be reduced into an ethogram of five foraging states: search, approach, attack group, attack individual, and capture [1] (Table 1). In a prototypical progression of a hunt, the wolf pack transitions through these states as follows: The wolf begins in the search state; in this state it travels somewhat at random, searching for prey. Upon discovery of prey, the wolves transition to the approach state. In this state the wolf pack orients itself directly toward the prey and moves toward them at moderate speed, not yet running. This allows the wolves to close their distance with the prey before the chase begins. As soon as the first prey animal spooks and begins to run away, the wolf pack transitions to attack group, assuming multiple prey animals are present, and begins the chase. The prey will usually split into groups running in all directions to generate confusion and force the wolves to choose which animals to follow. At this point in the hunt the wolves are scanning the herd for weak individuals to single out. Once such an individual has been selected, the wolf discovering it transitions to the attack individual state which is characterized by intensified pursuit and greater focus on the targeted prey individual. Other wolves may see this pursuit and join in, but this is not necessarily the case as explained later. The goal of this behavioral state is for the wolf to get close enough to the prey to begin biting it. Whether it is a single wolf or many, biting the prey signifies a transition to the capture state. The ultimate goal of the capture state is killing the prey. If the wolf is successful in this last state, the result is a kill; otherwise, the hunt may end in failure. This typical progression of a wolf hunt is presented in Figure 2.

TABLE 1  
Foraging States for Wolf Hunting Behavior

Foraging State	Description
<b>Search (S)</b>	Moving throughout the environment, searching for prey
<b>Approach (A)</b>	Closing distance on prey
<b>Attack Group (G)</b>	Scanning the herd for weak individuals
<b>Attack Individual (I)</b>	Focused pursuit of weak individual targeted for potential kill
<b>Capture (C)</b>	Grabbing and killing prey

It should be observed that these transitions do not necessarily occur in this order during a particular hunt, and many other transitions are possible between these foraging states, and the progression is often not linear and straightforward. In most cases, wolves primarily attacked groups after approaching, but “they also sometimes attacked elk groups immediately after discovering or watching the group” [6]. MacNulty et al. compiled statistical observational data regarding wolf hunt state transitions (Table 2) where the tabular values represent the probability of transition between states. Note that the transitions chosen for the description of the linear hunt above match those of highest probability in the table.

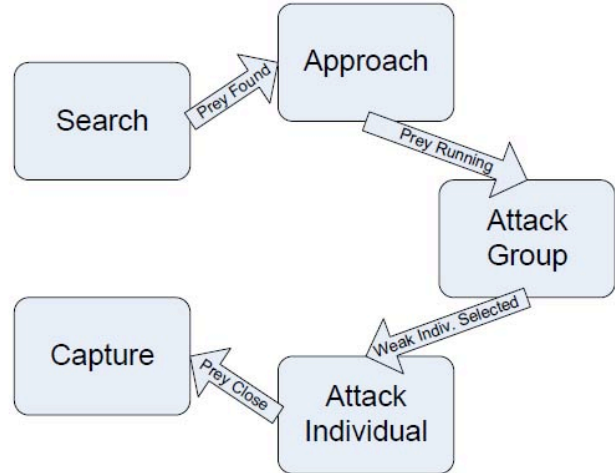


Figure 2. The progression of transitions between states seen in a prototypical successful hunt with wolves and elk.

TABLE 2  
Probabilities of Transitions Between States (From [6])

Preceding State	Following State					
	Search	Approach	Watch	Attack Group	Attack Individual	Capture
<b>Search</b>	.00	.68	.00	.31	.01	.00
<b>Approach</b>	.09	.00	.12	.69	.09	.01
<b>Watch</b>	.32	.35	.00	.27	.06	.00
<b>Attack Group</b>	.24	.09	.03	.13	.51	.00
<b>Attack Individual</b>	.16	.06	.02	.16	.08	.52

As mentioned earlier, when a wolf transitions to attack individual, other wolves may not necessarily join in the attack. Over two thousand hours of observed wolf behavior in Yellowstone Park seemed to prove that wolves do not coordinate their attacks, but rather act as individuals. The apparent coordination of their hunts is a result of byproduct mutualism [8], where multiple agents working toward the same goal at the same time appear to be helping each other and may in fact benefit from each other but not as a result of intentional assistance. Due to this, individual wolves progress through the foraging states on their own, where one wolf’s transition does not cause the transition of others, and disparity between states of wolves in the same hunt can be as far as one being in the Capture state while others are still in Search. The notion of an alpha wolf has been largely discredited [9].

### B. Heterogeneity in Age and Mass

As stated earlier, no two wolves are identical. Wolves differ over a number of dimensions such as age, mass, sex, experience, and personality, among others. The characteristics of these properties can either help or hinder a wolf’s hunting performance. Further, as the wolf hunt is broken down into a series of transitions between foraging

states, these different hunting tasks place different demands on a wolf's individual capabilities and yield different factors defining success or failure. As such, a variation in one of the above properties could cause a wolf to be weak at one hunting task while the same variation causes it to be strong for another.

Exit from each hunting task corresponds to a state transition. The transition from search is omitted in our design and implementation is omitted as we assume that individual variation will not cause a healthy wolf to miss the overall detection of prey in the presence of other wolves. The remaining hunting task transitions are labeled as follows: attacking ( $A \rightarrow G$ ), selecting ( $G \rightarrow I$ ), and killing ( $I \rightarrow C$ ).

Considering the possible dimensions of variation among wolves, this research explores the effects of variation with respect to age and mass only. Experience and personality appear to have an effect on a wolf's performance but they are difficult to quantify and to our knowledge no data exists to date on the effects of either of these. In fact, the data used for capturing the variations in age and mass is new [7,10] and represents some of the relatively few studies conducted to measure the effects of these variations on the predatory performance of large carnivores in general. Further, the gender of the animal is not considered as differences in performance based on sex appear to be completely explained by differences in mass between males and females [10]. Thus, only variation in mass and age of males is considered.

The prevailing view on variation in mass for large predators is simply that bigger is better [11]. A recent study challenged this view with the hypothesis that increasing predator size hinders foraging ability, stating that the bigger-is-better view did not consider the selecting task of the hunt, only the killing task [10]. Although the newer hypothesis was ultimately disproved in favor of the prevailing view [10], it did point out that heavier individuals showed a decrease in performance of the selecting task with an increase in size. This is due to the selecting task's demand on an individual's speed and agility which are hindered with increased mass. The heavier individuals had a slight advantage over their lighter counterparts in the task of attacking and a marked advantage in killing as their greater mass was essential for bringing down large prey. The improved performance in these other two tasks gave larger individuals the edge with an overall improvement in performance [10].

Age is also an important factor in a wolf's performance as a hunter and for the overall success of a wolf pack. In fact, the age structure of a wolf pack (i.e., the ratio of older to younger wolves) was found to be a top predictor of kill rate [7]. As with mass, variation in age has been found to affect wolves' performance differently for the different hunting tasks. However, according to studies by MacNulty et al. [7], for all tasks there was a peak age beyond which performance declined. The selecting task ( $G \rightarrow I$ ), which is considered the most strenuous task, is also the most sensitive to aging as it showed a sharp decline in performance for wolves aging beyond 3 years old. The peak ages for attacking and killing were found to be 1 year old and 2 year old, respectively,

with an aggregate peak age of 3 years old [7]. All wolves over this age are referred to as senescent while those between the ages of 2 and 3 will be referred to as peak wolves.

Even more interesting than the effect of age on individual wolf performance is the strong correlation between the age structure of a pack of wolves and the predatory performance of that pack. The greater the number of senescent individuals, the lower the overall performance of the pack. This is an interesting observation given the interplay between mass and age in wolves. To estimate the mass of the wolves being observed, MacNulty et al. [10] used the chart in Figure 3 that was created from measurements taken on the same wolves in YNP used in their other studies (only for males). This chart shows that a wolf's mass grows as it ages to a certain point and then begins to decline. This would suggest that if mass were the sole determinant of a wolf's predatory performance, older wolves would have higher performance until age 4 years of age and then decline. As described earlier, this is not the case and age itself has an effect on both the performance of an individual and the pack as a whole.

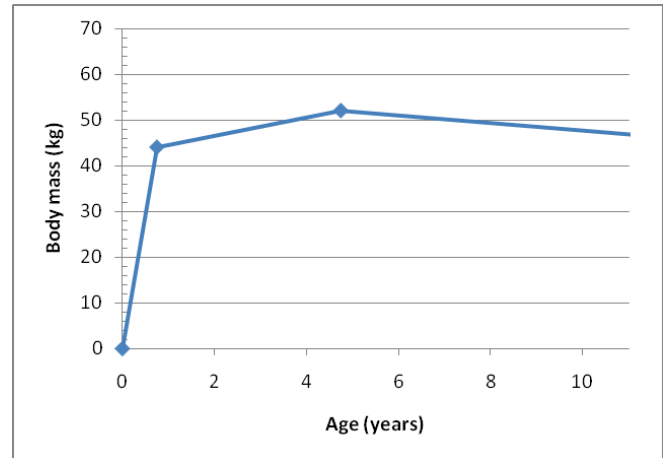


Figure 3. Trend line of observed body mass for male wolves versus age with population-averaged fitted value lines (derived from [10]).

### III. IMPLEMENTATION OF WOLF BEHAVIOR AND HETEROGENEITY

*MissionLab*<sup>1</sup>, a software package developed by the Mobile Robotics Laboratory at Georgia Tech [12,13], is used as the experimental infrastructure for the simulation results presented in this paper. *MissionLab* provides a graphical user interface where the user specifies behavioral states that control each robot's actions, and perceptual triggers that control the transitions between states, visualizing a finite state acceptor (FSA). The specific hunting behaviors described in this paper are compatible with *MissionLab*'s pre-existing behaviors, e.g., obstacle avoidance, wandering, or goal attraction. Combinations of behaviors (assemblages) can be created as part of the FSA, creating arbitrarily complex missions [14,15].

<sup>1</sup> *MissionLab* is freely available for research and educational purposes at: <http://www.cc.gatech.edu/ai/robot-lab/research/MissionLab/>



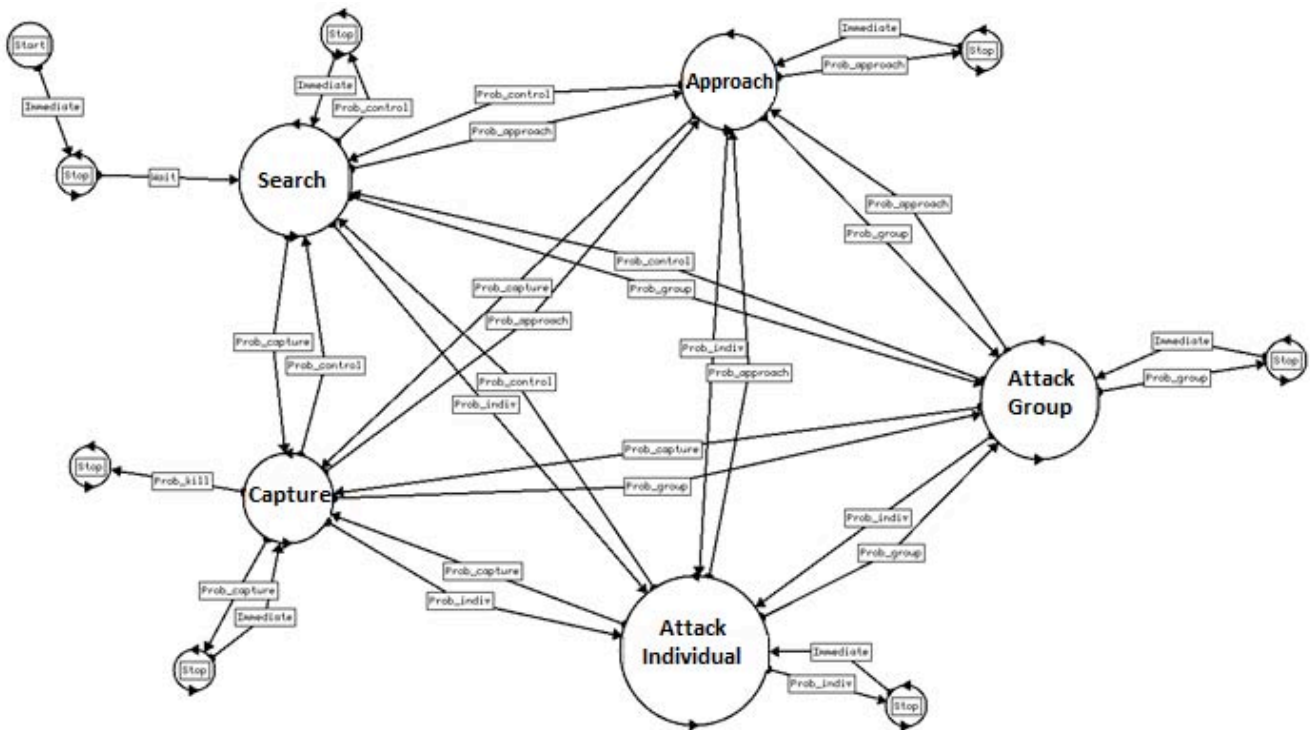


Figure 4. Finite State Acceptor for wolf hunting behavior with five foraging states (watch state removed) as well as initial and final states for experimentation purposes. The stop states connected to each foraging state are dummy states to facilitate the use of probabilistic triggers.

MissionLab’s use of FSAs facilitated the implementation of wolf hunting behavior as it was from the onset conceptualized as a set of states with various possible transitions. The FSA created in our earlier wolf behavior work [1] is depicted in Figure 4. Note that this FSA is a completely connected graph, although many transitions are not accessible to a wolf in a given situation. The availability of these transitions is handled by a system of releasers that depends on the presence of certain perceptual stimuli. For instance, the presence of prey is a releaser for the wolf to transition from search to any other state, as it would be impossible for the wolf to approach, attack, or capture prey without having detected a prey animal. Thus if this releaser is satisfied, transitions from S to A, G, I, C may be possible, otherwise they are not. Further, although this is the only releaser needed for  $S \rightarrow A$ , other transitions such as  $S \rightarrow G$  also require additional releasers to be present, namely, ‘prey running’ for  $S \rightarrow G$ .

In this manner, the releasers permit some transitions to occur and block others, but what happens when multiple transitions are simultaneously available? This is where the probability table related earlier becomes necessary. To decide which transition to take, a weighted roulette wheel is employed using the probabilities in Table 2 as the weights for all possible transitions. The details of this system and how it uses probabilities to trigger different state transitions is explained further in [1]. For the current discussion, it suffices to state that altering the probability table used in the weighted roulette wheel will alter the probability of a robot manifesting a specific behavior in a given situation. Heterogeneity was thus introduced to the robots by altering

the set of input probabilities for each individual.

It was decided that the effects of heterogeneity would best be explored by choosing a few wolf exemplars possessing the greatest variance in mass and age and then comparing the performance of each. The exemplars chosen were a heavy senescent male (60 kg, 5 years old) and a light young male (38 kg, 3 years old). The probability of transitions for these exemplar wolves were generated by the product of the original properties in [6] and a multiplier. The multiplier was created as the product of the ratio of successful transitions for the exemplar to that of the average wolf. Thus, a wolf of average mass and age would have a multiplier of 1.0 for all transitions, while the senescent male (5yo, 60kg) had a multiplier of 0.48 for selecting ( $G \rightarrow I$ ). Comparing transition success for an average wolf to that of the exemplars was facilitated from data obtained from Figure 2 in [10] and Figure 1 in [7]. Tables for these two classes of individuals are given in Table 3.

Although the elk being preyed upon may have defensive behavioral strategies and coordination, the focus of the research to date has been on the hunting behavior of wolves. Therefore, the elk’s behavior was greatly simplified, i.e., their reaction to approaching wolves is simply either stopping or running away in a direction opposite to the wolves’ approach. The elk’s state before the approach of wolves was varied to simulate situations where the elk either are initially stationary, moving back and forth between multiple grazing areas, or wandering around. Figure 5 illustrates an FSA representing the elk behavior for moving back and forth between two grazing areas, including a simple modification that switches between the elk stopping upon seeing wolves, or running away.

TABLE 3  
Probabilities of Transitions Between States For Heavy Senescent Male (60 kg, 5 years old)

Preceding State	Following State				
	Search	Approach	Attack Group	Attack Individual	Capture
Search	.00	.66	.33	.01	.00
Approach	.07	.00	.74	.06	.01
Attack Group	.36	.14	.20	.25	.00
Attack Individual	.10	.03	.10	.03	.74

Probabilities of Transitions Between States For Light Peak Male (38 kg, 3 years old)

Preceding State	Following State				
	Search	Approach	Attack Group	Attack Individual	Capture
Search	.00	.70	.29	.01	.00
Approach	.10	.00	.64	.10	.01
Attack Group	.21	.08	.11	.60	.00
Attack Individual	.21	.08	.21	.10	.40

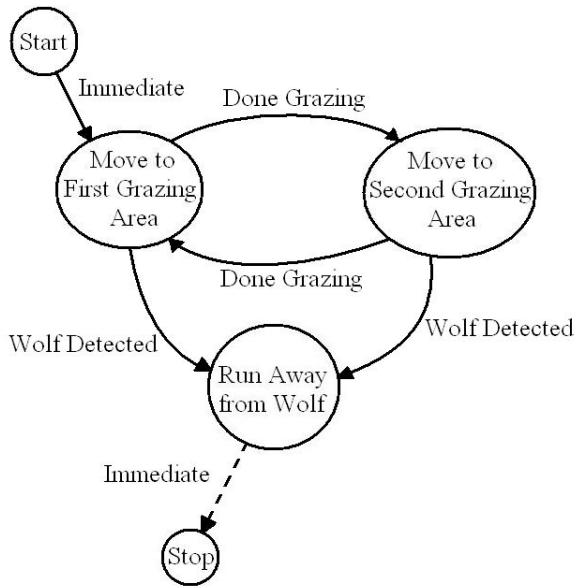


Figure 5. A Finite State Acceptor representing simplified Elk behavior. If the 'run away' behavior is desired, the dotted trigger to the final stop state is removed.

The quantitative results in this paper were obtained entirely using simulated robots; however, the wolf hunting behavior has been successfully implemented on the pioneer platform and demonstrated in our laboratory. To date, up to four real robot pioneers have been used to emulate hunting scenarios involving two wolves and two elk. The physical experiments displayed behavior that qualitatively matched the behavior of the simulated robots, scaled down and better suited to an indoor lab environment, rather than a large outdoor area. An example of such an experiment using the pioneer robots appears in Figure 6.




Figure 6a. This sequence of photos provides an example of a wolf hunt on the pioneer platform. This photo shows the starting location for a hunt.




Figure 6b. The elk robots initially head toward the first of two 'grazing areas' which they will move back and forth between until attacked by the wolves.




Figure 6c. As the elk robots head toward the second grazing area, the wolves attack.




Figure 6d. The elk robots run away from the wolves in an attempt to evade them. The wolves are both focused on the elk robot on the right.




Figure 6e. The elk robots ‘run’ in different directions and the wolves split up to chase both. The hunt ends with the wolf on the right ‘killing’ the elk it was chasing.

#### IV. SIMULATION RESULTS

Two phases of experiments were run, the first using two wolves and three elk while the second used six wolves and eleven elk. The goal of the first phase was to compare individual differences between the peak and senescent wolves whereas the goal of the second phase was to find an optimal pack composition.

The first phase had three scenarios: two peak wolves, two senescent wolves, and one of each with three elk included in each of these scenarios. The success of each scenario was defined by whether or not the wolves could make a successful kill in a reasonable amount of time. This time could be taken as the time to exhaustion for real wolves and was determined to be the time it took for elk to reach the far boundary of the *MissionLab* map. Beyond this time the success of the hunt would be drastically affected by the elk’s escape being thwarted by the boundaries of the map. Indeed, this boundary effect was seen to cause difficulties in the experiments of our previous work on wolf behavior [1]. Twenty runs were conducted for each scenario. The results of these experiments are shown in Table 4 below. Note that the pair of peak wolves performed better than the pair of senescent wolves and better than the mixed team. Further, in the runs with the mixed team, the peak wolf made 70% of the kills for that team.

TABLE 4  
Results From Phase One Simulations

Pack	Success
Scenario 1: All heavy (60kg), senescent (5yo)	30%
Scenario 2: All light (38kg), peak (3yo)	45%
Scenario 3: Half senescent, half peak	40%
Kills by peak wolf	75%

The second phase of experiments was conducted with six wolves and eleven elk using five different compositions of wolf packs: all heavy, senescent wolves (pack A), all light, peak wolves (pack B), half peak, half senescent (pack C), four senescent, two peak (pack D), and four peak, two senescent (pack E). Due to the increased number of robots in these experiments, the success or failure of the pack to make

a kill before a given exhaustion time was no longer a good metric by which to compare the predatory performance of packs. Instead, the map was made large enough to effectively be infinite and the experiments ran until a kill was made. Thus, the success of each pack was compared by means of the time to kill, taken to be from when the wolves began searching for the elk until any one of the elk was killed. Twenty runs were conducted for each of these packs. The results are given below in Table 5. If the order of the packs are rearranged to be in order of increasing numbers of peak individuals (decreasing number of senescent wolves), then the chart in Figure 7 to the right shows a chart of time to kill with an optimal performance observed with a composition of four peak wolves and two senescent wolves. Screenshots from exemplar experimental runs are also given in Figures 8 through 11 below where blue circles represent wolves, red circles represent elk, and the black lines trace the movements of both.

TABLE 5  
Results From Phase Two Simulations

Pack	Time to Kill (Average)
Pack A: All heavy (60kg), senescent (5yo)	1:06.6
Pack B: All light (38kg), peak (3yo)	1:00.6
Pack C: Half senescent, half peak	1:02.4
Pack D: Four senescent, two peak	1:03.5
Pack E: Four peak, two senescent	0:58.3

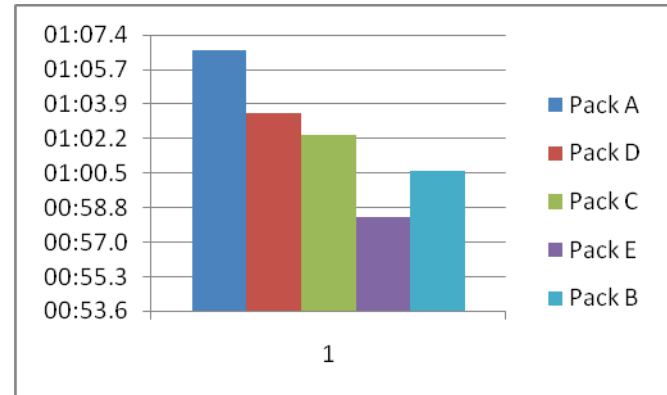


Figure 7. Comparison of time to kill for all five pack compositions

#### V. DISCUSSION

The first set of experiments demonstrated a general advantage of the younger, lighter wolves over their more mature counterparts. The younger group was 15% more likely to make a successful kill on any given run, a 50% improvement over the senescent wolves. From the transition probability tables for each class it is apparent that the largest disparity between the two groups involves the much higher probability for the peak wolves to complete the  $G \rightarrow I$  transition: 60% compared to 25%, and a much higher probability for the senescent wolves to complete the  $I \rightarrow C$  transition: 74% compared to 40%. The magnitude of these differences is nearly the same so it would be reasonable to assume that the two classes of wolves should perform

somewhat equally if these transitions are equally important. Thus, given the disparity in performance, it appears the two transitions are not equal in importance; namely, it appears the  $G \rightarrow I$  transition is more critical than  $I \rightarrow C$ . These two transitions correspond to the selecting and killing tasks, respectively, and the disparity in their importance has been noted in nature. The selecting task has been considered the most demanding and critical of the hunting tasks in a wolf hunt [7]. In nature, this is due to the speed and agility required of the wolves; in the simulations presented here, all robots essentially have the same speed and agility and thus should not be affected by this fact. The source of this disparity is then, rather, due to the parameters of the simulation detracting from the importance of the killing task. This task has less of an effect on the overall performance of the wolves because of two reasons: killing in simulation simply involved a successful transition based on probability rather than the often complicated dynamics of multiple wolves simultaneously biting at the prey animal. This simplification was necessary for the simulation as modeling the mechanics of the final stages of a wolf hunt, namely taking down the prey, was beyond the scope of this project. Second, the simulated wolves did not need to concern themselves with their own safety. In nature, the guiding principle behind a wolf's hunting behavior is to kill without being killed. More importantly, in the simulation the robot wolves do not need to concern themselves with this. These two differences between simulated and real world hunts led to the devaluation of the senescent wolves' advantage in the final transition of the hunt. Nonetheless the value of heterogeneity in robot packs is in evidence, but this research should not be construed as an attempted validation of the underlying biological model but rather as having a goal to provide guidelines for the design of heterogeneous robot teams.

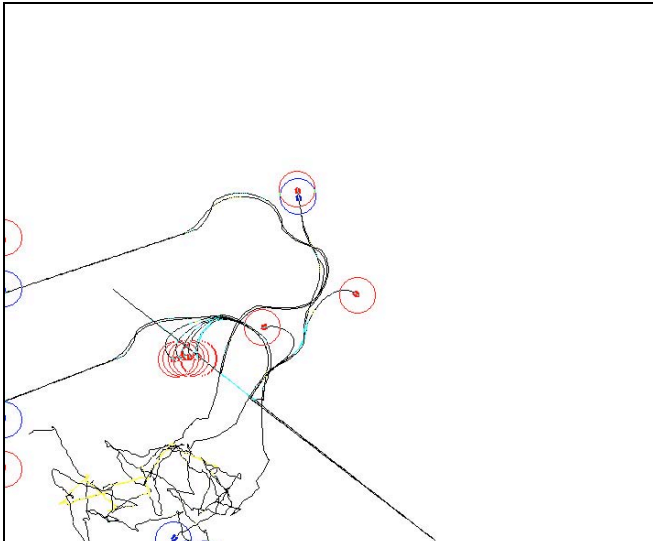


Figure 8. By increasing the number of simulated wolves and elk to six and eleven, respectively, the added chaos of the many simultaneous interactions required the expansion of the map to a virtually infinite space. The straight line is the initial migration of the prey between two grazing areas before the wolf encounter.

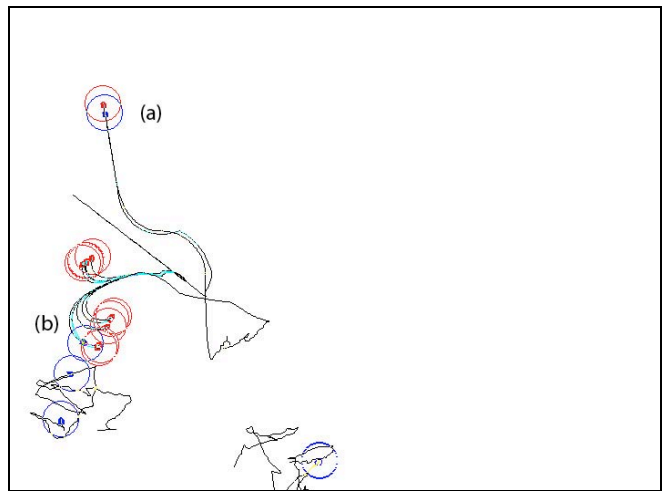


Figure 9. Example of run with pack E. This shows a peak wolf off in pursuit of an elk on its own towards the top (a) and another driving elk back toward senescent wolves in the middle (b).

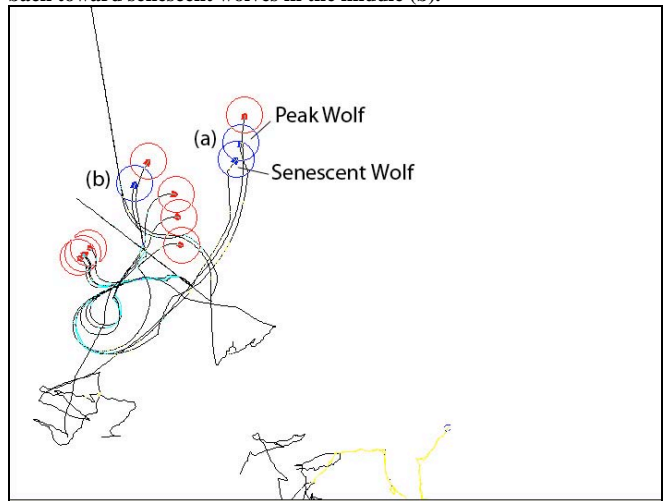


Figure 10. Continuing from the previous figure, a senescent wolf has joined a peak wolf near the top (a) in pursuit of a singled out elk while another peak wolf has begun pursuit of another elk to the left (b).

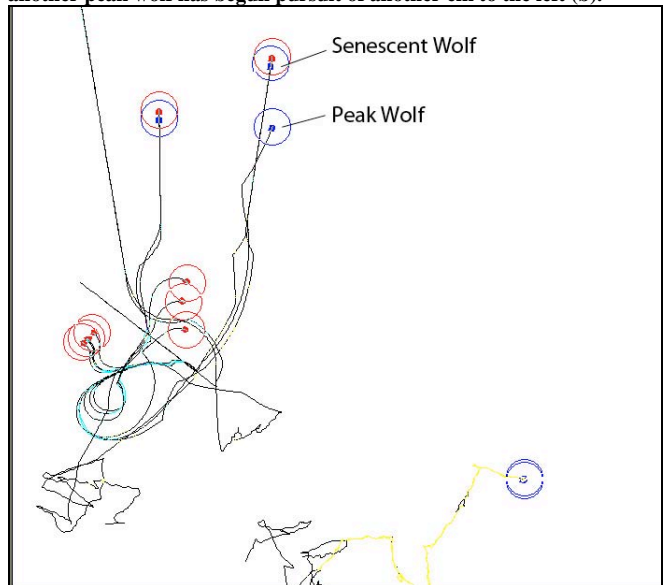


Figure 11. The peak wolf that had originally been leading the pursuit of the elk closest to the top failed, the senescent wolf that joined him is now leading the pursuit of that elk.



The results from the first phase of the experiment are enough to demonstrate the overall advantage of the peak wolves in the simulation environment. The second phase was intended to compare the effects of different compositions of wolf packs. The results show that packs with more peak wolves generally outperformed those with greater numbers of senescent wolves; however, pack E indicated an anomaly. The results show that this pack composition with four peak and two senescent wolves provided the best performance of the pack compositions evaluated. This was not entirely unexpected as there are a number of possible explanations for this. One is that the improved selecting task performance of the lighter peak wolves allowed them to succeed in situations where this task was difficult (i.e., pursuit of tightly grouped prey) and the improved killing performance of senescent wolves allowed them to quickly take advantage of opportunities for quick kills (i.e., the prey is driven to one wolf by another wolf). If the pack were entirely composed of one class of wolf it would only be able to take advantage of one of these situations, but not both. The superior performance of pack E over the other mixed packs (pack C and D) could possibly be explained by the greater prevalence of the former situation over the later. This could mean that pack E was able to take advantage of some relatively rare, yet greatly beneficial situations that pack B could not. Indeed, when the situation of senescent wolves making quick kills of elk driven towards them by other wolves did occur, the result was usually a quick kill which greatly helped drop the average time to kill for pack E. As this driving act by one wolf to another is believed unintentional, it exploits the notion of byproduct mutualism [8] discussed earlier.

Another, more general explanation is that having two different classes of wolves allows the pack to take advantage of the benefits of each. This would also explain why pack E outperformed pack B. Potentially pack E outperformed the other heterogeneous packs simply because it had more peak wolves. Nonetheless, more experimentation is required, over a broader range of pack sizes, to confirm this as a more general result.

## VI. SUMMARY AND CONCLUSIONS

Using a software implementation of wolf hunting behavior developed in previous work [1], the effects of heterogeneity among agents operating as a team was observed. The results demonstrated the expected result that simulated wolves configured to represent light, peak individuals would outperform their heavy, senescent counterparts. However, the results also demonstrated the somewhat unexpected result that a mixed team could outperform a team comprised of purely the superior class of peak wolf under certain circumstances. These results both demonstrate the relative fidelity of the system with what has been observed in nature as well as provide insight to purely hypothetical comparisons (e.g., packs of purely peak wolves do not exist). The results also demonstrate that by incrementally varying the composition of a heterogeneous

team, a high performing composition can be found without understanding the exact reasoning behind the underlying optimality. This has significant implications for the design choices for teams of robots operating in unstructured and dynamic environments.

## ACKNOWLEDGMENTS

The authors would like to thank Prof. Dan MacNulty for sharing his considerable insights and knowledge on wolf behavior for this project.

## REFERENCES

- [1] J. D. Madden, R.C. Arkin, MacNulty, D., "Multi-robot System Based on Model of Wolf Hunting Behavior to Emulate Wolf and Elk Interactions", *Proc. IEEE International Conference on Robotics and Biomimetics (ROBIO 2010)*, Tianjin, China, Dec. 2010.
- [2] A. Rahmani, M. Haque, and M. Egerstedt, "Biologically Inspired Coalition Formation of Multi-Agent Systems", *International Conference on Autonomous Agents and Multiagent Systems*, Toronto, Canada, May 2010.
- [3] B. A. Duncan, P. D. Ulam, R.C. Arkin, "Lek Behavior as a Model for Multi-Robot Systems", *Proc. IEEE International Conference on Robotics and Biomimetics (ROBIO 2009)*, Guilin, China, Dec. 2009.
- [4] A. Weitzenfeld, A. Vallesa, H. Flores "A Biologically-Inspired Wolf Pack Multiple Robot Hunting Model," in *Latin American Robotics Symposium and Intelligent Robotic Meeting (LARS 2006)*, Santiago, Chile, 26-27 Oct. 2006.
- [5] D.R. MacNulty, *Presentation and Interviews*, Atlanta, Georgia, May 6-7, 2010.
- [6] D.R. MacNulty, L.D. Mech, D.W. Smith, "A Proposed Ethogram of Large-carnivore Predatory Behavior, Exemplified by the Wolf," *Journal of Mammalogy*, Vol. 88, No. 3, pp. 595-605, June 2007.
- [7] D.R. MacNulty, D.W. Smith, J.A. Vucetich, L.D. Mech, D.R. Stahler, and C. Packer, "Predatory Senescence in Aging Wolves," *Ecology Letters*, 2009, 12: 1347-1356.
- [8] L., Dugatkin. *Cooperation among Animals: An Evolutionary Perspective*, Oxford University Press, 1997.
- [9] Wikipedia, "Pack (canine)", accessed May 12, 2011, [http://en.wikipedia.org/wiki/Pack\\_\(canine\)](http://en.wikipedia.org/wiki/Pack_(canine)).
- [10] D.R. MacNulty, D.W. Smith, L.D. Mech, L.E. Eberly, "Body Size and Predatory Performance in Wolves: is Bigger Better?" *Journal of Animal Ecology*, 2009, 78, 532-539.
- [11] A. Herrel, A.C. Gibb, "Ontogeny of Performance in Vertebrates," *Physiological and Biochemical Zoology*, 79, 1-6.
- [12] Georgia Tech Mobile Robotics Laboratory, "Manual for *MissionLab*," Version 7.0, 2007.
- [13] D. MacKenzie, R.C. Arkin, J. Cameron, "Multiagent Mission Specification and Execution," *Autonomous Robots*, Vol. 4, No. 1, Jan. 1997, pp. 29-57.
- [14] R.C. Arkin, *Behavior-based Robotics*, MIT Press, 1998.
- [15] R.C. Arkin, "Motor Schema-Based Mobile Robot Navigation," *International Journal of Robotics Research*, Vol. 8, No. 4, August 1989, pp. 92-112.
- [16] D. Smith, "Mollies Pack Wolves Baiting a Bison," National Park Service, [http://en.wikipedia.org/wiki/File:Canis\\_lupus\\_pack\\_surrounding\\_Bison.jpg](http://en.wikipedia.org/wiki/File:Canis_lupus_pack_surrounding_Bison.jpg)