Anomaly Based Intrusion Detection and Artificial Intelligence

Benoît Morel
Carnegie Mellon University
United States

1. Introduction

"The internet can be regarded as the most complex machine mankind ever built. We barely understand how it works, let alone how to secure it" [Schneier 2008]. The introduction of new technologies like the proliferation of new web applications or the increasing use of wireless, have exacerbated this fact. Cybersecurity, a spin-off of the phenomenal growth of the internet, has probably become the most complex threat to modern societies. The development of cybersecurity has been reactive and driven by the ingenuity and imagination of cyberattackers. In the words of Carl Landwehr in IEEE security and Privacy (Landwehr 2008), "defense has consisted in "fixing the plumbing". What we need is to put more Artificial Intelligence (AI) in cybersecurity". This is the theme of this chapter.

Cyberspace is a rather brittle infrastructure, not designed to support what it does today, and on which more and more functionality is build. The fact that the internet is used for all sorts of critical activities at the level of individuals, firms, organizations and even at the level of nations has attracted all sorts of malicious activities. Cyber-attacks can take all sorts of forms. Some attacks like Denial of Service are easy to detect. The problem is what to do against them. For many other forms of attack, detection is a problem and sometimes the main problem.

The art of cyber-attack never stops improving. The Conficker worm or malware (which was unleashed in Fall 2008 and is still infecting millions of computers worldwide two years later) ushered us in an era of higher sophistication. As far as detection goes, Conficker in a sense was not difficult to detect as it spreads generously and infected many honeypots. But as is the case for any other new malware, there are no existing tool which would automatically detect it and protect users. In the case of Conficker, the situation is worse in the sense that being a dll malware, direct detection and removal of the malware in compromise computers is problematic. One additional problem with Conficker is the sophistication of the code (which has been studied and reverse engineered ad nauseam) and of the malware itself (it had many functionality, was using encryption techniques to communicate (MD6) which had never been used before). It spreads generously worldwide using a variety of vectors, within networks, into a variety of military organizations, hospitals etc...). In fact the challenge became such that the security industry made the unprecedented move of joining forces in a group called the Conficker working group. The only indication that this approach met with some success is that even if the botnet that Conficker build involves millions of infected computers, that botnet does not seem to have been used into any attack, at least not yet....

Conficker is only but one evidence that cyber-attackers have reached a level of sophistication and expertise such that they can routinely build malware specifically for some targeted attacks (against private networks for example), i.e. malware that are not mere variations of a previous one. Existing tools do not provide any protection against that kind of threat and do not have the potential to do so. What is needed are tools which detect autonomously new attacks against specific targets, networks or even individual computers. I.e. what is needed are intelligent tools. Defense based on reactively protecting against the possibility of a re-use of a malware or repeat of a type of attack (which is what we are doing today) is simply inadequate.

With the advent of the web, the "threat spectrum" has broadened considerably. A lot of critical activity takes place through web application. HTML, HTTP, JavaScript among others offer many points of entry for malicious activity through many forms of code injections. Trusted sessions between a user and a bank for example can be compromised or hijacked in a variety of ways.

The security response against those new threats is tentative and suboptimal. It is tentative in the sense that new attacks are discovered regularly and we are far from having a clear picture of threat spectrum on web application. It is suboptimal in the sense that the "response" typically consists in limiting functionality (through measure such as "same origin policy", for example), or complicating and making more cumbersome the protocol of trusted session in different ways. The beauty and attraction of the web stem from those functionalities. This approach to security potentially stifles the drive for innovations, which underlie the progress of the internet.

Cybersecurity is a challenge, which calls for a more sophisticated answer than is the case today. In this chapter, we focus on intrusion detection. But there is a role for Artificial Intelligence practically everywhere in cybersecurity,

The aspect of the problem that Intrusion Detection addresses is to alert users or networks that they are under attack or as is the case with web application may not even involve any malware but is based on abusing a protocol. What kind of attributes should an Intrusion Detection System (IDS) have to provide that kind of protection? It should be intelligent, hence the interest in AI.

The idea of using AI in intrusion detection is not new. In fact it is, now decades old, i.e. almost as old as the field of intrusion detection. Still today AI is not used intensely in intrusion detection. That AI could potentially improve radically the performance of IDS is obvious, but what is less obvious is how to operationalize this idea. There are several reasons for that. The most important one is that AI is a difficult subject, far from mature and only security people seem to be interested in using AI in intrusion detection. People involved in AI seem much more interested in other applications, although in many ways cybersecurity should be a natural domain of application for AI. The problem may lie more with cybersecurity than the AI community. Cybersecurity projects the impression of a chaotic world devoid of coherence and lacking codification.

As a result of that situation, most of the attempts to introduce AI in intrusion detection consisted in trying to apply existing tools developed or used in AI to cybersecurity. But in AI tools tend to be developed around applications and optimized for them. There are no AI tools optimized for cybersecurity. AI is a vast field which goes from the rather "primitive" to the very sophisticated. Many AI related attempts to use AI in cybersecurity, were in fact using the more basic tools. More recently there has been interest in the more sophisticated approaches like knowledge base approach to AI.

In the spirit of the Turing test (Turing 1950), it is tempting to define what an AI based intrusion detector should accomplish, is to replicate as well as possible what a human expert would do. Said otherwise, if a human expert with the same information as an IDS is able to detect that something anomalous/ malicious is taking place, there is hope that an AI based system could do the job. Since cyber attacks necessarily differ somehow from legitimate activities, this suggest that an AI based detector should be also an anomaly-based detector, whatever one means by "anomaly" (we elaborate on that later in this chapter). A closer look at the comparison between human beings and machine suggests that there are irreducible differences between the two which translate in differences in the limit of their performance. Human beings learn faster and "reason" better. But those differences do not go only in favor of the human: machines compute faster and better...

Today's AI based IDS's are very far from the kind of level of performance that makes such comparisons relevant. To provide adequate protection to the increasing level of functionality and complexity that is happening in the internet, the AI systems involved in cybersecurity of the future would have to be hugely more sophisticated than anything we can imagine today, to the point of raising the issue of what size they would have and the amount of CPU they would need. Is it possible to conceive a future cyberworld where so much artificial intelligence could coexist with so much functionality without suffocating it? The answer has to be yes. The alternative would be tantamount to assume before trying that AI will be at best a small part of cybersecurity. Where would the rest, the bulk of cybersecurity come from?

In fact there is a precedent: the immune system. The immune system co-evolved with the rest of biological evolution to become a dual use (huge) organ in our body. There are as many immune cells in our body as nervous cells ($\sim 10^{12}$). The human body is constantly "visited" by thousands of "antigens" (the biological equivalent of malware) and the immune system is able to discriminate between what is dangerous or not with a high degree of accuracy. In the same way one could envision in the long run computers being provided with a "cyber-immune system" which would autonomously acquire a sense of situational awareness from which it could protect the users. This is at best a vision for the long run. In the short run, more modest steps have to be made.

The first detection of any attack is anomaly-based. Today most if not all of the time the anomaly-based detector is a human being. The interest in anomaly-based detection by machines has an history which overlaps the history of attempts of introducing AI in cybersecurity. In fact most of the attempts to introduce AI in intrusion detection was in the context of anomaly-based detection.

Basically all new attacks are detected through anomalies, and in most cases they are detected by human beings. Considering the variety of forms that attacks can take, it is rather obvious that anomalies can take all sorts of forms. Anomaly based Intrusion Detection has been a subject of research for decades.. If it has failed to deliver a widely used product, this is not for lack of imagination of where to look to find anomalies. One of the most promising attempts, which had an inspirational effect on the research in that field, was to use system calls.

The nemesis of anomaly-based detection has been the false positive. A detection system cannot be perfect (even if it uses a human expert). It produces false positive (it thinks it has detected a malicious event, which in fact is legitimate) and has false negative (it fails to detect actual malicious events). Often there is a trade-off between the two: when one puts the threshold very low to avoid false negative, one often ends up with a higher rate of false

positive. If a detector has a false positive probability of 1%, this does not imply that if it raises a flag it will be a false alert only 1% of the time (and 99% probability that it detected an actual malicious event). It means that when it analyzes random legitimate events 1% of the time it will raise a flag. If the detector analyzes 10,000 events, it will flag 100 legitimate events. If out of the 10,000 events one was malicious, it will raise an additional flag, making its total 101. Out of the 101 events detected, 1 was malicious and 100 were legitimate. In other words, out of the 101 alerts only one is real and 100 out of 101, i.e. more than 99% of the time the alert was a false positive.

Those numbers were illustrative but taken totally by chance. 1% is a typical performance for "good" anomaly based detection systems thus far proposed. The actual frequency of malicious activity in the traffic (if one neglects spam) is not precisely known, but malicious events are relatively rare. I.e. they represent between 0 and maybe 10^{-4} of the traffic. Before anomaly-based detection can be considered operational, one has to find ways to reduce the probability of false positive by orders of magnitude. It is fair to say that we are at a stage where a new idea in anomaly-based intrusion detection, inspired by AI or anything else, lives or dies on its potential to put the false positive under control. In this chapter, two algorithms or mechanisms are offered which can reduce the probability of false positives to that extent: one uses Bayesian updating, the other generalizing an old idea of von Neumann (von Neumann 1956) to the analysis of events by many detectors.

Those two algorithms represent the "original" or technical contributions of this chapter, but this chapter is also concerned more generally by the interface between AI and cybersecurity and discusses ways in which this interface could be made more active.

2. Framing the problem

a. The new Threat environment

The "threat environment" has evolved as has the art of cyber-attack. Buffer overflow vulnerabilities have been known for a long time - the Morris worm of 1988, that for many was the real beginning of cybersecurity, exploited a buffer overflow vulnerabilities. They became a real preoccupation a few years later and progressively people realize that most software written in C have exploitable buffer overflow vulnerabilities.

Buffer overflows are still around today. Although they have not been "solved" they now represent only one class in what has become a zoology of exploitable vulnerabilities. In most cases after those vulnerabilities are discovered, the vendor produces a patch, which is reverse engineered by hackers and an exploit is produced within hours of the release of the patch... Many well-known malware (Conficker is an example) exploit vulnerabilities for which there is a patch. They use the fact that for a variety of reasons, the patch is not deployed in vulnerable - of such attacks, where the attacker discovers the vulnerability before the vendor and susceptible computers are helpless. The attack in the fall 2009 against Google and a few more companies originating in China, called Aurora, was an example of an exploitable dangling pointers vulnerability in a Microsoft browser, that had not been discovered yet.

A good defense strategy should rely on the ability of anticipating attacks and produce patches in time. A really good defense system should be able to protect computers from the exploitation of yet undiscovered exploitable vulnerability.

With advent of the web new classes of vulnerabilities emerge. Some websites are not immune against code injection, which can have all sorts of implications. Some website are

vulnerable to Java-script instructions. This can be used for a variety of purpose, one being to compromise the website and makes its access dangerous to users. Protecting websites against all forms of code injection is easy in the case where it does not involve a lot of functionality. But interactive websites providing a lot of functionality are far more difficult to protect against every possible scenario of attack.

In the case of web application security, the Browser plays a central role. The interaction between users and servers go through the Browser, which in principle sees everything. In practice browsers have some security embedded in them, but not of the kind that could alert the user that he is victim of a cross site request forgery (CSRF) attack, for example. A really good defense system would be able to achieve a degree of situational awareness of what is taking place within the browser to detect that kind of attack and other forms of attack.

b. What are anomalies

The concept of anomalies is problematic, as is their relation with malicious activities (Tan and Maxion, 2005). By definition an anomaly is a "rare event", in other words, the concept of anomaly is statistical in nature. A noteworthy attempt to define anomaly was the idea of S. Forrest et al to make statistics of system calls (Hofmeyr et al. 1998). The idea was inspired by the concept of self and non-self ion immunology. The building blocks of proteins and antigens are amino acids. There are about 20 of them, some more essential than others. This means that there is an enormous variety of sequence of amino acids. Antigens are recognized by the immune systems as "non-self", i.e. having sequences that are not represented in the body. In principle the immune system attacks only the tissues which are non-self (This is what happens in the rejection of transplants). Auto-immune diseases would represent the "false positive" and they are relatively very rare. What is remarkable is that the distinction self non-self in immunology is based on short sequences (typically 9) of amino acids, called peptides.

The idea is that users can be recognized by the statistics of system calls, and that the equivalent of peptides would be short set of successive system calls. The number six (Tan and Maxion 2002) turned out to be "optimum". In that approach one can choose to define what is "anomalous", through its frequency of occurrence: 1%, 0.1%, .. The connection between abnormality and maliciousness is based on assumptions.

One advantage of this approach is that every user is supposed to be different. That puts potential attackers in situation of added complexity as it is difficult for them to fool many users with the same attack at the same time.

Among the other obstacles in using this approach is the fact that users change habits, the concept of what is normal is not constant. and that can potentially be exploited through so-called "mimicry attacks", i.e. manipulation of the concept of normality by a shrewd attacker. The fact that in modern computers there is a lot of activity taking place in the background, out of the control of the user introduces an additional noise. Furthermore that kind of approach has limited use for web security. In the context of web applications, the information to analyze statistically is buried in the set of HTTP requests that reach and are conveyed by the browser.

3. Review of previous relevant work

One can find many papers dealing with intrusion detection and using the word "AI" in their title. By AI, often is meant data mining, neural network, fuzzy logic (Idris et al 2005),

Hidden Markov Model (Choy and Cho, 2001), self-organizing maps and the like. Considering that all these papers deal with anomaly-based intrusion detection, the key figure of merit to gauge their contribution is whether their approach has the potential to tame the false positives. Those papers are remotely related to this chapter, as the problem of false positives is not as central and unlike this chapter, in those papers the machine learning and Knowledge base aspects of AI are not as prominent as in the discussion of this chapter. A lot but not all of the AI "machinery" is statistical (Mitchell 1997) in nature (and therefore is threatened by the curse of the false positives... There is branch of AI concerned by "reasoning" (Brachman et al. 2004, Bacchus et al 1999, Baral et al. 2000), making context dependent decision and the like. Among the papers dealing with AI in the context of intrusion detection, the paper of Gagnon and Esfandiari 2007 is probably the closest to this chapter. Its discussion is in fact less general than this chapter and is organized around a very specific use of a Knowledge based approach to AI. The discussion illustrates the challenges in trying to use sophisticated AI techniques in cybersecurity.

4. Reducing the false positives using Bayesian updating

As stated in the introduction the nemesis of anomaly based IDS systems is the probability of false positive. When the probability that an event is malicious does not exceed 10^{-4} , the probability of false positive should be less than that.

Little or no thought has been put in exploiting the fact that a cyber-attack is in general a protracted affair. In the same way that a human expert monitoring an suspicious events would see whether the evidence that what he is witnessing is indeed an attack or not, an IDS system could make a more protracted analysis of suspicion before raising a flag, thereby reducing the probability of false positive.

We sketch here how the math of such an iterated procedure would work, starting by spending some time defining what false positive means. It can mean more than one thing...

Let the Boolean variable ζ refer to whether one deals with a malicious event or not. By definition: $\zeta = 1$ means that the event is malicious. Otherwise $\zeta = 0$. The variable of interest is: $P(\zeta = 1)$, the probability that it was a malicious event. All the paraphernalia of data, measurements and detection, can be represented by another Boolean variable X . By definition $X = 1$ means that there is evidence for something malicious, i.e. something abnormal.

The famous Bayes theorem states that:

$$P(X = 1, \zeta = 0) = P(X = 1 | \zeta = 0)P(\zeta = 0) = P(\zeta = 0 | X = 1)P(X = 1) \quad (1)$$

In EQ1 there are three probabilities, which can be referred to as "false positive", but should be distinguished:

$P(X = 1, \zeta = 0)$ is the probability that, an attack is being detected while in fact no attack took place.

$P(X = 1 | \zeta = 0)$ is the conditional probability that even if there is no attack, the system of detection will detect one.

$P(\zeta = 0 | X = 1)$ is the conditional probability that when there is evidence of an attack, in fact this is a false alert.

From EQ 1, it is clear that they are three different numbers.

The conditional probabilities $P(X = 1 | \zeta = 0)$ and $P(X = 0 | \zeta = 1)$ are figures of merit of the detection system. They determined whether or not the information generated by the detection system should be used or not. The number of interest is: that an attack is taking place.

What is referred to as “false positive” in this chapter is $P(X = 1 | \zeta = 0)$, i.e. it is an attribute of the detection system. In the same way $P(X = 0 | \zeta = 1)$ represents the false negative, also an attribute of the detection system

One can then use the fundamental assumption underlying the so-called “Bayesian updating”: if at a given time the probability that there is a malicious event is $P(\zeta = 1)$, then after a new measurement where X is either 1 or 0, the new value of $P(\zeta = 1)$ is:

$$\tilde{p}(\zeta = 1) = \{X p(\zeta = 1 | X = 1) + (1 - X) p(\zeta = 1 | X = 0)\} \quad (2)$$

In order to have this expression in terms of “false positive” and false negative”, we rewrite EQ. 2, using EQ.1, as:

$$\tilde{P}(\zeta = 1) = \left\{ \frac{(1 - X) P(X = 0 | \zeta = 1)}{P(X = 0)} + \frac{X P(X = 1 | \zeta = 1)}{P(X = 1)} \right\} P(\zeta = 1) \quad (3)$$

$\vartheta = P(\zeta = 1)$ is a dynamical variable. Each time a measurement is made, the value of $\vartheta = P(\zeta = 1)$ is updated into $\tilde{\vartheta}$:

$$\begin{aligned} \frac{\tilde{\vartheta}}{\vartheta} = & \frac{(1 - X) P(X = 0 | \zeta = 1)}{\vartheta P(X = 0 | \zeta = 1) + (1 - \vartheta) P(X = 0 | \zeta = 0)} + \\ & + \frac{X P(X = 1 | \zeta = 1)}{\vartheta P(X = 1 | \zeta = 1) + (1 - \vartheta) P(X = 1 | \zeta = 0)} \end{aligned} \quad (4)$$

To show the potential power of using Bayesian updating, let assume that as a prior we take $\vartheta = P(\zeta = 1) \approx 10^{-4}$. We also assume that the detection system has 1% false positive ($P(X = 1 | \zeta = 0) = 0.01$), we also assume that $P(X = 1 | \zeta = 1) = 0.99$, and consistently in EQ.4 each measurement is suspicious, i.e: $X = 1$. The evolution of the value of $\vartheta = P(\zeta = 1)$ is shown in Figure 1. It takes 4 successive evidences of suspicion to put the probability that there is a malicious activity close to 1. The probability that the detector will make 4 mistakes in a row (if there is no correlation) is $(10^{-2})^4 = 10^{-8}$.

The possibility of using Bayesian updating in the context of anomaly-based detection has not yet been seriously contemplated. This is only one avenue toward making an AI based systems much less prone to false positives. Another is using several computers networked together.

Fig. 1. Evolution of $P(\zeta = 1)$ through Bayesian updating, using EQ.4 starting at $P(\zeta = 1) = 10^{-4}$, assuming $P(X = 1|\zeta = 0) = 0.01$ and $P(X = 1|\zeta = 1) = 0.99$ and assuming that at each measurement $X = 1$.

5. Reducing the false positives using networked computers

Another avenue, which offers a lot of promises too, is using the observation that what one computer may have difficulty to do, several computers networked intelligently could.

John von Neumann (von Neumann 1956) wrote a paper entitled "Probabilistic logics and the synthesis of reliable organisms from unreliable components", which supports this notion. The paper, which culminated several years of study was not about anomaly-based intrusion detection, but understanding how the brain works. The goal was to show how a logical system can perform better than its component and thereby establish some foundations for AI.

A way to interpret some of the results of von Neumann is that it is possible if one has a system involving a large number of components, to combine the components in such a way that they build a kind of information processor such that the resulting uncertainty on the outcome can in principle be made arbitrarily small if the number of components can be large enough.

Ostensibly the paper of John von Neumann (von Neumann 1956), addresses the question of how to reduce the error due to unreliable components to an arbitrary small level using multiplexing and large numbers. In practice, the ideas developed in that paper have the potential to be applied to a large variety of problems involving unreliable components and we think among others the problem of early detection of new malware. Here we described succinctly some relevant observations of von Neumann.

a. Logical 3- gates

A majority rule 3-gate receives information from three sources. The probability that the gate yields a false information is the probability that at least two of the three sources were providing a false information. If χ_i is the probability that line "i" gives a false positive, the probability that at least two of the three incoming lines give a wrong information and that the gate is sending a false positive signal is:

$$\pi_g = \chi_1\chi_2(1-\chi_3) + \chi_1\chi_3(1-\chi_2) + \chi_2\chi_3(1-\chi_1) + \chi_1\chi_2\chi_3 \quad (1)$$

Or equivalently:

$$\pi_g = \chi_1\chi_2 + \chi_1\chi_3 + \chi_2\chi_3 - 2\chi_1\chi_2\chi_3 \quad (2)$$

If one assumes that $\chi_i \approx 10\%$, then the probability of false positive of the system made of three detectors, feeding on a majority 3-gate will be $\pi_g \approx 3\%$ (Cf Figure 2).

Fig. 2. Output of majority rule gates: The green curve is for the case with three detectors assuming that: $\chi_1 = \chi_2 = \chi_3 = \xi$, i.e. that: $\pi_g = 3\xi^2 - 2\xi^3$. In that case: $\pi_{FP}^3 = 3\xi^2 - 2\xi^3$. The two other curves are for the case where there are nine detectors. The red curve corresponds to the simple majority rule π_{MR}^9 , the other one (blue) corresponds to the case where the nine detectors are distributed in three majority 3 rules feeding a majority 3 rule. I.e. it corresponds to: π_{FP}^9 .

b. With 3 N computers Logical 3-gates

Grouping the signal emanating from detectors in three and make them feed a majority rule gate would produce an aggregate with a somewhat improved probability of false positive (and this can be used for the false negative too).

For illustration let us assume that the number of detectors is nine, In the first scenario (construct of majority 3-gates), the probability π_{FP}^9 of false positive that nine computers (each with the same probability of false positive ξ) feeding three majority rule gates (each gate has a false positive probability $\pi_{FP}^3 = 3\xi^2 - 2\xi^3$), is therefore:

$$\pi_{FP}^9 = 3(\pi_{FP}^3)^2 - 2(\pi_{FP}^3)^3 = (3\xi^2 - 2\xi^3)^2 \{3 - 2(3\xi^2 - 2\xi^3)\} \quad (3)$$

The generalization of this formula to the case of 3N detectors is:

$$\pi_{FP}^{3N} = 3(\pi_{FP}^{3(N-1)})^2 - 2(\pi_{FP}^{3(N-1)})^3 \quad (4)$$

The speed at which the false positive rate decreases when N grows is shown in Table 1, where the individual probability of false positive is assumed to be 10% ($\xi = 0.1$). What in table 1 is referred to as $N=27$ in EQ. 4 would correspond to $3N=27$, i.e. $N=9$. Table 1 compares the situation of computers distributed into networked 3 gates, with the scenario where they build one logical N gates.

c. Logical N-gates

In this scenario (one majority rule gate), the probability of false positive π_{MR}^N has the general form:

$$\pi_{MR}^N = \sum_{i > \frac{N}{2}}^N \binom{N}{i} \xi^i (1-\xi)^{N-i} \quad (4)$$

In that scenario the overall probability of false positive decreases with N even faster than in the scenario of the networked 3-gates, as illustrated in Table 1.

When the number of computers increases, the improvement increases as well and it increases fast, in particular in the majority rule case. For example for $\xi = 0.1$:

	$\pi_{MR}^N = \sum_{i=\frac{N}{2}+1}^N \binom{N}{i} \xi^i (1-\xi)^{N-i}$	$\pi_{FP}^N = 3 \left(\pi_{FP}^{\frac{N}{3}} \right)^2 - 2 \left(\pi_{FP}^{\frac{N}{3}} \right)^3$
$N=3$	0.028	0.028
$N=9$	0.00089	0.0023
$N=27$	5.6×10^{-8}	0.0000159
$N=81$	3.5×10^{-20}	7.6×10^{-10}
Set-up	Majority rule	3-gates

Those results assume that the probabilities of false positive of the different detectors are independent. This is clearly not always the case. This idea inspired from von Neumann could benefit significantly anomaly-based network intrusion detection.

d. Operationalizing such ideas and the need for more AI

If one could exploit the full implications of Bayesian updating and/or when possible use logical N-Gates, the fact that anomaly-based detection generate intrinsically too many false positive, would not constitute an insuperable obstacle to build a full anomaly-based system.

Logical N-gates and network security:

The multi computer approach inspired from von Neumann would be appropriate for network intrusion detection. If several computers detect anomalies simultaneously and they are appropriately connected, this could lead to a powerful system of detection with few false positives and few false negatives at the same time.

Ghostnet (and its follow up "Shadows in the Cloud") refers to a Trojans which penetrated several networks associated with government agencies, most notoriously the network of the Dalai Lama in 2008 and of Indian agencies involved in national Security in 2009/2010. In both cases it was traced back to China. Ghostnet was eventually discovered when the Dalai

Lama began to suspect that his network must have been penetrated by the Chinese and asked infowar in the university of Toronto to investigate. Using honeypot they uncovered the presence of a Trojan, which was spying on the e-mails and reporting to servers scattered in the world. The investigation established that the compound of the Dalai Lama was only one of several networks that had been penetrated. A close monitoring of the traffic coming in and out of the network, by the computers of the networks, could have detected some suspicious queries. But the probability that those suspicious queries were false positive would have been large. If the evidence of those suspicions had been sent to a centralized server, by an algorithm similar to the logic N-gates scenario, it may have been able to establish the suspicion with far more certainty, much earlier.

The same kind of argument can be made about malware like Agent.btz which "traumatized" the US military and malware like Silent Banker that roam in the networks of Banks. In each case an individual computer would not be able to do a very good job at detecting a malicious activity with high level of certainty. But those malware do not infect only one computer. They need to infect quite a few, which therefore could cooperate to establish the presence of the malware.

Operationalizing Bayesian updating:

Bayesian updating is somewhat reminiscent of the implications of the observation that if one uses more than one measurement, the two best measurements may not be the best two (Cover 1970). A way to operationalize the Bayesian updating technique would be for example through a tool making periodic assessments of whether a sequence of events involves an increasing number of evidences that it is suspicious or not. For example, the tool could be embedded in the Browser of a client monitoring all the HTTP requests. If the tool detects suspicious activity it would trigger this updating procedure by analyzing subsequent events and see whether the suspicion tends to increase or not.

Ideally the tool would be designed in such a way that it would be able to "reason" about those events and analyze them. The important part here is that the tool would use a protracted analysis of the event to reach a decision about the event. Its reasoning would be probabilistic, but not necessarily statistically based.

6. Web applications

Although the web is only one aspect of the internet, web applications are becoming a dominant feature of the internet and this trend is growing. From the perspective of cybersecurity, the world of web applications is very complicated and seems to offer an infinite numbers of opportunities for abuse. Some exploitable vulnerabilities are difficult to understand or anticipate as they result from technical details of protocols, implementation of application or are consequences of abusing functionalities which otherwise are very useful or valuable (vanKesteren et al. 2008). Each time a new vulnerability is discovered, suggestions are made on how to avoid them (Barth et al. 2008b, Zeller and Felten 2008). Those suggestions are often not very attractive because they are based on reducing some functionality or they include adding complications in the implementation of applications. To the credit of system administrators, many of them spontaneously find ways to avoid potentially exploitable vulnerabilities. This is one reason why it is not so easy to find popular websites with obvious cross-site scripting (XSS) or cross site forgery request (CSRF) vulnerabilities (Zeller and Felten 2008). On the other hand, new forms of attacks appear regularly (for example "ClickJacking"

(Grossman 2008), login CSRF (Barth et al. 2008) and more will appear. Still in the same way that the semantic web is based on the culture of AI, the new level of complexity of cybersecurity accompanying this development, would benefit from relying more on AI.

a. The example of Cross Site Request Forgery (CSRF)

In a CSRF attack, the attacker manages to pose as the legitimate user to a trusted website (Zeller and Felten 2008). CSRF is in fact not a new form of attack. In 1988 it was known as “confused deputy”. For a long time it was a “sleeping giant” (Grossman 2006), which came to prominence only recently.

CSRF can take many forms, some of them not so easy to understand. But a simple instantiation of CSRF would run the following way. A user has a trusted session (trust being guaranteed by cookies) with his bank website. If without having logged out from the session, the user goes to a malicious website and is induced to click on a link, a CSRF could occur. If HTTP request the user makes an HTTP Get request to the bank website, the browser of the user will make the query to the bank website. Since the cookies of the session are still active, the website will not be able to realize that the query technically originates from the malicious site and will execute it and it could be a instruction to transfer money from the account of the user. This is one (there are others) of the possible abuses of HTTP requests. This is an unfortunate consequence of what otherwise makes HTTP such a powerful protocol allowing a lot of functionalities in web applications.

In order for the attack to be successful, not only should the user omit to log off from the trusted session with the bank, but the attacker should know all the coordinates of the bank and user. There are several ways to do that. One, which is simple to understand is if the website of the bank has been compromised in the first place by another form of popular web attack: Cross Site Scripting (XSS) (Foggie et al. 2007). Then the user can find himself been send to a spurious website and induce into But there are many other ways to lure a hapless user into going a malicious website or let an attacker hijack a trusted session.

A few suggestions have been made for defense against CSRF, either on the server side (Zeller and Felten 2008) or on the user side (for example RequestRodeo (Johns and Winter 2006)). But “to be useful in practice, a mitigation technique for CSRF attacks has to satisfy two properties. First, it has to be effective in detecting and preventing CSRF attacks with a very low false negative and false positive rate. Second, it should be generic and spare web site administrators and programmers from application-specific modifications. Basically all the existing approaches fail in at least one of the two aspects” (Jovanovic et al. 2006).

Would an expert monitoring each HTTP request and everything that goes through the browser always be able to realize that a CSRF attack is unfolding? The answer is not obvious. But it is safe to say that in most cases he would. That suggests that a AI-based defense system located within the browser could in principle also detect attacks.

b. Web Application Firewalls (WAF)

Firewalls have been part of the arsenal of cyberdefense for many years. The simplest and also the most reliable ones deny access based on port number. The filtering can be more sophisticated, like being based on a deeper analysis of the incoming traffic, like deep packet inspection.

Web applications firewalls (WAF) cannot rely on port number as most web applications use the same port as the rest of the web traffic, i.e. port 80. WAFs are supposed to tell the

difference between benign and malicious web applications. This has to be made through deep packet inspection.

The idea of firewalls operating at the application layer is not new. They were introduced as “third generation” firewalls in the early 1990’s. They are used to protect data bases against SQL injections, for example. WAFs, are sometimes treated as a specialized form of application layer firewalls. WAFs began to enter the market at the end of the 90’s and tended to find their niche around specific applications. However sophisticated as they sometimes seem or are made to seem, as of today WAFs are not the reliable and performant tools that web security requires.

WAFs are in a sense very illustrative of what this chapter is about: to become what cybersecurity requires, WAFs need more Artificial Intelligence. One reason WAFs progress so slowly is that putting AI in security tools in general and in WAFs in particular is difficult. AI tends to be developed around specific applications. Many areas of applications have inspired aggressive AI research. Cybersecurity is not one of them, at least not yet, although it seems a very natural area of application as it is about computers. Human beings are completely in control of the rules and protocol and they could be designed to facilitate the use of AI.

7. Artificial Intelligence

a. The need for a new paradigm for defense

Instead of being adaptive, defense is purely reactive. In most cases it involves or essentially consists in limiting or foregoing some functionality. When a new attack has been discovered, more often than not the “security” solution consists in reducing the functionality. One major reason to turn toward AI, is to put an end at the present situation.

The move from DNS to DNSSEC illustrates somewhat the problem with the present approach. It has improved the security of the internet. But the cost is a complicated system of keys, whose renewal opens the door for scenarios of failures, which did not exist before. With DNSSEC the internet is less vulnerable to malicious exploitations of the weaknesses of the DNS system, but the use of a cumbersome system of authentication for all the servers involved does not make it more reliable.

The world of web applications is growing fast in importance and size, but in parallel it raises increasing concerns about security, which will not be solved adequately within the present paradigm of defense.

Same Origin Policy (SOP) is another example of the “old-fashioned” approach to cybersecurity. SOP (a policy adopted by most browsers) was designed to prevent the possibility that scripts originating from other than one site can be run on a web site (admittedly this has potentially dangerous consequences). Not only attacks such CSRF show that it is possible to circumvent the same origin policy, but that policy blocks other functionalities, which could be useful. In the words of Douglas Crockford: “[The Same Origin Policy] allows dangerous things while preventing useful ones”. The way out of that dilemma may lie in a much more intelligent defense.

In the case of CSRF, the proposed defenses are either in the website (alerting the system administrator that the website in its present design allows CSRF attacks) or in the client side. Typically the solutions suggested either reduce the functionality of the website, changes the protocol of trusted session by requiring more authentication, or (as is the case with request rodeo) offers a tool which limits partially the access to websites from the clients. In other

words, the solutions tend to make the protocols safer by making them more cumbersome and the protection often involves a reduction of functionality, i.e. it goes exactly against the logic, which underlies the success of the internet.

What an AI-based approach potentially offer is a way to address the multiple threats associated with cybersecurity, without having to rely on an increasing list of changing rules or security procedures for diagnostic and recovery procedures. If security tools were expert systems, which could not be abused as easily, the situation would be very different. Ideally they would understand what users are trying to do and make sure that this is what is happening. They would develop a sense of “situational awareness”, from which they would be able to tell malicious activity from legitimate ones. They would be able to make context dependent determinations.

b. Prospects of AI in cybersecurity

If AI means introducing intelligence in an automated system, there is no doubt that the future of cybersecurity lies in AI. But AI is at the same time an advanced field and at a very early stage of development.

The limits of the possible with AI are not known. The limits of the capabilities of AI are a moving frontier. In a “post biological intelligence” world (P. Davies, 2010), the division between natural and artificial intelligence will be blurred.

We are still far away from that world, but it is not too early to envision it. And the question is how should AI be introduced in the world of cybersecurity with maximum effect in the short term. Should we have a vision of AI-based cybersecurity as a cyber-equivalent of what happen with the immune system during biological evolution? I.e. of the creation over time of a large and complex organ inseparable from the rest of the organism? Should the first phase attempt to build the equivalent of a rudimentary immune system, with the vision of an eventual large and sophisticated one? Or should the search be more random and based on trying to introduce more intelligence in security tool whenever possible and wherever possible? In fact the two approaches differ only on paper. The immune system must have developed out of a random search as we are told the rest of biological evolution, leading in the long run to high levels of organization.

To what extent does AI in its present state provide a framework to start building such a system? It is impossible and not useful here to try and describe a field like AI. On the one hand it is a large body of academic knowledge (Russel and Novig 2003). When it comes to its applications, it looks more like a vast and fragmented field. Through expert systems AI has found applications in numerous fields from medical diagnosis to helping manufacture to finance management to fault analysis to advanced optimization, and to a too limited extent to cybersecurity.

Of the many techniques used in AI, when it comes to anomaly-based intrusion detection the techniques, which seem the most natural are either “statistics” based (Mitchell 1997) or “knowledge-based” (Kerker and Srinivas 2009).

The whole area of machine learning tends to make heavy use of statistics. The a priori caveat with that kind of approach in the context of intrusion detection is the possibility that the problem (or curse) with false positive re-emerges. If it is possible to set-up the system in such a way it can “reason” probabilistically (Pearl 1988) about events along the lines of the iterative Bayesian updating described previously, this problem may turn out manageable. Statistically based machine learning traditionally needs huge amount of data. This may turn problematic in many situations of interest for intrusion detection.

This points to the fact that there are fundamental cognitive differences between human beings and machines. Human beings need much less data to “learn” than machines and get a better power of discrimination. One implication of that remark is to invalidate partially the assumption that what human beings can do, machines will also be able to do.

Still an approach based on statistical learning in cybersecurity is not hopeless, quite the opposite. But this suggests that the best use of AI may not be to try to find a way to have machines replicating what human beings do.

An alternative to statistical learning is Knowledge Based systems (KBS) (kerkhar, 2009), although that approach raises also challenging issues. KBS tends also to be specialized. The system can acquire its knowledge in a variety of ways. It can be made to learn. A lot rides on the way knowledge is stored and represented. Those systems can reason and make inferences. In principle they could be used to make autonomous determination of whether a malicious attack is unfolding.

In practice in the case of web applications, for example, the information they have is what the browser sees: http requests, they can parse ad nauseam, website contents etc... One immediate challenge is to set up a knowledge base, which can make sense of such information.

Other approaches used in AI may turn out quite powerful in cybersecurity. Intrusion detection has some features in common with problem solving. Techniques using formal logic and theorem proving may turn out to be quite useful. If it were possible to reformulate the problem of intrusion detection as solving a logical problem, we would know better what the limits of the possible are.

As of now probabilistic reasoning seems to be the most natural and easiest way to introduce AI in intrusion detection, but this may not be the only one in the long run.

In the context of cybersecurity, AI applications could take several forms. But it is clear that to be useful any AI will have to be used intensively. Even if the processing power of computers is increasing impressively, one has to be concerned by the potential CPU overhead associated with any intensive AI technique. Considering that there is hardly any alternative in the long run to AI in cybersecurity, one has to be prepared to see cybersecurity to be part of the rest of cyber in the same way as the immune system is part of the animal's organisms. Instead of being a protection added at the end, it will be an integral part of the system, and as is the case with the immune system, it could be made “dual use”. I.e. its function may not be limited to protection.

8. Conclusions

Cybersecurity is a real challenge and the future of the internet partially depends on how that challenge is met. For a long time it has been clear that the cybersecurity response to the fast evolving threat needs to be much smarter than has been the case. The alternative is bound to lead to the situation predicted by Jonathan Zittrain [Zittrain 2008]: “If the problems associated with the Internet [...] are not addressed, a set of blunt solutions will likely be applied to solve problems at the expense of much of what we love about today's information ecosystem”.

This chapter focused on anomaly-based intrusion detection. But the role of AI in cybersecurity should not be seen as limited to that. Some cybersecurity problems needs urgent attention: for example the BGP (Border Gateway Protocol). In that case as was the case with the DNS system, the origin of the problem has to do with authentication. It seems

that there is little that a human expert can do let alone an AI system without an authentication protocol. Today that means a cumbersome system of keys, which slows down everything and introduces new failure modes. Furthermore as we saw recently with MD5, encryption protocols get eventually broken.

BGP does not have a system of authentication based on keys. It has no authentication system at all. It is vulnerable to any rogue routers. That is not sustainable. Eventually BGP may also use an authentication system based on keys, with the same negative consequences. A very nice scenario would be that AI could offer other ways toward authentication, which is one of the major basically unsolved problem in cybersecurity.

There is an imperative to put far more artificial intelligence in cybersecurity. The question is how best to do it.

Artificial intelligence is a vast and advanced field still relatively immature and definitely not optimized for cybersecurity. Specific work will be needed in artificial intelligence to facilitate its application to cybersecurity. Progress on that front will go faster if the possibility of applying AI techniques in cybersecurity inspires more interest to the AI community.

Cybersecurity could turn out a very good field of application for AI. It has the computer to computer interaction dimension and some of the problem solving culture (Newel and Simon 1972) developed in AI may find a natural area of application there.

One obstacle to the development of the interface between the two communities (security and AI) is the way security world operates. In cybersecurity, there is no real repository of knowledge. The knowledge exists, there is a lot of it, but it is scattered and not codified. Instead of looking at AI to try and find some "tools" which could be applied directly to cybersecurity, security people should have a harder look at their field, to make it more easily penetrable to outsiders, like AI people. As long as the approach to security is organized around the specific of attacks and the defense consists in looking for ways to prevent them, through some tweaking of existing protocols or functionalities, this will not happen soon. A possibility would be to be able to develop some sense of "situational awareness" which could be communicated or inoculated to the AI based tools.

AI systems and human beings differ in significant ways. The implications of those differences will be clearer when operational AI based tools become far more common in cybersecurity. Although the limit of the possible in AI is really ill-defined, it exists.

9. References

- Bacchus, Fahiem, Halpern, Joseph Y., and Levesque, Hector J., 1999, "Reasoning about noisy sensors and effectors in the situation calculus", *Artificial Intelligence*, 111(1-2): 171-208.
- Baral, Chitta and Gelfond, Michael, 2000, "Reasoning agents in dynamic domains", in *Logic-Based Artificial Intelligence*, Jack Minker, ed., Dordrecht: Kluwer Academic Publishers, 257-279.
- Baral, Chitta, McIlraith, Sheila, and San, Tran Cao, 2000, "Formulating diagnostic reasoning using an action language with narratives and sensing", in *KR2000: Principles of Knowledge Representation and Reasoning*, Anthony G. Cohn, Fausto Giunchiglia, and Bart Selman, eds., San Francisco: Morgan Kaufmann, 311-322.

- A. Barth, C. Jackson, and J. C. Mitchell. Robust Defenses for Cross-Site Request Forgery. In Proceedings of 15th ACM Conference, CCS, 2008
- Christopher M. Bishop. Pattern recognition and machine learning. Springer Verlag, Berlin, Germany, 2006
- Brachman, R. J., Levesque, H. J.: Knowledge Representation and Reasoning, Morgan Kaufmann, San Francisco, chapters 10 and 11 (2004).
- Alan Belasco et al. (2004). "Representing Knowledge Gaps Effectively". In: D. Karagiannis, U. Reimer (Eds.): *Practical Aspects of Knowledge Management, Proceedings of PAKM 2004, Vienna, Austria, December 2-3, 2004*. Springer-Verlag, Berlin Heidelberg.
- Jongho Choy and Sung-Bae Cho, Anomaly Detection of Computer Usage Using Artificial Intelligence Techniques_ R. Kowalczyk et al. (Eds.): PRICAI 2000 Workshop Reader, LNAI 2112, pp. 31–43, 2001. Springer-Verlag Berlin Heidelberg 2001
- Roberto Cordeschi: The role of heuristics in automated theorem proving. J.A. Robinson's resolution principle, *Mathware & Soft Computing* (1996) 3: 281-293
- Marco Cova, Davide Balzarotti, Viktoria Felmetsger, and Giovanni Vigna: Swaddler: An Approach for the Anomaly-Based Detection of State Violations in Web Applications, C. Kruegel, R. Lippmann, and A. Clark (Eds.): RAID 2007, LNCS 4637, pp. 63–86, 2007. Springer-Verlag Berlin Heidelberg 20
- T. Cover. The Best Two Independent Measurements are Not the Two Best. *IEEE Trans. on Systems, Man and Cybernetics*, SMC-4(1):116–117, January 1974
- Davies Paul: Eerie Silence, Renewing our Search for Alien Intelligence, Penguins Book 2010.
- Stefan Edelkamp and Carsten Elfers and Mirko Horstmann and Marcus-Sebastian Schröder and Karsten Sohr and Thomas Wagner, Early Warning and Intrusion Detection based on Combined AI Methods, Bremen, 2009.
- Scott E. Fahlman: *Marker-Passing Inference in the Scone Knowledge-Base System*, J. Lang, F. Lin, and J. Wang (Eds.): KSEM 2006, LNAI 4092, pp. 114 – 126, 2006. Springer-Verlag Berlin Heidelberg 2006
- Fahlman, S. E.: *NETL: A System for Representing and Using Real-World Knowledge*, MIT Press, Cambridge MA (1979)
- Feigenbaum, E.A.: Some Challenges and Grand Challenges for Computational Intelligence, *Journal of the ACM*, Vol. 50, No. 1, January 2003, pp. 32–40.
- Fogie, S., Jeremiah Grossman, Robert Hansen, Anton Rager, and Petko D. Petkov. XSS Attacks: Cross Site Scripting Exploits and Defense. Syngress, 2007
- S. Forrest, S.A. Hofmeyr and A. Somayaji, "Computer immunology," *CACM*, vol. 40, no. 10, pp. 88–96, October 1997
- Francois Gagnon and Babak Esfandiari, "Using Artificial Intelligence for Intrusion Detection", in *Frontiers in Artificial Intelligence and Applications*, Vol. 160, "Emerging Artificial Intelligence Applications in Computer Engineering", Edited by Ilias Maglogiannis, Kostas Karpouzis, Manolis Wallace, John Soldatos, IOS Press, ISBN 978-1-58603-780-2, 2007, pp. 295 - 306.

- A.K. Ghosh, A. Schwartzbard and M. Schatz, "Learning program behavior profiles for intrusion detection," *Proc. Workshop on Intrusion Detection and Network Monitoring*, pp. 51-62, Santa Clara, USA, April 1999.
- J. Grossman. CSRF, the sleeping giant.
<http://jeremiahgrossman.blogspot.com/2006/09/csrf-sleeping-giant.html>, Sep 2006
- S. A. Hofmeyr, S. Forrest, A. Somayaji, Intrusion Detection using Sequences of System Calls, *Journal of Computer Security*, 6:151--180, 1998
- Norbik Bashah Idris and Bharanidharan Shanmugam, Artificial Intelligence Techniques Applied to Intrusion Detection IEEE Indicon 2005 Conference, Chennai, India, 11 - 13 Dec. 2005, pp52-55.
- Norbik Bashah, Idris Bharanidharan Shanmugam, and Abdul Manan Ahmed, Hybrid Intelligent Intrusion Detection System, *World Academy of Science, Engineering and Technology* 11 2005, pp.23-26
- M. Johns and J. Winter. RequestRodeo: Client Side Protection against Session Riding. In F. Piessens, editor, *Proceedings of the OWASP Europe 2006 Conference*, refereed papers track, Report CW448, pages 5 - 17. Departement Computerwetenschappen, Katholieke Universiteit Leuven, May 2006.
- N. Jovanovic, E. Kirda, and C. Kruegel. Preventing Cross Site Request Forgery Attacks. *Securecomm and Workshops*, 2006, pages 1-10, Aug. 28 2006- Sept. 1 2006
- Kerkar RA and Sajja Priti Srinivas: "*Knowledge-based systems*", Jones & Bartlett Publishers, Sudbury, MA, USA (2009)
- Landwehr, Carl, Cybersecurity and Artificial Intelligence: From Fixing the Plumbing to Smart Water, *IEEE, Security and privacy*, September/October 2008, p.3
- Lee, W., S. J. Stolfo, *Data Mining Approaches for Intrusion Detection*, Columbia University, 1996
- Loveland D. (1978), *Automated theorem proving: a logical basis*, Amsterdam, North Holland.
- McCarthy, John, 1959, "Programs with common sense", in *Proceedings of the Teddington Conference on the Mechanization of Thought Processes*, London: Her Majesty's Stationary Office, 75-91.
- McCarthy, John, 1979, "First order theories of individual concepts and propositions", in *Machine Intelligence 9*, J.E. Hayes, D. Mitchie, and L.I. Mikulich, eds., Chichester, England: Ellis Horwood, 129-148.
- Meltzer B. (1969), The use of symbolic logic in proving mathematical theorems by means of a digital computer, in Bulloff J.J., Holyoke T.C., Hahn S.W. (eds), *Foundations of mathematics*, Berlin, Springer, 39-45.
- Meltzer B. (1971), Prolegomena to a theory of efficiency of proof procedures, in Findler N.V., Meltzer B. (eds), *Artificial intelligence and heuristic programming*, Edinburgh, Edinburgh University Press, 15-33.
- Minsky M. (1975), A framework for representing knowledge, in Winston P. (ed.), *Psychology of computer vision*, New York, McGraw-Hill, 211-280.
- Tom M. Mitchell. *Machine learning*. McGraw-Hill, New York, NY, 1997.
- Moore, Robert C., 1985, "A formal theory of knowledge and action", in *Formal Theories of the Commonsense World*, Jerry R. Hobbs and Robert C. Moore, eds., Norwood, New Jersey: Ablex Publishing Corporation, 319-358.

- B. Morel, Anomaly-Based Intrusion detection using Distributed Intelligent Systems, Proceedings of the 3rd Conference on Risks and Security of the Internet and Systems, Tunis, October 28-30, 2008, Tozeur, Tunisia
- J. von Neumann, "Probabilistic logics and the synthesis of reliable organisms from unreliable components", in C. E. Shannon and J. McCarthy, editors, *Annals of Math Studies*, numbers 34, pages 43-98. Princeton Univ. Press, 1956
- Newell A., Shaw J.C., Simon H.A. (1958), Elements of a theory of human problem solving, *Psychological Review*, 65, 151-166.
- Newell A., Simon H.A. (1965), Simulation of human processing of information, *American Mathematical Monthly*, 72, 111-118.
- Newell A., Simon H.A. (1972), *Human problem solving*, Englewood Cliffs, Prentice-Hall.
- Nilsson N.J. (1971), *Problem solving methods in Artificial Intelligence*, New York, McGraw-Hill.
- Judea Pearl. Probabilistic Reasoning in Intelligent Systems: Networks of Plausible Inference. Morgan Kaufmann, San Mateo, CA, 1988.
- Pearl, Judea, 2000, *Causality: Models, Reasoning, and Inference*, Cambridge, England: Cambridge University Press, ISBN 0-521-77362-8.
- T. Pietraszek Using Adaptive Alert Classification to Reduce False Positives in Intrusion Detection E. Jonsson et al. (Eds.): RAID 2004, LNCS 3224, pp. 102-124, 2004. Springer-Verlag Berlin Heidelberg 2004
- Mario Castro Ponce, Intrusion Detection System with Artificial Intelligence, FIST Conference - June 2004
- Robinson J.A. (1965), A machine oriented logic based on the resolution principle, *Journal of the Association for Computing Machinery*, 12, 23-41.
- Robinson J.A. (1967a), Heuristic and complete processes in the mechanization of theorem-proving, in Hart J.T., Takasu S. (eds.), *Systems and computer science*, Toronto, University of Toronto Press, 116-124.
- Russell, Stuart and Norvig, Peter, 2003, *Artificial Intelligence: A Modern Approach*, Englewood Cliffs, New Jersey: Prentice Hall, 2 ed.
- Blake Shepard et al. (2005). "A Knowledge-Based Approach to Network Security: Applying Cyc in the Domain of Network Risk Assessment". In: *Proceedings of the Seventeenth Innovative Applications of Artificial Intelligence Conference*. Pittsburgh, Pennsylvania, July 2005.
- Schneier, Bruce. On Security, 2008
- K. C. Tan, R. Maxion, The Effects of Algorithmic Diversity on anomaly detector performance", International Conference on Dependable Systems and networks, Yokohama Japan, 2005, p.216
- K.C. Tan, R. Maxion: "Why 6? Defining the Operational Limits of stide, an Anomaly- Based Intrusion Detector", Proc. IEEE Symposium on Security and Privacy, (2002)
- Turing, A. M. 1950. Computing machinery and intelligence. *Mind* 59, 433-460.
- Anne van Kesteren et al. Access control for cross-site requests.
<http://www.w3.org/TR/access-control/>. 2008
- C. Warrender, S. Forrest and B. Pearlmutter, "Detecting intrusions using system calls: Alternative data models," *Proc. IEEE Symposium on Security and Privacy*, pp. 133-145, May 1999

William Zeller and Edward W. Felten; Cross-Site Request Forgeries: Exploitation and Prevention, Princeton (2008);
<http://citp.princeton.edu/csrf/>

Jonathan Zittrain: The Future of the Internet and how to stop it, Blog, 2008

Intrusion Detection Systems

Edited by Dr. Pawel Skrobanek

ISBN 978-953-307-167-1

Hard cover, 324 pages

Publisher InTech

Published online 22, March, 2011

Published in print edition March, 2011

The current structure of the chapters reflects the key aspects discussed in the papers but the papers themselves contain more additional interesting information: examples of a practical application and results obtained for existing networks as well as results of experiments confirming efficacy of a synergistic analysis of anomaly detection and signature detection, and application of interesting solutions, such as an analysis of the anomalies of user behaviors and many others.

How to reference

In order to correctly reference this scholarly work, feel free to copy and paste the following:

Benoît Morel (2011). Anomaly Based Intrusion Detection and Artificial Intelligence, Intrusion Detection Systems, Dr. Pawel Skrobanek (Ed.), ISBN: 978-953-307-167-1, InTech, Available from:

<http://www.intechopen.com/books/intrusion-detection-systems/anomaly-based-intrusion-detection-and-artificial-intelligence>

INTech
open science | open minds

InTech Europe

University Campus STeP Ri
Slavka Krautzeka 83/A
51000 Rijeka, Croatia
Phone: +385 (51) 770 447
Fax: +385 (51) 686 166
www.intechopen.com

InTech China

Unit 405, Office Block, Hotel Equatorial Shanghai
No.65, Yan An Road (West), Shanghai, 200040, China
中国上海市延安西路65号上海国际贵都大饭店办公楼405单元
Phone: +86-21-62489820
Fax: +86-21-62489821

© 2011 The Author(s). Licensee IntechOpen. This chapter is distributed under the terms of the [Creative Commons Attribution-NonCommercial-ShareAlike-3.0 License](#), which permits use, distribution and reproduction for non-commercial purposes, provided the original is properly cited and derivative works building on this content are distributed under the same license.