



Enlisting Robots for Guard Duty

An Artificial Intelligence Approach to Border Security

What is the optimal way to deploy a robot patrol squad along a security fence or around an army base? *Profs. Gal Kaminka and Sarit Kraus*, and their crack BIU computer science research team are investigating how to achieve successful "robotic guarding". Their findings bode well for Israeli defense and are triggering the formation of the world's first robot border patrol.

Translated and adapted from the Hebrew article by Yisrael Binyamini, published in *Galileo*, Vol. 121

Penetrating a Guarded Compound

It's a scene familiar to moviegoers everywhere: the hero or villain lies in the shadows near a fence, carefully counting the number of guards on duty and analyzing their movements. If he makes his move in the right place at just the right time, he has a chance of breaking through. If not, the guards will capture him and all hope will be lost.

Choosing the optimum break-in spot for the above scenario is really a matter of mathematics. That's why a team of Bar-Ilan researchers has taken up the challenge of designing knowledge-sharing robots that can work together to breach the fence – or guard it successfully.

"If it takes approximately three minutes to cut and cross the fence, intruders should find a spot where it is unlikely for a guard to pass during those three minutes," says Prof. Gal Kaminka, an expert in artificial intelligence, robotics and applied philosophy who heads the Bar-Ilan lab known as MAVERICK (an acronym for Multi-Agent Virtual Environments Robots Intelligence Cooperation and Knowledge). "Let's assume that – because of the limited number of guards – each spot in the fence may be left unguarded for a time interval up to three minutes in length. Under these circumstances, while penetration is possible, it's the responsibility of the guards to make sure that it's difficult for the intruder to predict the fence's most vulnerable spots."

Kaminka explains how this cat-and-mouse game is based on a balance of knowledge and assumptions. The intruder might want to try his luck at the last spot a guard recently passed, assuming that the guard won't return too soon. But the guards know that this would be a logical choice, so they will try to outfox the intruder by surprisingly reversing their direction. This, however, does not guarantee protection: if the would-be intruder observes the guards' behavior long

enough, he would be able to estimate the probability of direction reversal at any given moment. If the guards can be sufficiently random in the way they reverse their steps, this probability factor is moot: the intruder will not have enough information to predict the direction of the guards' movement.

The above scenario was recently dealt with in the MAVERICK laboratory. Prof. Kaminka, along with his departmental colleague, Prof. Sarit Kraus, and doctoral students Yehuda Elmaliach and Noa Agmon, have been exploring the challenges involved in using robots – rather than human guards – in border areas. This article presents some of their main findings about how successful "robotic guarding" can be achieved.

In the Name of Protection

Enlisting robots for guard duty is an attractive idea. Robots can be equipped with cameras that transfer photographs or video to the center of command where the human guards sit, protected both from bad weather and from enemy attack. Moreover, the human guards can get a better overview of the situation, by watching simultaneously transferred images from several robots. But who decides how the robots actually do their guarding? While some may assume that the humans should be in charge, MAVERICK studies have shown that leaving the movement decisions in the hands of the robots actually improves the chances of protecting the compound successfully.

The main challenge is engineering the guard robots' movements in order to maximize the chance of discovering a break-in attempt.

To create successful robotic guarding, the main challenge is to engineer the robots' movements in order to maximize the chance of catching the intruder as he or she attempts to break through the fence.

»

>>

The first, most simple model to consider is a fence that can be circled by one-way continuous movement. For demonstration purposes, let's assume that a robot is guarding a large military base, that the perimeter fence surrounding the base is three kilometers long, and that the robots can move at a speed of up to 200 meters per minute (12 km/h). As such, a robot can complete a full circuit of the fence in fifteen minutes. If there are three robots and the distances between them are equal, then the distance between each pair of robots will be one kilometer, which a robot can cover in five minutes.

Under these circumstances, it is obvious that the intruder must not select a point on the fence, where the closest robot to it is 2.5 minutes away. This is bad news for the intruder because – remember – in our scenario, it takes three minutes to break through the fence. On the other hand, these numbers also add up to bad news for the defenders: There are many spots that no robot will reach within four minutes.

How so? Let's assume for a moment that every spot in the fence is visited by a robot every four minutes. The length of the fence is three km, so that three robots combined cover three km every four minutes. Under these circumstances, at least one robot will be required to cover one km or more in four minutes. In order to achieve this, a



velocity of fifteen km/h is necessary. However, the maximum velocity of these robots is only 12 km/h. This conclusion is not dependent on the robots' movement engineering, their shift of directions or the distance between them.

If the intruder were able to predict the movements of the robots, he could easily choose one of the spots that would not be guarded over the next four minutes, and penetrate the compound without fear of getting caught.

Coordinated Choreography

An important question related to planning the movement of robotic guards is whether the movements of the robots should be coordinated. Alternatively, perhaps it would be preferable for each robot to operate independently. The answer reached by the Kaminka and Kraus research team is unanimous: coordinated movement is preferable.

In a scenario where the robots' movements are not coordinated, the distances between the robots will not remain constant. Thus, once a distance of more than 1,200 meters is formed between a pair of robots, a vulnerable spot will appear at equal distances between the two robots – a spot where penetration can succeed because even if the robots are moving towards that specific spot, they will need more than three minutes to get there (giving the intruder time to get through). Therefore, all the intruder needs is to wait patiently until such a situation occurs. In order to prevent a security breach based on exaggerated distance between pairs of robots, they need to move in coordination with each other.

Another reason that coordination between robots is necessary is that one cannot assume that the robots' velocity is constant and accurate. Inter-robot communication would allow every robot to report

its location, and to maintain an appropriate distance between robots. It would also make it possible to discover malfunctions (if, for example, one of the robots ceases to transmit), and re-coordinate the optimum distance that should be maintained between robot guards when there is one less robot on patrol. While a reduced number of robot guards would obviously be a positive development for any would-be intruder, it would not necessarily help him break through the fence. This is because the event would not be predictable, and the defenders would likely hurry to repair the robot before he could take advantage of the situation. The researchers have developed methods for maintaining a constant optimum distance between robots, as well as methods for coping with robot malfunction.

One may assume that the intruder has all the knowledge he or she needs. However, they don't know what the defenders themselves don't know: when a random reversal of direction will occur.

Another important question that must be asked is: is it better for the robots to move in the same direction all the time, or should they occasionally change the direction of their movement at random? Obviously, limiting the movement to one direction is not recommended; in this scenario, all an intruder would need would be to allow the robot to pass by, and then he could break through the fence, confident that he has almost five minutes before the next robot arrives.

This is the basis of an axiom widely accepted in the world of security – that one must assume the intruder has all possible knowledge about the protection strategy. If this is the assumption, and since there aren't enough robots to cover every spot along the fence in frequency of more than once every five minutes, the guards will be forced to randomly change directions. While the intruder has all the knowledge he needs including probability of a change of direction, he doesn't know what the

defenders themselves don't know: When exactly will these changes in direction occur?

This is why using random decisions is so important. An intruder operating in this environment can never be certain that the robot passing him won't "decide" at the next moment to retrace its steps, because such a "decision" is completely random.

As we have already seen, the robots should always maintain a constant distance between them. This means that the random changes must also occur simultaneously by all the robots. The need for coordinated movement and direction changes highlights one of the advantages of robots: unlike human guards, robots are able to maintain a coordinated choreography.

Incidentally, the strategy described here is not applicable for a central command scenario, in which one of the robots "rolls a dice" about how to move, and informs the others of his decision. This is because, if this commanding robot breaks down, the entire system collapses with it. Therefore, a decentralized mechanism of command, control and coordination is required so that a shutdown of some of the robots still allows the other robots to cooperate between themselves and re-divide the missions. This type of interaction is characteristic of all "multi-agent environments."

Rock-Paper-Scissors

How should the robots determine the best probability for changing directions? The MAVERICK team has studied this challenge, as well.



Let's say a robot "rolls a dice" every minute to decide whether or not to retrace its

»

>>

steps. Should he change direction only if the dice shows six (probability of 1/6), or when the dice shows five or six (probability of 1/3), or any other probability? Choosing the probability is performed using a method known as “max-min” (maximum-minimum) decision-making.

First, when the probability to change direction is given, it is compared to a calculation of the probability of detection for every spot along the fence – in other words, the chances of a robot breaking through each particular spot within the next three minutes. This calculation takes into consideration the location of two robots situated to the left and right of the same spot, when one of them is approaching it and the other is moving away, as well as the probability that the direction of the robots’ movement will change at a certain moment during the next three minutes. In the second stage of this calculation, the robot determines the spot where the probability of detection is minimal. This is the ultimate spot for the intruder (don’t forget that the intruder can also perform these calculations, according to the assumption that he knows the probability of the robots changing direction).

These first and second stages are repeated for all possible values of the probability of shifting direction. For each value the robot will find the detection probability should the intruder attempt to penetrate at a particular spot.

In the third stage the robot shall choose the shifting direction probability so that the detection probability will be maximal. In other words, the robot will determine the intruder’s ultimate spot, then reduce as much as possible his chances for breaking through there.

Under the conditions described above, it can be proven that this calculation method provides the best

If the intruder doesn’t know anything about the guards’ movement, a constant direction movement is preferable.

management of robot guard movement. This is true even when taking into consideration the delays caused by every direction shift. This is compatible with our intuitive assumption that it is better to shift direction randomly, in order to surprise the opponent.

The Uncertainty Scenario

At this point the MAVERICK lab team raised a new question: what happens if the intruder doesn’t really know everything there is to know about the movement of the guards? And what happens if such an intruder chooses to make his move at a random spot, anywhere along the fence?

To many people working in this field, the answer to this question is counter-intuitive: if the intruder doesn’t know anything about the movement of the guards, it is better for the guards to employ a constant direction of movement. This is because, when the robot shifts directions, he re-scans a part of the fence that he has just scanned a few minutes earlier. The outcome is that – for a short time period – certain parts of the fence are “guarded better,” meaning they are scanned more often in comparison to other parts.

Should the intruder choose one of the spots that are left less closely guarded as a result of these direction shifts, his success is almost guaranteed. In this scenario the intruder has an equal probability of choosing every spot along the fence; therefore, the guards must guarantee that all spots are properly and equally guarded. This aim is achievable only if the robots never change direction. Thus we have seen that the best protection method against a “smart” intruder is not the best method against an “ignorant” one.

This finding leads to a situation reminiscent of the children's game "Rock-Paper-Scissors": we can only choose the best method for dealing with an opponent if we know how the opponent behaves. This comparison is not airtight: as opposed to the children's game, when attempting to breach a fence, it's not advisable for a "smart" intruder to decide to behave like an "ignorant" intruder even if he comes across a guarding method that includes random direction shifting (a situation in which the detection probability is different in every spot, making it better to choose a spot with relatively low detection probability). However, on the defending side, there isn't a single optimum strategy for all kinds of intruders. Therefore, the guard is required to perform an educated assessment of the nature of the opponent.

Experiments in Virtual Defense

In order to test how various protection methods can be coordinated with different types of intruders, the researchers at MAVERICK conducted an experiment based on a simple computer game. The participants watch a circular fence on a computer screen and are asked to choose the point at which one should attempt to break through. In the first stage of the game, the participant sees the location of the guards as an immobile photograph, so he is not able to predict their behavior. In the second stage, the participant is asked to watch the robots' movement for a few seconds before choosing what he considers the optimum penetration spot. In the third stage, the participant observes the robots' movement for three minutes – during that time he can choose several points at which to infiltrate (see following link: <http://www.cs.biu.ac.il/~sadvov/fence.swf>).

The first stage of the game casts the participant in the role of an intruder with limited knowledge: he knows the actual location of the robots, but has no way of predicting their behavior from the moment they

start to move. The second and third stages allow the intruders to acquire more and more knowledge concerning the robot guards' movements.

In setting up the game, the MAVERICK researchers posited three methods for protecting the compound. In the first, "deterministic" method, there was no shift of direction, which meant that the success rates for an intruder with a lack of knowledge were extremely low. In the second method, the robots engaged in movement including random direction shifts chosen by the "max-min" probability method. This reduced, as much as possible, the success rates of a knowledgeable intruder who chooses a spot in hopes of maximizing his probability of success. The last method was devised to include movement with random direction shifts in lower probability, as a way to minimize the difference in success rates that would be experienced when facing different types of intruders.

As predicted, the deterministic model of movement proved the best protection strategy for the first stage of the experiment. For the second stage, the movement with random direction shifts was to some extent preferable, by considering both types of intruders; and for the third stage, the "max-min" method was most effective for intruders with a maximal knowledge.

This experiment demonstrated that it is possible to describe the behavior of human intruders according to the simple models, which are then used in the planning of robot movement control. However, we still need to set policy by determining which type of intruder the robots are likely to face, and which defensive strategy is most appropriate. Recent work by recently graduated PhD student Noa Agmon, advised by Kaminka and Kraus, has tackled this challenge directly. Dr. Agmon has developed algorithms that consider the uncertainty faced by the intruder when observing the guarding robots, and the effects of this

>>

uncertainty on the intruder's decision on where to attempt crossing.

What Happens at the End of the Fence?


The scenarios described above all concern a circular fence. The MAVERICK research team has also developed strategies for planning robot patrols along an “open” fence – a fence that is not a closed circle but rather, a connecting line between two end points.

When a robot (or a human guard) reaches the end of such a fence, he must turn around and go back. At that moment, he will return to cover spots visited only a short time earlier. Therefore, in the transition between a circular fence and an open one, it becomes more difficult to maintain a unified scanning frequency, as the sections of the fence closest to the ends are necessarily scanned at a different rate than the middle section. For instance, if the robot moves from one end to the other in ten minutes, then it will visit the spot in the middle of the fence every ten minutes, but will visit a spot near one of the ends at drastically different intervals, for example, two minutes between visits, followed by an 18 minute gap before the next visit occurs. If we have at our disposal three robots guarding a three km fence, the simplest solution is to divide the

fence into three sections of one km each: if we mark both ends of the fence with the letters A and D, then the first robot will guard section AB, the second will guard section BC, and the third will guard section CD (Illustration 1). Even for this type of condition, the MAVERICK researchers show that a coordinated and synchronized movement strategy is preferable to a program where each robot moves independently without coordination. The fact is, even if the length of the sections being patrolled is equal, and the robots' velocity is equal, as well, coordination is necessary in order to prevent the accumulation of small delays over time.

Even with coordinated movement, there are disadvantages to this system. The above-mentioned problem of a varying

Illustration 2




Three robots guarding a section of a fence. Each robot is assigned to a section that overlaps with the others.

time-gap between visits to spots not located in the middle of the fence is exacerbated when each section exclusively “belongs” to a different robot – because each section has its own middle and end points. Another disadvantage to consider is the increased time lost during direction shifts – when the robot turns around to change directions, the fence is not being scanned.

Responding to these disadvantages, recently graduated PhD student Yehuda Elmaliach, advised by Kaminka, suggested dividing the fence into

Illustration 1



Three robots protecting a section of the fence. Each robot was assigned its own section of the fence to guard.

overlapping sections (Illustration 2). In this method, the first robot will move between A and C, the second robot between B and D, and the third robot between C and D. Since the movement is coordinated, the third robot will finish the movement before the others, and wait for them in spot D until they finish scanning longer sections.

At first glance it might appear that this strategy “wastes” the potential of the third robot, because it remains in one spot for half the time. Nevertheless, it appears that some of the disadvantages associated with the simple division strategy are lessened here. Direction shifts occur less often, so less time is wasted. At the same time, the uniformity of fence scanning, and the frequency with which each spot is visited, is improved.

Clearly this improvement is achieved for the section between B and D, at the expense of section AB. As the fence is divided into a larger number of sections, the increase of guarding efficiency is correspondingly greater. This also makes it possible to consider greater overlap between the robots’ individual sections. MAVERICK researchers developed detailed mathematical models explicating the advantages and disadvantages associated with each of these possibilities, and recommend increasing the protection of section AB by placing a fixed camera in spot A.

In order to examine whether these models are suitable to handle setbacks and errors that might occur in the real world, the laboratory examined the predictions of the models in comparison to the experiments conducted on actual robots. Simple robots – originally designed for vacuum cleaning – were adapted for the experiments. The robots were programmed to operate according to both methods presented here, and the uniformity of the visitation


rates in each spot was measured. The experimental results perfectly matched the prediction given in the mathematical models.

Theory and Practice

MAVERICK research focuses on important theoretical challenges in artificial intelligence. However, the lab’s scientific team insists that theoretical studies should be integrated into practical applications based on real-world problems. In a society like Israel, where security concerns make border and fence patrols an everyday fact of life, a model like this may lead to more efficient use of military and civilian resources, and better protection for the robots’ human handlers. ❖

For additional reading: See the MAVERICK website: <http://u.cs.biu.ac.il/~maverick>

Prof. Gal Kaminka heads Bar-Ilan University's MAVERICK group, which conducts research in autonomous agents and multi-agent systems, robotics, and state-of-the art technological opportunities. He



previously served as an Adjunct Assistant Professor at Carnegie Mellon University, where he conducted his postdoctoral research. Prof. Kaminka, who received his PhD in Computer Science from the University of Southern California, has won a number of awards and is a member of the Association for Advancement of Artificial Intelligence (AAAI). ❖

»

>>
>>

Computerized "ARMOR" for LAX Security: Bar-Ilan Professor's Anti-Terror Shield Cited for Excellence by LA World Airports Police Division



Los Angeles International Airport is a prime target for terrorists, who would like nothing more than to perpetrate a major attack on the West Coast of the United States. Now, breaking through the airport's anti-terror defenses has been made significantly harder, thanks to a computerized system designed in collaboration with researchers from the University of

Southern California and Prof. Sarit Kraus, an expert in artificial intelligence from the Department of Computer Science at Bar-Ilan University, and a member of the University's *Leslie and Susan Gonda (Goldschmied) Multidisciplinary Brain Research Center*.

Prof. Kraus recently received a special citation from the city of Los Angeles, World Airports Police Division, for her contribution to overall airport security at Los Angeles International Airport. The system – called "ARMOR," an acronym for Assisted Randomized Motoring Over Routes – was the brainchild of an international team headed by USC Professor Milind Tambe, together with Kraus, as well as USC Prof. Fernando Ordonez and students. Implemented and used at the airport since November 2007, ARMOR has resulted in significantly higher success rates for security-related tasks, including the location and seizure of hidden weapons and narcotics, and the identification of suspicious individuals. Citation signatories, the Deputy Executive Director and Chief of Police of the Los Angeles World Airports Police Division, wrote of Prof. Kraus: "To merit this

commendation you have performed an exceptional service to the Airport Police Division, the Los Angeles World Airports and the City of Los Angeles. Your outstanding service facilitates the critical link between the laboratory and the operational world. Thank you for your outstanding contribution to the security of our nation."

The performance of the ARMOR system has been noted throughout the homeland security community. It has been presented as a model for emulation to the United States' Transportation Security Administration, the Joint Chiefs of Staff Level IV Antiterrorism Seminar, and the full Congressional Committee on Homeland Security, as well as several countries around the world.

Randomization and the "Rational" Terrorist

The ARMOR system is based on the randomization of countermeasures – those activities through which security forces attempt to foil the efforts of would-be terrorists. Targeting suspicious activities that may take place well before an attempted attack, the goal of countermeasures is to reveal terrorists' efforts to gather information about airport vulnerabilities and security protocols.

"Unlike Ben Gurion Airport, The Los Angeles International Airport is serviced by several major access roads, and there aren't enough resources to stop and check every car," says Prof. Kraus. "Similarly, the resources required for sniffing out drugs are also limited – dogs and their handlers can't cover every corner of every terminal at all times. The answer is smart randomization – setting up checkpoints and canine patrols in an unpredictable pattern that provides sufficient protection for the most vulnerable areas, while preventing terrorists and drug dealers from predicting where the next inspection will take place."

>>

Security Through “Games Theory”

According to Prof. Kraus, one of the advantages of the system is its theoretical basis in games theory – an area of science that focuses on how perceived advantage affects human decision-making. “ARMOR provides a randomized schedule of where to set up check points, and to send canine patrols, but security officers also have the ability to override and change this automated schedule,” she explains. “If an officer makes these changes in a non-randomized, predictable pattern – a pattern that could potentially tip off terrorists or drug dealers – a warning is issued by the system. This improves security, and promotes trust of the system by its human users.”

In her current research, Prof. Kraus is working with the USC team on scenarios in which a potential attacker has only limited knowledge about the security protocols and vulnerabilities of a particular target site. “Our current model focuses on an opponent that has complete knowledge about how an airport – or other site – is being protected, and acts rationally, based on this information. If we can expand this model to include the attacker whose actions are less firmly based on the weighing of rational options, we will be able to improve security still further.”

Prof. Sarit Kraus, an internationally acknowledged authority on multi-agent systems for artificial intelligence, conducts joint research with colleagues at Harvard University, the Stanford Research Institute, Tel-Aviv University, Hebrew University, the University of Southern California, and the University of Liverpool. She is a Professor of Computer Science at Bar-Ilan University and a member of Bar-Ilan’s Leslie and Susan Gonda (Goldschmied) Multidisciplinary Brain Research Center. Author of Strategic Negotiation in Multiagent Environments (2001) and co-author of Heterogeneous Active Agents (2000), both published by MIT Press.

Prof. Kraus has published over 200 papers in leading journals. Among her many honors, Prof. Kraus received the 2007 ACM/SIGART Autonomous Agents Research Award for her pioneering work on the development of techniques for computational negotiation and automated coalition formation. In 2008 the European Coordinating Committee for Artificial Intelligence named her a ECCAI Fellow, a program founded in order to recognize individuals who have made significant, sustained contributions to the field of artificial intelligence. ❖



shutterstock/JASAP