The Predator's Purpose: Intention Perception in Simulated Agent Environments

Amani R. Maina-Kilaas
AMISTAD Lab
Harvey Mudd College
Claremont, CA, USA
amainakilaas@hmc.edu

Cynthia Hom
AMISTAD Lab
Harvey Mudd College
Claremont, CA, USA
chom@hmc.edu

Kevin Ginta
Math and Computer Science
Biola University
La Mirada, CA, USA
kevin.ginta@biola.edu

George D. Montañez
AMISTAD Lab
Harvey Mudd College
Claremont, CA, USA
gmontanez@hmc.edu

***Abstract*—We evaluate the benefits of intention perception, the ability of an agent to perceive the intentions and plans of others, in improving a software agent's survival likelihood in a simulated virtual environment. To model intention perception, we set up a multi-agent predator and prey model, where the prey agents search for food and the predator agents seek to eat the prey. We then analyze the difference in average survival rates between prey with intention perception—knowledge of which predators are targeting them—and those without. We find that intention perception provides significant survival advantages in almost all cases tested, agreeing with other recent studies investigating intention perception in adversarial situations and environmental danger assessment.**

***Index Terms*—awareness, intention perception, attention perception, simulation, predator-prey model, intention trilogy**

I. Introduction

You notice a small bug crawling on your bathroom floor. It walks in roughly a straight line, turns seemingly at random, and continues walking straight in its new direction. It repeats this behavior in a robotic, space-filling, randomized pattern. What is the purpose of its meandering? Is it looking for food, or a potential mate? You know nothing of the insect or its species, so you simply continue to observe. Its motion algorithm and circuitry seem to produce simple randomized behavior.

As far as you can tell, it pays no attention to you, even when you approach and decide to remove it from your home. A smarter, larger animal would become alarmed at your proximity, but the bug doesn't realize the danger posed by your sudden attention. If it sensed your intentions it would likely flee.

As an engineer, you ask yourself how you would create an algorithm for this sort of "danger detection." Would you simply test for sudden motion, or use features better correlated with intentionality? Leaves caught in the wind can be fast moving, but typically pose no threat. Some form of eye-tracking or motion-tracking, allowing one to judge path trajectories, could result in improved differentiation between threats and non-threats. Estimating the target of an eye gaze seems promising; there is likely a reason humans feel threatened when others stare at them, especially large crowds. We sense danger when the attention of many others is directed at us. It seems plausible that *some* ability to differentiate intention of agents versus non-intentional behavior, based on visual or auditory clues, can confer survival advantages.

In this study, we perform a series of tests in a simulated environment inhabited by virtual agents to measure the survival advantages gained by agents who are able to perceive intention. We seed the simulated world with small, fast agents searching for food and more powerful agents seeking to eat the smaller agents. We measure survival rates when large agent (*predator*) intention is perceptible to the smaller agents (*prey*), above and beyond merely sensing their general proximity, allowing them to sneak around agents and flee when in danger. We compare this to the survival rates of agents who cannot perceive intention, but can still sense proximity. Lastly, we measure survival rates when neither proximity nor intention is perceptible. We find that intention perception leads to statistically significant survival advantages in almost all cases tested.

The rest of the paper is organized as follows. Section II highlights related work, while Section III gives an overview of the agents and their behavior. We present the main results of the paper in Section IV and discuss them in Section V. Lastly, we draw conclusions and address future work in Section VI.

II. Related Work

Several other studies have used virtual agents to understand intention perception [1]–[3]. Heinze explored the use of virtual agents in intention perception studies, and proposed various methodologies [1]. Qi and Zhu developed a framework and algorithm for virtual agent intention perception, which they tested in a driving simulation comprised of multiple intention-aware agents [3]. Other researchers have emphasized the utility of intention-aware agent environments, and have discussed the complications of many-agent systems [2]. More specifically, Kaminka investigated mechanisms which allowed an agent to acquire, maintain, and infer knowledge of other agents. Kaminka also addressed several challenges in agent modeling, especially monitoring selectivity in the context of potentially many agents.

978-1-7281-8393-0/21/$31.00 ©2021 IEEE

Observing the interactions between two groups of agents is significant to other areas such as evolutionary psychology. As Barrett suggests, the influence of the predator-prey dynamic covers several domains of human psychology and behavior including perception, emotion, learning, inference, and reasoning [4].

The dynamics of predator-prey models have been widely studied [4]–[12]. There has been notable interest in understanding behavioral changes in agents when factors such as levels of stamina, hunger, and proximity awareness fluctuate [12]. These variants have been closely observed and simulated to provide insight regarding the reasoning of certain actions that the agents choose to take in any given situation.

Others have studied how predators can affect prey population density by stimulating costly defensive strategies such as limiting energy income (similar to our prey's food) and increasing vulnerability to other predators [8]. However, the importance of intimidation and consumption effects remains an open question.

III. METHODS

A. Simulated Predator-Prey Model

Our research begins to address the aforementioned open question. For a similar predator-prey simulation, one biologist suggested implementing "targeted kills" in which a predator selects a prey to target and chase [12]. Thus, we allow our predators to target the prey that they intend to eat. We then test whether a prey's ability to detect a predator's intention gives a measurable survival advantage over simple predator proximity awareness and over even simpler predator obliviousness.

We model our predators and prey as follows. Prey survival involves eating food and steering clear of predators. In our simulation, prey use observations to compute a probability distribution of possible directions depending on the favorability of each direction. They then sample a direction to move along according to that distribution. Prey know the locations of all food within their range of awareness and are more likely to move towards food if it is close or if they are hungry. Proximity-aware prey also know the locations of nearby predators and are more likely to move away from a predator if it is close. Intention-aware prey know when nearby predators are targeting them and will emphasize running away from those specific predators over others. Furthermore, prey are given incentive to continue with their current direction to help with indecision. When prey are targeted, they move as fast as they are able, and when not targeted, move with a speed in relation to the probability of the chosen direction.

In our simulation, a predator's only objective is to catch prey; they do not worry about being eaten themselves. Predators know the locations of all prey within their range of awareness and stochastically select one to target, with closer prey being more likely to be selected. Once a prey is targeted, predators move as fast as they can to eat the prey. At each moment during a chase, predators have a small chance of re-evaluating which prey to target—perhaps there is now a closer prey which would make more sense to chase. If the predator

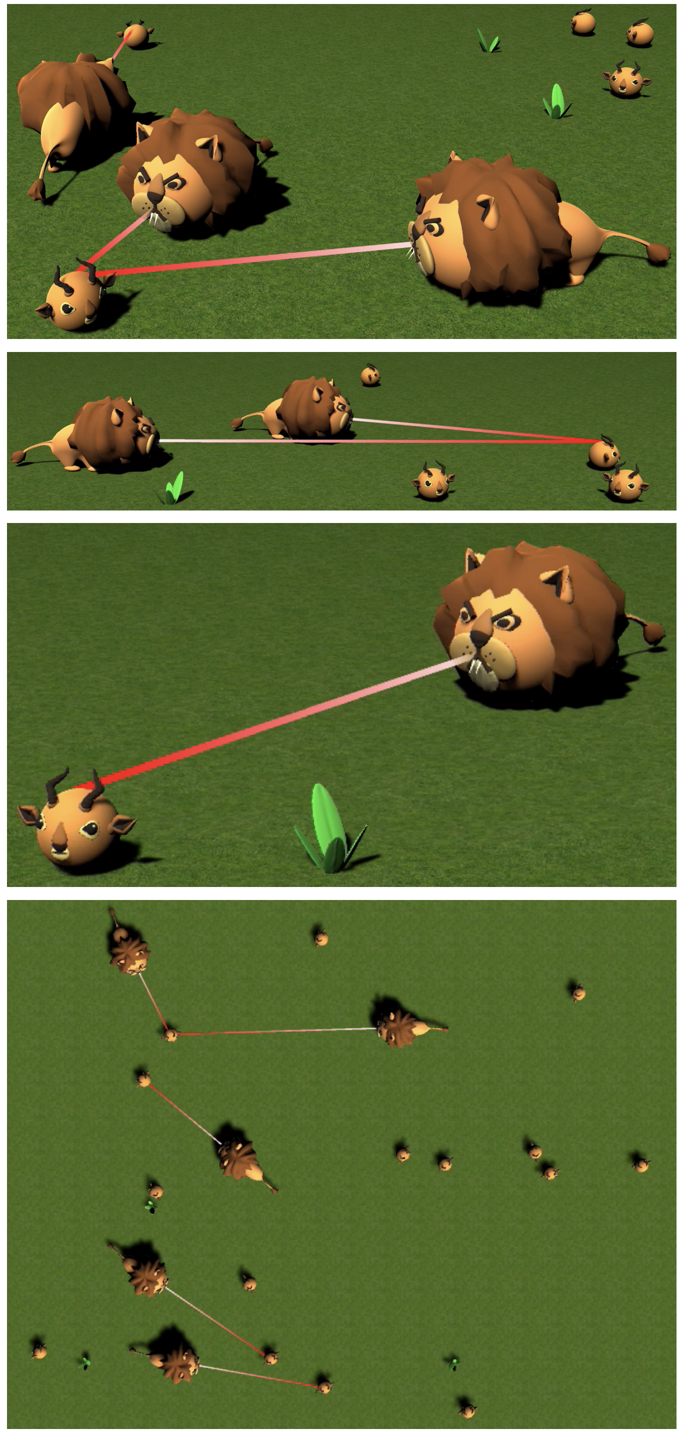

Fig. 1. Our predators are portrayed as lions, our prey as antelope, and our food as vegetation. A red line between predator and prey indicates that the predator is targeting that prey.

does not see anything, it moves randomly with emphasis around its current direction and with a speed relating to the chosen direction's probability. The prey and predator decision processes are outlined in pseudocode in Algorithms 1 and 2, which use notation from Table II in the Appendix.

Both predators and prey have measures of stamina to slow them down if they move too fast for too long and are given incentive to stay close to the center of the environment.

Because intention-aware prey focus on running away from predators and take fewer risks than other prey, one might worry that any benefits observed are due to caution rather than knowledge of intention awareness. To address this concern, we implement a "cautious" prey which acts on paranoia rather than intention perception. We retrieve data on how often intention-aware prey are targeted by predators and for how long, and use this to occasionally make the cautious prey think that they are being targeted by a random predator, regardless of whether that is true. In practice, this means that cautious prey are proximity-aware prey that will sometimes put emphasis on running away from particular predators—like the intention-aware prey—except that these predators may or may not actually be a threat. This decouples the knowledge from the behavior and strengthens the link between intention awareness and increased survival rate.

Further detail about these parameters and algorithms are given in the Appendix, and the full implementation of this simulation is available at https://github.com/AMISTAD-lab/predators-purpose-source.

Algorithm 1 Direction Algorithm

if agent is a predator that is targeting a prey **then**
  Return the direction to the targeted prey.
else
  for all $i = 0, 1, \ldots, 11$ **do**
    Set $\phi_i \leftarrow 30i + 15$ degrees.
    Set $p_i \leftarrow \sum_{s \in S} \Delta P_s(\phi_i)$, where S denotes the set of all stimuli that the agent is aware of and $\Delta P_s(\phi_i)$ is calculated using Eq. 5 in the Appendix.
  end for
  Normalize the distribution of the p_i's.
  Sample $i_{chosen} \in \{0, 1, \ldots, 11\}$ using the distribution of the p_i's and save the corresponding probability p_{chosen}.
  Sample direction ϕ_{chosen} using a uniform distribution over range of angles $30(i_{chosen})$ to $30(i_{chosen} + 1)$ degrees.
  Return $\phi_{chosen}, p_{chosen}$.
end if

B. Experimental Setup

Our agents are simulated as spheres in the PyBullet physics engine and rendered afterwards in the Unity3D game development engine. We randomly spawn several prey, predators, and food in a flat, circular, obstacle-free environment. At each time step, the agents make a decision regarding their direction and their speed and move forward accordingly.

A prey is considered eaten when it comes into contact with a predator. The prey eats food in a similar fashion, by coming into contact with the food objects.

Algorithm 2 Speed Algorithm

Initialize $V_{max} \leftarrow V_m$ (where V_m is the agent's maximum speed).
if agent is tired ($S_{t'} < S_{tired}$ for some previous step t' and $S_t < S_0$ for all $t > t'$, where S_t denotes the stamina at step t) **then**
  Set $V_{max} \leftarrow V_{tired}$.
end if
if agent is a predator targeting a prey **then**
  Return $\frac{2}{3} V_{max}$.
else if agent is a prey being targeted **then**
  Return V_{max}.
else
  Set $V_{ideal} \leftarrow \frac{p_{chosen} \cdot V_m + V_{current}}{2}$.
  Sample speed v using a normal distribution with $\mu = V_{ideal}$ and $\sigma = V_m/6$.
  Return $\min(|v|, V_{max})$.
end if

Information about the prey, such as the amount of food eaten and their lifespan, is recorded during every run. These results are statistically analyzed to assess the relative survival rates for intention-aware, proximity-aware, unaware, and cautious prey, and their typical mode of death.

We vary several parameters in our experiments. The first are prey and predator **sight distance**, which gives the radius of each agent's field of view. Then there is **sight angle**, which gives the angle that their field of view spans. The **speed fraction** is the ratio of the predator's speed to the prey's speed. Finally, we have the **Maximum Fasting Interval (MFI)**, which is number of time steps a prey can go without eating food before it starves.

For computational efficiency, we set each parameter to a default value and vary them individually. In our values we use a distance reference based on prey body width, so that the diameter of a prey is 1 unit and that of the larger predator is 2 units.

Prey typically have a very wide field of view and strong additional senses to help with awareness of their surroundings (e.g., mice and gazelle). In our simulation all knowledge is through sight, and so to account for this situational awareness of food through smell or predators through hearing, we give our prey a fixed sight angle of 360°. However, a wide field of view generally comes with a shorter range of sight, and so we give prey a default sight distance of 10—enough to identify food and predators but small enough to contain primarily relevant information.

Predators, on the other hand, tend to have more narrow fields of view but very sharp eyesight (e.g., eagles and lions). In accordance with this, we give them a sight angle of 90°, which is wide enough to not incapacitate the the predators and narrow enough to prevent full side-to-side vision. We also note that in exploratory simulations, the sight angle did not appear to have a large effect since predators frequently change direction in order to find and focus on particular prey. To incorporate the sharp eyesight of predators, we give them a sight distance of 20 (doubling that of prey) which is generally sufficient to identify most prey in their field of view.

We give prey a fixed maximum speed of 31.25 prey lengths per 100 time steps, which is somewhat arbitrarily chosen to yield a reasonable duration for viewing the simulation. More importantly, predators have a default speed fraction of 0.8 so that they are a little slower with a maximum speed of 25 prey lengths per 100 time steps. Additionally, we test scenarios where predators move faster than prey (see Fig. 3).

For the Maximum Fasting Interval, we set this to a default of 2000 steps as we found that this provided a substantial cost to not eating while not being too punishing.

Note that while these are the default values, each parameter is individually varied to test other scenarios, such as when predators are faster than their prey. These parameters and values are summarized in Table I. See Section A of the Appendix for further details.

TABLE I
DEFAULT PARAMETER VALUES.

Description	**Value**
Prey sight distance	10
Predator sight distance	20
Prey sight angle	360
Predator sight angle	90
Speed fraction	0.8
Prey to predator ratio	4
Maximum Fasting Interval (MFI)	2000

IV. RESULTS

Our results show that intention perception provides prey with a measurable, statistically significant survival advantage in almost all cases. Not only are more prey with intention perception alive at each time step on average than others, but they also have the longest lifespan across most parameters.

Each of the line graphs in Fig. 3 was generated from 1,000 independent runs (each consisting of 10,000 time steps) per set of parameters, except for the *cautious* results, which were computed based on 100 runs per set. Each line denotes the uniform average of all runs, and is surrounded by a (near-imperceptible) 95% confidence interval indicating the statistical significance of the results. Similarly, the stack plots in Fig. 2 were each generated from 1,000 independent runs per intention tier, except for the cautious plot which was computed from 100 independent trials.

In Fig. 2, we see that at default values, a greater proportion of intention-aware prey are alive at each time step than that of any others. Fig. 3 shows that the intention-aware prey also have the longest lifespans as we vary most parameters—only briefly are they overcome by the prey with just proximity awareness and the prey with caution in the speed fraction plot when the speed fraction is less than 0.5. However, as soon as the fraction increases past 0.5, predators become fast enough to become a significant threat, such that prey with intention and proximity awareness begin to live much longer than all other prey.

In almost all plots, we see that cautious prey have the second longest lifespans, followed by proximity-only prey and

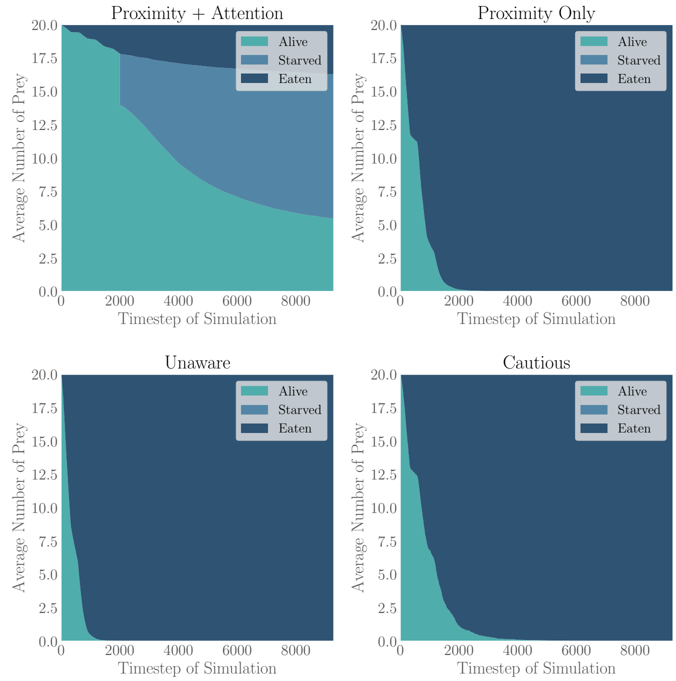

Fig. 2. The effect of intention perception as seen through the status of prey progressing through a simulation with the default parameters.

then unaware prey. The exceptions to this are of course speed fraction, as well as predator sight distance, where cautious and proximity-only prey have very similar lifespans at low predator sight distance values.

We also note that a greater proportion of prey with intention perception die from starvation than any of the other three types of prey (see Fig. 2). This is why the average lifespan of prey with intention perception increases rapidly with respect to the MFI, the maximum duration a prey can stay alive without eating, until MFI reaches 4000 and the intention-perception prey have enough time to eat while simultaneously fleeing predators (see Fig. 3). However, the trade-off between starvation and predator avoidance is worth it, as at all maximum fasting intervals under standard conditions, the prey with intention perception still survive much longer than those without.

We notice some specific trends across the six varied parameters. At low MFI and prey sight distance values, all prey have similarly low survival rates because they all either starve almost immediately or are unable to see predators, respectively. On the other hand, at low speed fraction and predator sight distance values, all prey have similarly high survival rates because predators are either too slow or unable to see, respectively, preventing them from catching most prey.

At high values of most parameters, prey with intention perception live far longer than the others. However, at high speed fraction values, predators are fast enough that even prey with intention perception are eaten before they can escape, resulting in short lifespans for all types of prey.

We also notice peaks in the graphs for prey sight distance, speed fraction, and prey to predator ratio. For prey

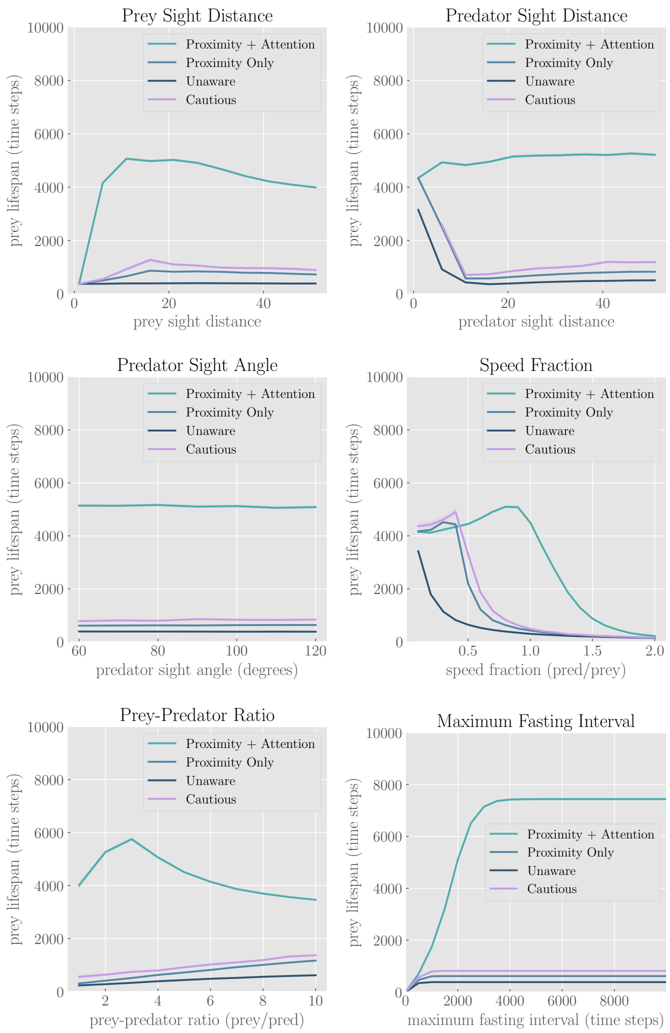

Fig. 3. The effect of varying different parameters on prey lifespan. All lines include near-imperceptible 95% confidence intervals. Note: There were too few observations with which to estimate the tracking distribution for the first point in Predator Sight Distance for the cautious prey, and therefore the point was left out.

sight distance, it seems that too much information negatively impacts survival. All three tiers with awareness have a peak after a certain sight distance—perhaps beyond that point the additional information muddles their priorities and makes them less likely to choose the optimal direction.

For the speed fraction, we expect a downward trend as predators become faster and eat more prey. This does appear in the graph, but not until a certain point. This is because once we factor in hunger, slightly faster predators can chase prey towards food that relatively idle prey might not find. The peak is where the main cause of death shifts from starvation to being eaten.

For the prey-to-predator ratio, when there are more prey, a smaller percentage will be eaten. However, the prey that avoid the predators are likely on the outskirts of the terrain and will starve because they are not eating. It is also harder to find food when there are more prey. The peak is where we find the balance of being less likely to starve but not too likely to be eaten.

Lastly, the lack of a relationship between prey lifespan and predator sight angle is likely due to a predator's behavior when it cannot see prey, as it turns in a random direction every frame until it finally locks onto a prey. As such, low values do not have a huge effect on the predator's ability to find and target prey.

V. Discussion

Across the varied parameters of our experiments, we see that intention perception almost always provides the greatest survival advantage, followed by caution, proximity awareness, and finally, no awareness. It makes sense that caution provides a slightly greater survival advantage over proximity awareness, as cautious prey have both proximity awareness along with a randomly occurring speed boost provided by their cautious nature. However, the average lifespan of cautious prey is almost always shorter than that of intention prey, establishing that the survival advantage of prey with intention perception is due to intention perception specifically, as opposed to mere caution.

There is one case in which prey with only proximity awareness and prey with only caution have longer average lifespans than prey with intention perception—low speed fraction values. Speed fraction is the ratio of predator speed to prey speed. At low speed fraction values, prey with intention perception run away from predators that are often too slow to catch them, unnecessarily preventing them from eating food which leads to their starvation. Although one might think that this would also happen to cautious and proximity-only prey, cautious prey sometimes think that a faraway predator is targeting them, such that they mainly use proximity awareness to decide upon a direction. Proximity-only prey, by nature, also only fear predators that are close to them, allowing them to focus on eating food unless in immediate danger. As such, prey with intention perception do not survive the longest when the prey are very fast but the predators are very slow.

We also note that although prey with intention perception survive the longest overall, prey with intention perception that do not survive are more likely to die from starvation than prey of any other type. This is because their intention perception ability makes them prioritize fleeing predators that intend to eat them over eating food themselves.

Overall, intention perception does confer statistically-significant survival advantages in most cases, with only one observed exception (when predators are very slow compared to prey). Furthermore, the advantages of intention perception were not simply due to caution. We saw that decoupling the "cautious" aspect from the "informational" aspect of intention perception demonstrated that the same gains were not realized by agents equipped with caution alone. Finally, although prey with intention perception are the most likely to starve, they are the least likely to be eaten, and thus have the highest overall survival rate.

VI. Conclusion

The goal of our study was to investigate whether a prey's ability to detect a predator's intention gives a measurable survival advantage over simple predator proximity awareness and over even simpler predator obliviousness. To test the hypothesis that intention awareness is advantageous, we created a multi-agent predator-prey simulation and analyzed the survival of prey with and without intention perception. The results of our study establish that intention perception provides significant survival advantage for agents, as our intention-aware prey had the highest survival rates in almost all of the experiments conducted.

However, prey with intention perception do starve the most, since they often prioritize avoiding predators over food consumption. This behavior leads to lower survival rates in certain situations where predators are less of a threat. Yet, intention perception still provides the greatest survival advantage overall, as it allows prey to observe when predators are targeting them and act accordingly.

We show that it is the knowledge gained through intention perception, and not caution in general, that greatly increases survival rates. In our experiments, we found statistically significant differences in the survival rates between intention-perception prey and cautious prey, which have the same level of caution as intention-perception prey but lack true information about predator targeting. This provides additional evidence of the value of intention perception.

While our results are strikingly clear and unambiguous, we remind the reader that these are *virtual* experiments in *simulated* multi-agent environments, and thus might not be fully representative of actual predator-prey scenarios. All models are wrong, yet some are useful: we believe that our experiments provide insight into the advantages of intention perception in software agents, and perhaps might even suggest similar survival advantages for living organisms. Without further research in actual predator-prey systems, however, these suggestions remain no more than tantalizing possibilities. Extensions to this work include using evolutionary algorithms to *evolve* predator and prey strategies, rather than using hand-designed motion algorithms. Future work to probe advantages of intention perception in other multi-agent situations remains ongoing.

Acknowledgment

Special thanks to Cindy Lay for aid in the preparation of this paper, and to Jerry Liang, Aditya Khant, Kyle Rong, and Tim Buchheim for assistance in experimental set-up. This research was supported in part by the National Science Foundation under Grant No. 1950885. Any opinions, findings, or conclusions are those of the authors alone, and do not necessarily reflect the views of the National Science Foundation.

References

[1] C. Heinze, "Modelling Intention Recognition for Intelligent Agent Systems," Defence Science and Technology Organisation, Salisbury (Australia) Systems, Tech. Rep. DSTO–RR–0286, 2004.

[2] G. Kaminka, J. Wendler, and G. Ronen, "New Challenges in Multi-Agent Intention Recognition," in *Proceedings of the Fall Symposium on Intention Recognition for Collaborative Tasks*, 2001.

[3] S. Qi and S.-C. Zhu, "Intent-Aware Multi-Agent Reinforcement Learning," in *2018 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2018, pp. 7533–7540.

[4] H. C. Barrett, "Adaptations to Predators and Prey," *The Handbook of Evolutionary Psychology*, p. 200, 2005.

[5] M. L. Rosenzweig and R. H. MacArthur, "Graphical Representation and Stability Conditions of Predator-Prey Interactions," *The American Naturalist*, vol. 97, no. 895, pp. 209–223, 1963.

[6] M. P. Hassell, *The Dynamics of Arthropod Predator-Prey Systems*. Princeton University Press, 1978.

[7] B. Basener and D. S. Ross, "Booming and Crashing Populations and Easter Island," *SIAM Journal on Applied Mathematics*, vol. 65, no. 2, pp. 684–701, 2004.

[8] E. L. Preisser, D. I. Bolnick, and M. E. Benard, "Scared to Death? The Effects of Intimidation and Consumption in Predator-Prey Interactions," *Ecology*, pp. 501–509, 2005.

[9] K. Hawick, H. James, and C. Scogings, "A Zoology of Emergent Patterns in a Predator-Prey Simulation Model," in *Proceedings of the Sixth IASTED International Conference on Modelling, Simulation, and Optimization, Gabarone, Botswana*, 2006, pp. 84–89.

[10] R. Gras, D. Devaurs, A. Wozniak, and A. Aspinall, "An Individual-Based Evolving Predator-Prey Ecosystem Simulation Using a Fuzzy Cognitive Map as the Behavior Model," *Artificial life*, vol. 15, no. 4, pp. 423–463, 2009.

[11] R. Cressman and J. Garay, "The Effects of Opportunistic and Intentional Predators on the Herding Behavior of Prey," *Ecology*, vol. 92, no. 2, pp. 432–440, 2011. [Online]. Available: http://www.jstor.org/stable/41151152

[12] D. Richards, M. J. Jacobson, J. Porte, C. Taylor, M. Taylor, A. Newstead, I. Kelaiah, and N. Hanna, "Evaluating the Models and Behaviour of 3D Intelligent Virtual Animals in a Predator-Prey Relationship," in *Proceedings of the 11th International Conference on Autonomous Agents and Multiagent Systems-Volume 1*. Citeseer, 2012, pp. 79–86.

Appendix

A. The Agents

Our predators and prey have measures of hunger and stamina to enhance realism. *Hunger satiety*, denoted by H, starts at an initial value H_0, decreases by ΔH_t every time step, and increases by ΔH_{eat} each time they eat. Moving forward, we will often refer to hunger satiety as simply hunger, but note that this measure decreases over time contrary to actual hunger, since it represents satiety (fullness). This hunger factor is used to increase the likelihood of prey moving towards food. Hunger has a minimal value of $+\Delta H_t$ and any deaths that should have resulted from starvation are accounted for after the simulation, with the default maximum fasting interval set at 2000 time steps. It is also important to note that one food spawns every 70 time steps and that the amount of food is capped at the initial number of prey.

Stamina, denoted by S, starts at an initial value S_0 and changes according to how fast agents move with respect to their threshold speed V_{thresh}:

$$\Delta S = S_{weight} \cdot \left(\frac{V_{thresh} - V}{V_{thresh}} \right) \tag{1}$$

where S_{weight} is a value that varies depending on whether the change is positive or negative and V is the current speed.

When an agent's stamina measure reaches their tired stamina S_{tired}, their speed, originally restrained by a maximum speed V_m, becomes restrained by a tired speed V_{tired} until they restore their stamina back to S_0.

TABLE II
HUNGER AND STAMINA CONSTANTS.

Agent	Name	Notation	Value
Both	Initial hunger satiety	H_0	1.0
Prey	Hunger satiety decrement	$H_{t,py}$	0.0005
Pred.	Hunger satiety decrement	$H_{t,pd}$	0.0003
Both	Hunger satiety increment	H_{eat}	0.5
Prey	Initial stamina	$S_{0,py}$	1.5
Pred.	Initial stamina	$S_{0,pd}$	1.0
Both	Tired stamina	S_{tired}	0.2
Both	Stamina weight (+)	$S_{weight,+}$	0.01
Both	Stamina weight (-)	$S_{weight,-}$	0.0025
Prey	Max speed	$V_{m,py}$	31.25
Pred.	Max speed	$V_{m,pd}$	*
Both	Threshold speed	V_{thresh}	$\frac{V_m}{2}$
Both	Tired speed	V_{tired}	$\frac{V_m}{3}$

*Note: This value is a fraction of the prey max speed and varies between simulation runs (default is $0.8 \cdot V_{m,py}$).

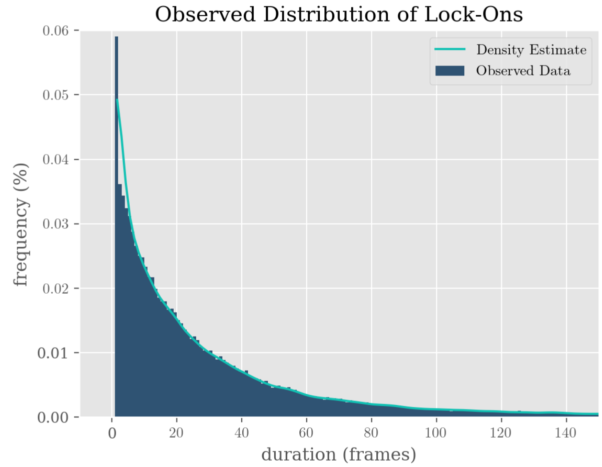

Fig. 4. Distribution of tracking durations under default parameters.

Both prey and predators decide their direction and speed using the direction and speed probability distribution algorithms as outlined by Algorithms 1 and 2 in Section III. The specific processes for sampling direction and speed are further detailed in Sections C and D. Note that when prey are targeted, their speed is instead set to their maximum (restrained by stamina). Also note that predators only sample direction and speed when they do not see any prey. When prey become visible, the predator randomly selects a prey with each one having probability proportional to $^1/_d$, where d is the distance to the prey. When a prey is targeted, the predator direction is always towards the targeted prey and the predator's speed is set to the maximum it can be (restrained by stamina). Note that predators of any fixed speed move much faster when not faced with indecision, as the motion becomes entirely in one direction, and thus we set the max speed to $^2/_3\, V_{max}$ when they target a prey as a form of counterbalancing this effect. At each time step when a prey is targeted, the predator has a 5% chance to reevaluate the nearby prey and choose again.

The environment is a flat square but the agents are discouraged from exploring all of it. The side length of the environment is $E = 250$, but we limit the reasonably accessible terrain to a circle with radius $R_{terrain} = {}^3/_{10}\, E = 75$. Predators are only allowed to target prey up to a maximum chase radius, given by $R_{chase} = {}^7/_{10}\, R_{terrain} = 52.5$, with a growing probability of "untargeting" the prey:

$$P_{untarget} = \frac{d}{(R_{chase})^{15}} \tag{2}$$

where d is the distance to the center of the environment. Furthermore, all agents are placed within a given spawn radius $R_{spawn} = {}^7/_{10}\, R_{chase} = 36.75$ and are encouraged to stay close by the direction probability distribution algorithm. Note that we cap $P_{untarget}$ at 1 to ensure it is a valid probability.

B. Values for Cautious Prey

To create the cautious prey, we used observed data gathered from 100 simulations for each parameter configuration (seed), most of which had 20 intention-aware prey. We averaged each prey's number of "lock-ons" per lifespan, giving the probability of a prey getting targeted each frame. These probabilities are in Table III. Note that the "val" column represents the value of the parameter of interest (with all other parameters set to their defaults), and that the entries under the parameters indicate the probability of a lock-on in each frame as a percent, for that configuration. For example, when $d_{py} = 1$, the probability is 0.00453. Note that d_{py} and d_{pd} refer to prey and predator sight distance, respectively; n_{py}/n_{pd} refers to prey to predator ratio; v_{pd}/v_{py} refers to speed fraction and has parameter values in 10%'s; and θ_{pd} refers to the predator sight angle and has parameter values in degrees. With this information, we use the probability of a prey getting targeted to determine on each frame whether the prey should think a predator is targeting them, and if so, we choose a predator at random and sample a duration of this fake lock-on from an estimated distribution model. To model this distribution of tracking durations we used nonparametric kernel density estimation (KDE) based on observed tracking durations for intention perception agents, with reflection around the $y = 1$ boundary to ensure validity of sampled points. Fig. 4 shows the distribution of tracking durations for the default configuration.

C. Direction Probability Distribution

To explain our direction probability distribution, we must first define our weighting system.

We use $D_\phi(\phi_1, \phi_2)$ as a measure of similarity between two angles, defined as the length of the chord between them on a unit circle, so zero is most similar and two is most different with the angles differing by 180°. With this, we weight similar angles as more favorable according to our angle weight formula:

$$W_\phi(\phi_1, \phi_2) = \left(D_\phi(\phi_1, \phi_2) + \frac{1}{2} \right)^{-3} \tag{3}$$

TABLE III
OBSERVED TARGETING FREQUENCIES (%).
DOTS INDICATE INAPPLICABLE PARAMETER VALUES.

val	$\mathbf{d_{py}}$	$\mathbf{d_{pd}}$	$\theta_{\mathbf{pd}}$	$\mathbf{n_{py}/n_{pd}}$	$\mathbf{v_{pd}/v_{py}}$
1	.453	.001	·	.739	.830
2	·	·	·	.508	.747
3	·	·	·	.472	.682
4	·	·	·	.477	.635
5	·	·	·	.447	.606
6	.303	.063	·	.421	.565
7	·	·	·	.397	.525
8	·	·	·	.371	.480
9	·	·	·	.355	.399
10	·	·	·	.334	.333
11	.501	.282	·	·	.249
12	·	·	·	·	.205
13	·	·	·	·	.179
14	·	·	·	·	.170
15	·	·	·	·	.189
16	.560	.400	·	·	.215
17	·	·	·	·	.240
18	·	·	·	·	.296
19	·	·	·	·	.374
20	·	·	·	·	.426
21	.396	.491	·	·	·
26	.316	.568	·	·	·
31	.295	.645	·	·	·
36	.295	.690	·	·	·
41	.298	.734	·	·	·
46	.296	.769	·	·	·
51	.312	.794	·	·	·
60	·	·	.377	·	·
70	·	·	.400	·	·
80	·	·	.443	·	·
90	·	·	.465	·	·
100	·	·	.487	·	·
110	·	·	.536	·	·
120	·	·	.561	·	·

For a weighing of physical proximity, we use our distance weight formula, D_x, as defined by

$$D_x(X_1, X_2) = \frac{d_{sight}}{d} \tag{4}$$

where d_{sight} is the maximum distance the agent can see and d is the Euclidean distance between the two positions X_1 and X_2. The sight distance of the agents is determined in the simulation, as well as the sight angle for the predator. The sight angle for the prey, however, is set to $360°$ as prey generally have wider range of view and often use scent and hearing for greater proximity awareness.

To determine the probability distribution that an agent uses to move, we assign probabilities to 12 bins of angles (evenly dividing $0°$ and $360°$) by iterating through the visible stimuli and adding the probability increase associated with it.

For a given stimulus (i.e., food), denote by F its corresponding stimulus factor, by ϕ_s the favorable angle for the stimulus (e.g., towards if food, away if a predator), and by X_s the position of the stimulus. Then the probability increase for angle ϕ of an agent at position X_0 is given by

$$\Delta P(\phi) = F \cdot W_\phi(\phi, \phi_s) \cdot D_x(X_0, X_s) \tag{5}$$

except in two unique cases. In the case of adding probability for the current angle, we forgo the distance weight component. In the case of adding probability towards the center of the environment, we use replace D_x with

$$D_{center} = \frac{d}{(R_{spawn})^4} \tag{6}$$

where d is the distance to the center and R_{spawn} is the radius of the circle in which we place our agents (so that agents do not drift too far from the central location). All of the probabilities are then normalized and a bin is randomly selected. We then sample a specific angle within that bin using a uniform distribution.

All of the stimulus factors for prey and predators are given in Table IV. Note that the prey's center of environment factor is not applied when it can see predators, as it should still prioritize running away.

TABLE IV
STIMULUS FACTORS.

Agent	Stimulus	Value
Prey	An ordinary predator	1.5
Prey	A targeting predator	6.0
Prey	Food	$1.5 \cdot \frac{1}{H}$
Prey	Center of environment	5.0
Pred.	Center of environment	20.0
Both	Current direction	30.0

D. Speed Probability Distribution

First we find the normalized probability of the bin selected by the direction probability distribution algorithm, p_{chosen} and compute a chosen speed

$$V_{chosen} = p_{chosen} \cdot V_m \tag{7}$$

To minimize drastic changes in speed, we average the chosen speed with the current speed and use that as the ideal speed V_{ideal} (see Algorithm 2 in Section III).

The actual speed is chosen randomly according to a normal distribution centered on V_{ideal} with a standard deviation of $\frac{V_m}{6}$, and the absolute value is taken (to avoid "negative" speeds). It is then capped at V_{tired} if the agent's stamina is less than S_{tired}, and at V_m otherwise.