Multi-robot System Based on Model of Wolf Hunting Behavior to Emulate Wolf and Elk Interactions

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

***Abstract*— Wolves are one of the most successful large predators on earth. Their success is made apparent by their presence in most northern ecosystems. They owe much of this success to their generalized hunting behavior which allows them to quickly and effectively adjust to different species of prey. The success of this hunting behavior for wolves is the inspiration for a project to bestow this behavior onto a system of robots with the hopes that they might utilize the apparent strengths of the behavior to achieve their own success.**

I. Introduction

As part of a project for the Office of Naval Research, models of behavior from biology are being used to develop heterogeneous unmanned network teams (HUNT) of robots. An earlier study in this project used lekking behavior from prairie chickens to develop a basis for structuring groups of robots [1]. The results of the lekking study were used as a starting point for the current study with wolves. Our group includes a biologist (D.R. MacNulty), who specializes in wolf behavior and has conducted extensive studies of wolves in Yellowstone National Park (YNP). The model of wolf behavior used for this project was based on observations from these studies.

Figure 1. Elk are the primary food source of wolves in Yellowstone National Park and wolves hunting elk are the focus of this project. [2]

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 5th ST NW, Atlanta, GA, 30332. email {jmadden,arkin}@gatech.edu

D.R. MacNulty is with the Department of Ecology, Evolution, & Behavior University of Minnesota, 1987 Upper Buford Circle, St. Paul, MN, 55108. email: macn0007@umn.edu

This is not the first project to use wolf behavior as a model for robots. Weitzenfeld et al. created packs of robot wolves where alpha wolves would lead and beta wolves would follow [3]. Our project breaks from this work in two significant ways. First, Weitzenfeld assumed a tight structure to exist in the coordination of wolf packs; however, direct observations of wolves hunting elk in YNP indicate no obvious pattern of coordinated hunting behavior [4]. For this reason, our wolf model has been given no hard constraints to keep them together. Second, Weitzenfeld also assumed roles such as alpha wolf and beta wolves, to control the position and actions of each individual throughout a hunt. The observations on which our models are built show that roles do exist but they may change in an ad hoc manner and are based on physical abilities and not on a pre-existing dominance hierarchy. The field observations, from which this understanding of wolf hunting behavior is based, indicate that there is a lack of explicit coordination between the wolves. Their group behavior is evidently not a well structured set of strategies but rather generalized 'rules of thumb' that are used to react to the prey's escape behavior in order to minimize the risk of injury to themselves [4].

II. Wolf Behavior

A. Individual Properties

Wolves are able to consume a variety of prey – from mice to moose – because of their generalized skull morphology. And the apparent lack of coordination could be an advantage in that it allows wolves to hunt over a range of conditions irrespective of any requirement to coordinate. They use a few basic heuristics ('rules of thumb'), e.g., attack while minimizing the risk of injury with no overall hard behavioral constraints on actions [4]. This makes their behavior very flexible and allows them to quickly and easily make the transition between different species of prey, such as elk in the summer and bison in winter when elk migrate. One observation in Yellowstone National Park, involved a pack of wolves that had been hunting bison, moved into a new valley, and immediately started hunting elk. This serves as a testament to the adaptability of wolf hunting behavior, and a powerful clue regarding their success in such varied environments.

B. *Breakdown of a Wolf Hunt*

As is the case for most large carnivores, the predatory behavior of wolves is composed of multiple phases of behavior or foraging states. Traditionally, only three states are considered: search, pursuit, and capture [5]. In this research, however, a modified ethogram with six states has been adopted: search, approach, watch, group attack, individual attack, and capture, as proposed in [6]. Here, MacNulty concluded that the additional states represent "functionally important behaviors", and for robotics, this more detailed ethogram lends itself more easily to software implementation. The wolf packs studied to form this ethogram were located in Yellowstone National Park, hunting elk and American Bison. The focus of the first phase of the robotics implementation involves a model of wolves hunting elk. The following is a description of a typical hunt with wolves and elk. A diagram showing the typical progression through foraging states is given below in Figure 2.

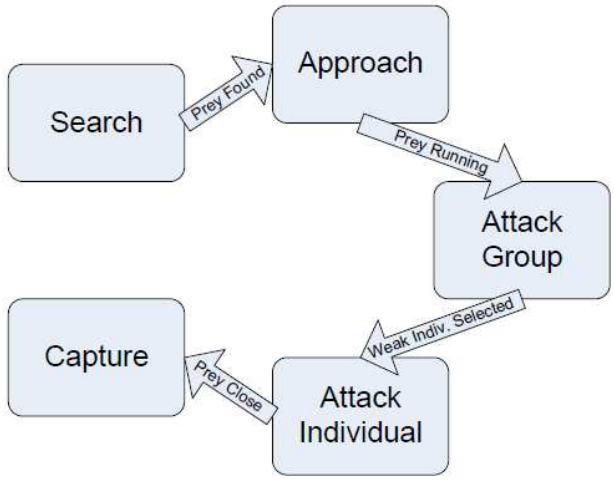

Figure 2. The progression of transitions between states seen in a typical hunt with wolves and elk.

TABLE 1
Foraging States for Wolf Hunting Behavior

Foraging State	Description
Search	Traveling without fixating on and moving toward prey
Approach	Fixating on and traveling toward prey
Attack Group	Running after a fleeing group or lunging at a standing group while glancing about at different group members (i.e., scanning)
Attack Individual	Running after or lunging at a solitary individual or a single member of a group while ignoring all other group members
Capture	Biting and restraining prey

When a hunt is initiated, the wolf pack heads out from its den or resting site and begins searching for prey. Hunger motivates the initiation of a hunt [8]. What direction the wolves go and to what extent they are willing to travel are dependent on their experience of prior successes and failures. As they search they make use of their strong senses, using the wide range of their lateral vision and their movable ears, to scan the landscape for potential prey. Once prey has been located, they start approaching.

Assuming that that the pack has located a relatively stationary herd of elk, the wolves approach at moderate speed. In general, wolves do not sneak up on their prey, nor do they target a specific individual from the herd until after the herd begins running. Species that use this approach strategy are known as cursorial predators and it is the principal difference separating their hunting behavior from that of other large predators such as lions [6]. In response to approaching wolves elk will either stand their ground or to run away. Elk most commonly run away which usually leads to the 'attack group' state.

As the prey quarry run away, they split up into groups headed in different directions and the wolves must also split up to follow as many as they can. During this stage of the hunt the wolves are scanning through the groups of prey, trying to locate the weakest individual that will provide the best opportunity for a kill. An advantage of running the animals to exhaustion is that it creates opportunities for the prey animals to make a fatal mistake (i.e., tripping). It also provides a useful test of performance by which the wolves can evaluate which animal is the weakest [8]. When a weak animal is detected by a wolf, that wolf then transitions to the 'attack individual' state.

The 'attack individual' state is characterized by intensified pursuit and greater focus on the targeted prey individual. Other wolves may see the pursuit of this wolf and join in, but that is not necessarily the case. Coordination of multiple wolves (or lack thereof) is discussed in the next section. The goal of this behavioral state is for the wolf to get close enough to the prey to begin biting it in an attempt to bring it down. Whether it is a single wolf or a number of wolves, biting the prey signifies a transition to the capture state.

The ultimate goal of the capture state is killing the prey. If the prey animal is small (i.e., a calf) the first wolf may attack the throat directly since it can easily handle the animal by itself. If the prey is larger and there are many wolves, they will often bite at the hind legs and rump attempting to slow their prey down before grabbing the neck. This project is not concerned with the mechanics of how wolves bring down prey but it is important to note that there are differences in attacking different prey. If the prey truly was a weak individual, the wolves will most likely complete a successful kill, but if they had misperceived a strong animal as weak, they may fail and either give up on the hunt or transition back to an earlier state.

The narrative of a hunt that has just been related gives a general idea of how many specific individual hunts progress through these foraging states; however, it is often not this clear cut. Many other transitions are possible aside from the

seemingly linear straightforward progression from search, to approach, then attack group, attack individual, and finally capture. For instance, 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 their statistical observational data of state transitions (Table 2) where the tabular values represent the probability of transition between states. Notice that the transitions chosen for the description of the linear hunt above are those of highest probability in the table.

TABLE 2
Probabilities of Transitions Between States From [6]

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

Thus far, the 'watch' state has been neglected as it is a rare state for wolves to enter when attacking elk; as seen in the table above, the highest probability of entering the 'watch' state is 12% from 'approach'. For this reason, the 'watch' state has been left out of the ethogram for our robotics implementation described in Section III.

C. *Coordination or Lack Thereof*

Wolves are generally perceived by the public to be highly coordinated hunters using strategies and teamwork to bring down large prey. Over two thousand hours of observed wolf behavior in Yellowstone Park seem to prove otherwise [4]. According to these observations, wolves not only show no signs of planned strategies but also little to no noticeable communication while hunting. This is evidenced by the fact that wolves hunting the same herd do not make transitions between states together (i.e., one may find a weak prey and transition to attack that individual while the others remain in an attack group). The disparity in these transitions goes so far as to see one wolf having killed an animal and begin eating it while the others persist in the 'attack group' state. Furthermore, in this last example, the wolf that made the kill did not appear to make any attempt to signal the others of its success.

The seemingly coordinated wolf hunting behavior is most likely the result of "byproduct mutualism" where each individual is simply trying to maximize its own utility. It is hypothesized that wolves see the fact that other wolves are chasing an elk as a sign of weakness of that prey animal and from that stimulus determine that they have the best chance of a meal if they join in the pursuit of that animal. Even far greater size of their prey does not force wolves to rely on teamwork; according to MacNulty, some aggressive wolves would attack even large bison alone. It is possible that such wolves simply assume the others will help them, or they are unaware that they need the others to help them take down the large prey because this is most often the case. It may not, however, be required that wolves need help to take down any of their usual prey. It is proposed that one of the biggest reasons that large terrestrial predators do not use group coordination is that they do not necessarily need it. Solitary hunters have a high success rate, roughly 21% for most large carnivores [MacNulty unpublished data].

III. Implementation of Wolf Behavior

Experiments for this project were conducted with simulated robots in *MissionLab*[1], a software package developed by the Mobile Robotics Laboratory at Georgia Tech [9,10]. *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, yielding a finite state acceptor (FSA). The behaviors created for the current project can be combined with other pre-existing behaviors such as obstacle avoidance, moving toward an object, or noise (random wandering). This allows for assemblages of behaviors to be created and connected in the FSA to create arbitrarily complex missions [11,12].

Reducing the overall hunting behavior of wolves into the five foraging states related earlier, facilitated implementation where each behavioral state represents a corresponding state for the robot. These states, together with a few others added for initial configuration and termination of experiments, were used to create the FSA shown in Figure 3. The perceptual conditions that must be met in order for a wolf to transition from one state to another are known as releasers. For instance, for the wolf to switch from the search state to the approach state the wolf must of necessity have found prey to approach; therefore, we say the presence of prey is a releaser to transition to the approach state. These are encoded as perceptual triggers in *MissionLab*. A list of the releasers used in this implementation and the transitions they facilitate are given in Table 3. The system of releasers would normally be enough to define the transitions in a *MissionLab* FSA except that often, multiple transitions are possible from the same state to many others. In nature, what decides which transition is chosen is a combination of situational factors such as the number of wolves in the pack, the number of prey individuals in the herd, terrain features, as well as the wolf's individual attributes such as age, weight, and personality (i.e., aggressive individuals are more likely to move more quickly toward capture). While these factors will be incorporated directly in later work, for now their affect was indirectly computed by using the probabilities of transitions of observed wolf behavior described earlier in Table 2.

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

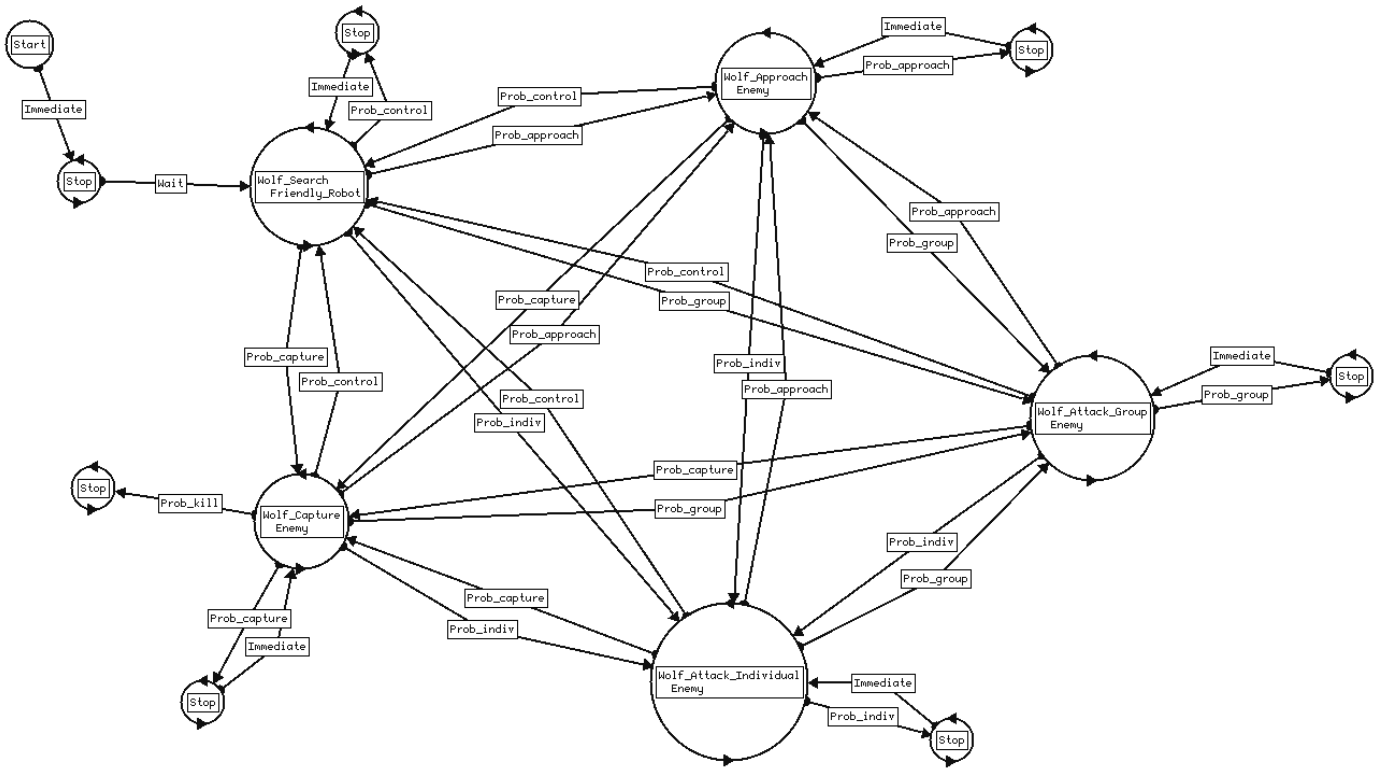

Figure 3. Finite State Acceptor for Wolf behavior with five foraging state (watch state removed) as well as initial and final states for experimentation purposes. The stop states connected to each foraging state are dummy states to facilitate the probabilistic trigger.

TABLE 3
List of Releasers and Transitions

Releaser	Transitions possible
Prey Found	S → A, G, I
Prey Lost	A, G, I, C → S
Multiple Prey	S, A, I, C → G
Prey Running	A → G, I
Prey Stopped	G, I → A ; A, I → C
Weak Individual Identified	G → I
Prey Close	G, I → C

Search, Approach, Attack Group, Attack Individual, and Capture abbreviated: S, A, G, I, C, respectively.

These probabilistic triggers for the FSA in Figure 3 were created in the following fashion: for each state there is a 'Control' trigger leading from that state to the search state and then another trigger leading to every other state to yield a complete graph encompassing all possible transitions between the 5 major behavioral states. The probabilities from Table 2 were entered as parameters into the 'Control' trigger at build time, and at run time this trigger would check which transitions had their releasers satisfied. A weighted roulette wheel was then created by normalizing the satisfied probabilities such that they added to one and then a random number generated between zero and one would decide the transition to take. Once a transition is selected, the condition for the corresponding trigger is satisfied and that behavioral transition occurs. If the transition selected was from any state to the search state, the control trigger's condition would be satisfied and a transition to search would occur. Also incorporated into the control trigger was a timer to force the wolves to stay in a state for a minimum length of time based on the observed average time wolves spent in each state (D.R. MacNulty, unpublished data). Without this hysteresis feature the wolves would constantly alternate back and forth between states for which the releasers are present (a form of behavioral dithering).

Each state in the diagram above is a combination of the constituent pre-existing behaviors: move to object, wander, and avoid obstacles. Move to object creates an attraction vector from the robot to the object selected. In the search state the selected object was friendly robots, which in this case represent other wolves, in all other states the selected object was elk. *MissionLab* uses built-in vector and simulation specific functions, implemented in C++ code. The move to object behavior creates a vector directed from the center of the robot to the center of the selected object so long as the selected object is within range to be detected.

The wander behavior creates vectors of random direction. While in the search state, this was useful to give the wolves the ability to explore their environment for prey. In all other states the wander behavior was used to help the wolf overcome situations of indecision which may occur, for instance, when a wolf is exactly the same distance between two elk.

Finally, the avoid obstacles behavior was added to prevent

collisions. As the robot detects obstacles, repellant vectors are created, radiating away from the obstacle. This allows the robot to move around obstacles while searching and in pursuit of prey. The final vector that determines the robot's movement is the resultant of the vectors created by these constituent behaviors. The degree to which each behavior had an effect on the resultant vector is dependent on the gain, entered by the user, for each behavior. All of the parameters, including gains, for each state can be found in a table in the appendix.

Although the elk being preyed upon may have defensive behavioral strategy and coordination, the focus of the research to date has been on the hunting behavior of wolves. Therefore, the behavior of the elk was simplified with their reaction to the approaching wolves as simply either stopping or running away in a direction opposite to the wolves' approach. To create a range of test scenarios, the elks' behavior before the approach of wolves was varied to simulate situations where the elk are initially stationary, moving back and forth between multiple grazing areas, or wandering around. An example of the FSA to control the elk behavior for moving back and forth between two grazing areas is given in Figure 4, showing a simple modification that switches between an elk stopping upon seeing wolves, or running away.

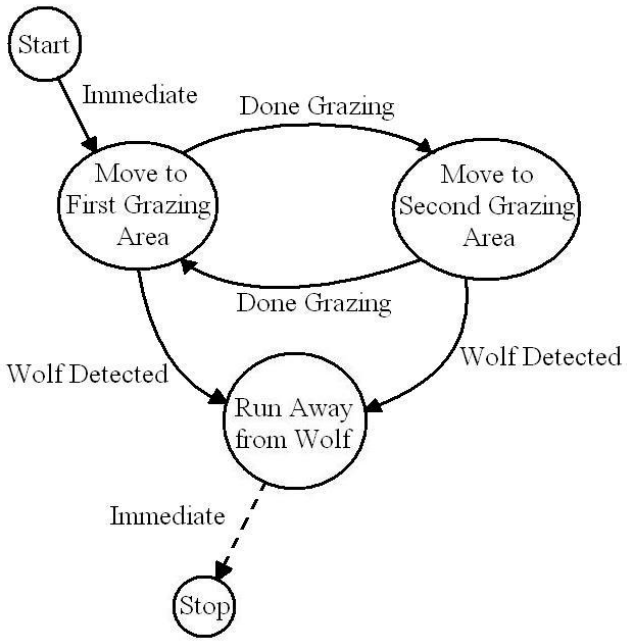

Figure 4. Finite State Acceptor of Elk behavior. If the 'run away' behavior is desired, remove the dotted trigger to the final stop state.

This project was conducted entirely with simulated robots; however, future work expects to move the system to physical robots. The platforms expected to be used are: (1) WowWee Rovio Wi-fi robots, and (2) iRobot Creates with the Element BAM (Bluetooth Adapter Module) as pictured below.

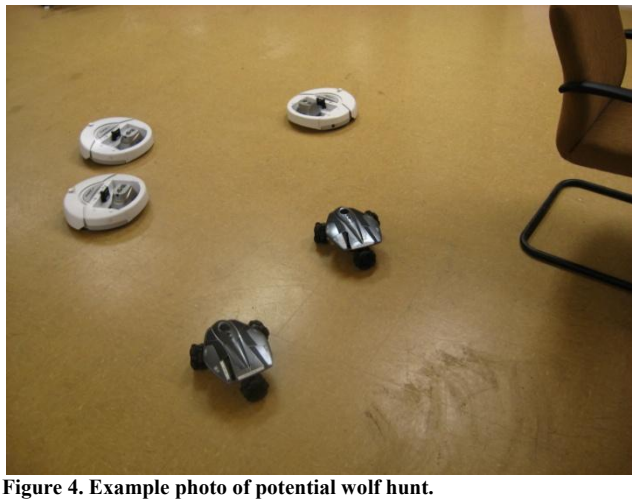

Figure 4. Example photo of potential wolf hunt.

IV. SIMULATION RESULTS

The *MissionLab* wolf pack simulations examined multiple scenarios that were commonly observed with wolf hunts in Yellowstone National Park. The underlay used for the simulations, is from the Lamar Valley in Yellowstone National Park where many of the actual observations of wolf hunts were taken.

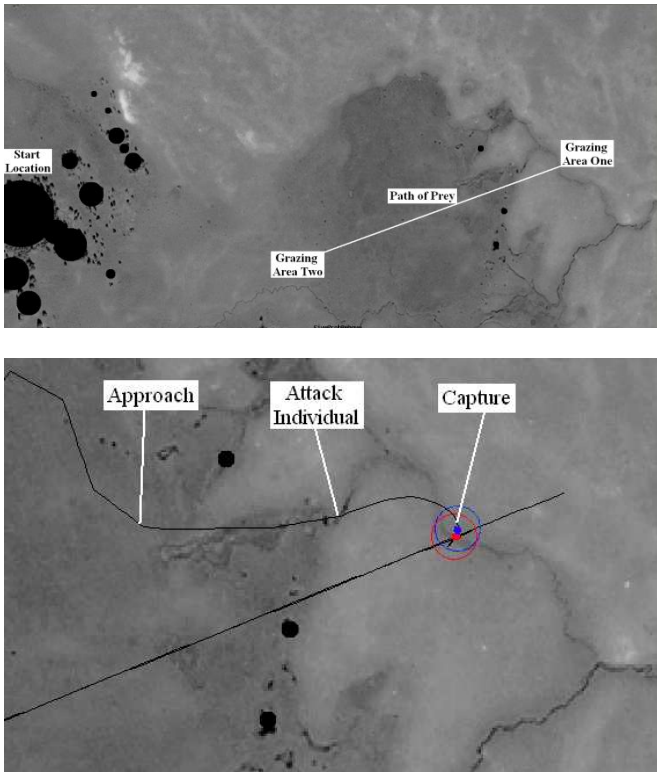

Figure 5. (Top): The hunting area used for experiments with starting location on the left, and grazing areas in the right and bottom of the map. (Bottom): Close-up showing transitions from search to capture for one wolf hunting one elk with stop behavior. The Straight Diagonal line is the prey's path between grazing areas. Underlay used for simulation of wolves in the Lamar Valley, Yellowstone National Park where many observations were recorded [13].

The first scenario was one on one between a wolf and an elk, with the elk moving back and forth between grazing areas and stopping when it perceived the wolf. The progression of the wolf through the foraging states was recorded for each run so that frequency of transitions could be tabulated for comparison between states and with the original observed probabilities. An example of a hunt run with these parameters is given in Figure 5 with starting location, grazing areas, and transition points in the hunt labeled. The next scenario was created by modifying the prey's behavior to run away, rather than stop, when it perceived the wolf. An example of this, also showing the wandering pattern of the wolf, is given in Figure 6. A third scenario had the same 'run away' behavior for the prey but with one wolf and three elk Figure 7. A final scenario involved two wolves and three elk as seen in Figure 8. For this scenario, the number of runs that both wolves ended up killing the same elk is compared to that of the wolves killing different elk. This comparison is made for both situations when the wolves discovered the elk together and when the wolves discovered elk separately. Twenty runs were completed for each scenario and the tabulated results of these runs are given in Table 4.

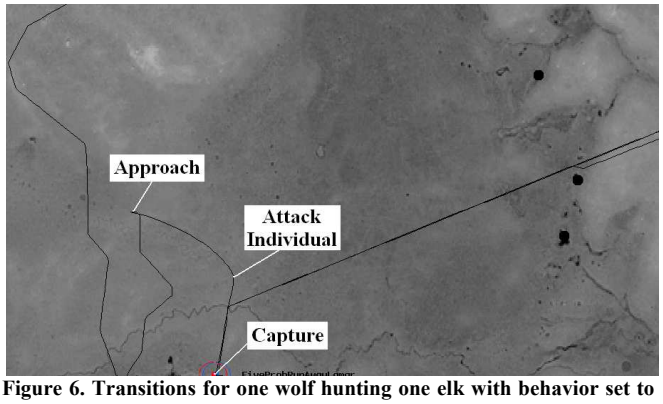

Figure 6. Transitions for one wolf hunting one elk with behavior set to run away from wolf. Random behavior of wolf search can also be seen previous to transition to approach.

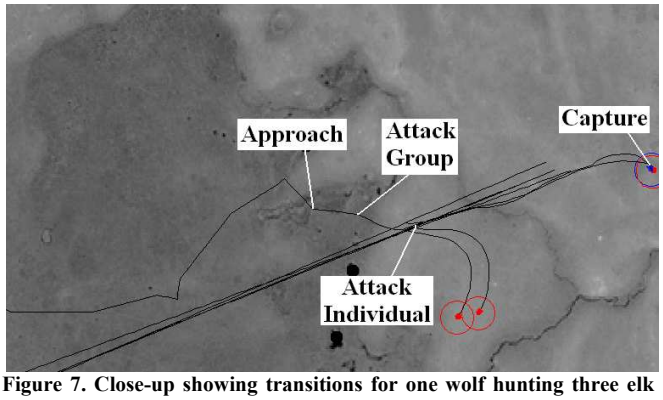

Figure 7. Close-up showing transitions for one wolf hunting three elk with behavior set to run away from wolf.

TABLE 4
Results From Wolf Simulations

One Wolf and One Elk (Stop)

Preceding State	Following State: Search	Approach	Attack Group	Attack Individual	Capture
Search	--	.95	.00	.05	.00
Approach	.00	--	.00	.88	.12
Attack Group	.00	.00	.00	.00	.00
Attack Individual	.08	.12	.00	.15	.67

One Wolf and One Elk (Run Away)

Preceding State	Following State: Search	Approach	Attack Group	Attack Individual	Capture
Search	--	1.00	.00	.00	.00
Approach	.07	--	.00	.93	.00
Attack Group	.00	.00	.00	.00	.00
Attack Individual	.07	.21	.00	.23	.49

One Wolf and Three Elk (Run Away)

Preceding State	Following State: Search	Approach	Attack Group	Attack Individual	Capture
Search	--	.73	.24	.03	.00
Approach	.04	--	.75	.21	.00
Attack Group	.07	.26	.13	.54	.00
Attack Individual	.09	.15	.22	.13	.41

Result of hunt	Runs	%
Discover elk together, kill same elk	4	20
Discover elk together, kill different elk	9	45
Discover elk separately, kill same elk	3	15
Discover elk separately, kill different elk	4	20

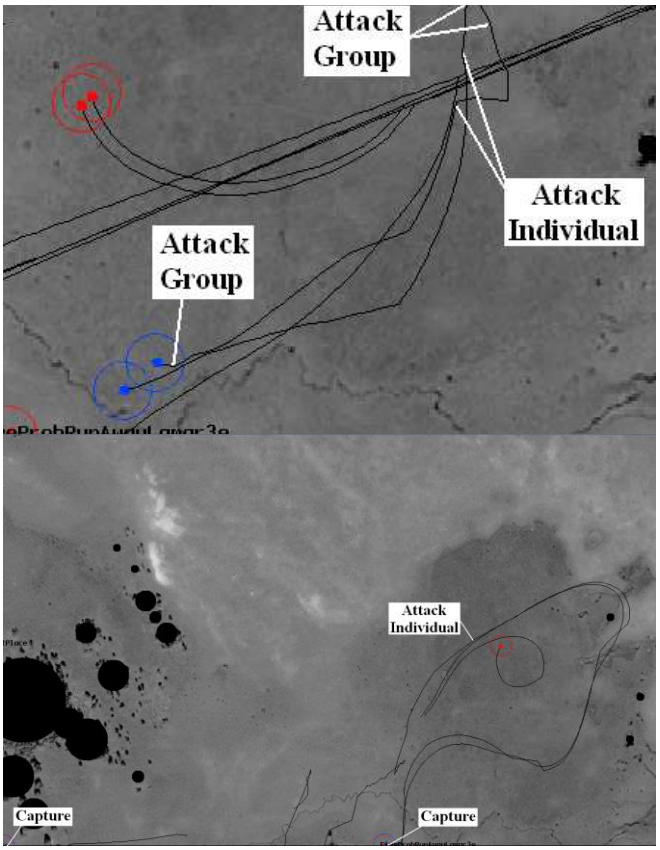

Figure 8. (Top): Close-up of transitions with two wolves approaching three elk together and separating during pursuit. (Bottom): Complete view of hunt after the two wolves split up and kill different elk.

V. Discussion of Results

Many of the resulting probabilities of transitions from the first scenario vary from the observed probabilities because in this scenario the prey would never run away whereas in the wild, running is the most common reaction for the prey. Comparison of the resulting transition probabilities from the first and second scenarios show that by changing the prey behavior from stopping when approached by a wolf to running away, the change in transitions is most notable in the transitions leading to attack. The first and third scenarios show similar differences for the same reason. Comparison of the second and third scenario reveals that adding multiple elk to the hunt has a large effect on the transitions for the obvious reason that the attack group state is only possible in the third scenario where there are multiple elk. This is the most realistic scenario as the vast majority of the observations in YNP were wolves hunting multiple elk. For this reason, the results from only the third scenario are compared to the observed data. The probabilities of transitions were similar to those in the observed data with the error for the primary four transitions (S→A, A→G, G→I, I→C) at 5%, 6%, 3%, and 11% respectively. Some transitions showed higher errors. Simulation results showed much lower probabilities for all transitions leading to the search state. This is most likely due to the hunts being confined within boundaries that often affected the elk's attempts to run away.

The final scenario involved two wolves and three elk. The purpose of this scenario was to examine how multiple wolves react to multiple prey. The wolves were given a slight attraction to one another through the move-to-object behavior to simulate actual hunts where the wolves generally start relatively together. This behavior made the discovery of prey with the wolves together more common. When the wolves came across the prey together, they most often killed different prey individuals. This may have been due to the confined space of the map allowing the capture of individuals they may have otherwise been lost and forced the pursuer to join their pack member. When they discovered the prey separately the results were split for killing the same or different individuals.

VI. Summary and Conclusions

By modeling wolf hunting behavior as a set of foraging states per the ethogram of MacNulty et al. 2007, and using a system of releasers and roulette wheel with probabilities of transitions, we were able to simulate interactions between wolves and elk with a relatively high fidelity to what is observed in the wild. The probabilities of transitions over all scenarios were similar to those in the observed data. This is not surprising as the observed data was used in the system that generated these results. Although there would seem to be an advantage in structured attack strategies, the high variability in behavior of the wolf's prey as well as the chaos inherent in attempting to locate, chase down, and kill one from a herd of hundreds of running elk would quickly cause strict strategies to breakdown. The loose adoption of general rules gives wolves the ability to react quickly and effectively. The results of simulations done in this project showed that the wolves were in fact reacting to the prey's behavior as evidenced by the change in transitions due to the prey stopping or running when attacked by the wolf.

The purpose of this study is to determine if high fidelity biological models can provide utility for a range of multi-robot applications, in this case with an emphasis on pursuit-evasion tasks. In particular, noting that byproduct mutualism can produce very robust results for biological groups has implications for the ability to reduce communication and planning requirements for robot groups, while still achieving purposeful missions. We believe our results to date support this goal. It is also worth noting that while the probabilities for transitions are currently invariable as defined directly from the wolf model, future research could address variations that may adapt to a range of different situations.

Appendix

This appendix provides the formulas for behaviors and transitions as well as the associated parameters used in the simulated wolves.

a) Probabilistic Transition: Transitions based on the existence of releasers and probabilities to simulate factors such as prey group size, wolf pack size, and environmental properties.

$$P_{State_n} = \frac{P_{State_{ni}}}{\sum P_{possible}}$$

$$State_{next} \begin{cases} State_1, & R < P_{State_1} \\ State_2, & R < P_{State_1} + P_{State_2} \\ \dots & \dots \end{cases}$$

where:

P_{State_n} = Probability of transitioning from the current state to $State_n$

$P_{State_{ni}}$ = Input probability of the above transition taken from MacNulty et al. 2007

$\sum P_{possible}$ = Sum of probabilities of transitions with releasers satisfied

$State_{next}$ = Resulting state of the transition

R = Random number between 0.00 and 1.00

b) Move-to-Object: Variable attraction to selected object. Used for attraction between pack members and attraction from wolf toward prey.

$V_{magnitude}$ = Adjustable gain value

$V_{direction}$ = Direction from the center of the robot to the center of the object, moving toward the object

c) Avoid-obstacle: Repel from object with variable gain and sphere of influence. Used for collision avoidance.

$$V_{mag} \begin{cases} \infty, & d \le r \\ \frac{max - d}{max - r}, & r < d \le max \\ 0, & d > max \end{cases}$$

$V_{direction}$ = Direction from the center of the robot to the center of the obstacle, moving away from obstacle

where:

max = Maximum obstacle detection sphere

d = Distance of robot to obstacle

r = Radius of obstacle

d) Noise: Random wander with variable gain and persistence. Used for exploration and to overcome local maxima, and minima.

$V_{magnitude}$ = Adjustable gain value

$V_{direction}$ = Random direction that persists for specified number of steps

Parameter	Value	Units
Wolf search assemblage		
Move to object gain	.2	
Selected object	Friendly	
Wander gain	.7	
Secondary wander gain	.3	
Avoid obstacle gain	.5	
Avoid obstacle sphere	10	m
Avoid obstacle safety margin	4	m
Wolf approach assemblage	.7	
Move to object gain	Enemy	
Selected object	.4	
Wander gain	.5	
Avoid obstacle gain	10	m
Avoid obstacle sphere	4	m
Avoid obstacle safety margin		
Wolf attack group assemblage		
Move to object gain	.8	
Selected object	Enemy	
Wander gain	.3	
Avoid obstacle gain	.3	
Avoid obstacle sphere	10	m
Avoid obstacle safety margin	4	m
Wolf attack individual assemblage		
Move to object gain	1	
Selected object	Enemy	
Wander gain	.1	
Avoid obstacle gain	.3	
Avoid obstacle sphere	10	m
Avoid obstacle safety margin	4	m
Wolf capture assemblage		
Move to object gain	.1	
Selected object	Enemy	
Wander gain	0	
Avoid obstacle gain	.1	
Avoid obstacle sphere	1.5	m
Avoid obstacle safety margin	.5	m

REFERENCES

[1] B. A. Duncan, P. D. Ulam, R.C. Arkin, "Lek Behavior as a Model for Multi-Robot Systems"

[2] D. Smith, "Wolves Chasing an Elk," National Park Service. Dec. 1, 2007, <http://www.nps.gov/yell/photosmultimedia/photogallery.htm?eid=379961&aid=547&root_aid=547&sort=title&startRow=10#e_379961>

[3] 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.

[4] D.R. MacNulty, Smith, D.W., Vucetich, J.A., and Packer, C., "Nonlinear Effects of Group Size on the Success of Wolves Hunting Elk", (in review), 2010.

[5] C.S. Holling, "The Functional Response of Predators to Prey Density and its Role in Mimicry and Population regulation," *Memoirs of the Entomological Society of Canada,* 45: 1-62.

[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] R. Abrantes, "The Evolution of Canine Social Behavior 2nd Edition," Ann Arbor, MI: Wakan Tanka Publishers, 2005.

[8] L.D. Mech, L. Boitani, "Wolves: Behavior, Ecology, and Conservation," University of Chicago Press. 2003.

[9] Georgia Tech Mobile Robotics Laboratory, "Manual for *MissionLab*," Version 7.0, 2007.

[10] D. MacKenzie, R.C. Arkin, J. Cameron, "Multiagent Mission Specification and Execution," *Autonomous Robots*, Vol. 4, No. 1, Jan. 1997, pp.29-57. Also appears in Robot Colonies, (eds.) R.C. Arkin, G. Bekey, Kluwer Academic Publishers, 1997.

[11] R.C. Arkin, "Behavior-based Robotics," MIT Press, 1998.

[12] R.C. Arkin, "Motor Schema-Based Mobile Robot Navigation," *International Journal of Robotics Research*, Vol. 8, No. 4, August 1989, pp. 92-112.

[13] Figure reprinted from Google, in accordance with their guidelines posted online. ©2009 Google – Imagery ©2009 DigitalGlobe, GeoEye, USDA Farm Service Agency, Map data ©2009 Tele Atlas.