A Network Security Classifier Defense:

Against Adversarial Machine Learning Attacks

Michael J. De Lucia
Network Science Division
U.S. Army Research Laboratory
Aberdeen Proving Ground, MD
michael.j.delucia2.civ@mail.mil

Chase Cotton
Department of Electrical and Computer Engineering
University of Delaware
Newark, DE
ccotton@udel.edu

ABSTRACT

The discovery of practical adversarial machine learning (AML) attacks against machine learning-based wired and wireless network security detectors has driven the necessity of a defense. Without a defense mechanism against AML, attacks in wired and wireless networks will go unnoticed by network security classifiers resulting in their ineffectiveness. Therefore, it is essential to motivate a defense against AML attacks for network security classifiers. Existing AML defenses are generally within the context of image recognition. However, these AML defenses have limited transferability to a network security context. Unlike image recognition, a subject matter expert generally derives the features of a network security classifier. Therefore, a network security classifier requires a distinctive strategy for defense. We propose a novel defense-in-depth approach for network security classifiers using a hierarchical ensemble of classifiers, each using a disparate feature set. Subsequently we show the effective use of our hierarchical ensemble to defend an existing network security classifier against an AML attack. Additionally, we discover a novel set of features to detect network scanning activity. Lastly, we propose to enhance our AML defense approach in future work. A shortcoming of our approach is the increased cost to the defender for implementation of each independent classifier. Therefore, we propose combining our AML defense with a moving target defense approach. Additionally, we propose to evaluate our AML defense with a variety of datasets and classifiers and evaluate the effectiveness of decomposing a classifier with many features into multiple classifiers, each with a small subset of the features.

WiseML '20, July 13, 2020, Linz (Virtual Event), Austria

ACM ISBN 978-1-4503-8007-2/20/07…$15.00
https://doi.org/10.1145/3395352.

CCS CONCEPTS

• Security and privacy~Intrusion/anomaly detection and malware mitigation~Intrusion detection systems • Security and privacy~Network security • Computing methodologies~Machine learning~Machine learning algorithms~Ensemble methods

KEYWORDS

Adversarial Machine Learning, Machine Learning, Network Security, Cyber Security, Cyber Defense

ACM Reference format:

Michael J. De Lucia and Chase Cotton. 2020. A Network Security Classifier Defense: Against Adversarial Machine Learning Attacks. In *Proceedings of ACM Workshop on Wireless Security Machine Learning (WiseML'20). ACM, New York, NY, USA, 6 pages.*

1 Introduction

The discovery of practical adversarial machine learning (AML) attacks against machine learning-based wired and wireless network security detectors has driven the necessity of a defense. Without a defense mechanism against AML, attacks in wired and wireless networks will go unnoticed by network security classifiers resulting in their ineffectiveness. Therefore, it is essential to motivate a defense against AML attacks for network security classifiers.

An increased reliance on the use of machine learning in network security detectors stimulates the risk of adversarial employment of AML, to evade detection. Historically, traditional network security detectors have encountered a similar threat whereby an adversary perturbs malicious network traffic with the intent of avoiding exposure. The defender subsequently analyzes the attack and develops a countermeasure to detect the attackers malicious network traffic. This process is a vicious, repetitive cycle between the attacker and the defender. Similarly, the same process is

likely to occur with the deployment of machine learning-based network security detectors.

Our previous work [9] demonstrates the effectiveness of an AML attack against an existing network scanning classifier. The AML attack in our previous work is a mimicry attack, where the adversary's goal is to make their malicious network traffic appear as benign. Thus, the adversary perturbs their malicious traffic to mimic the feature values of benign network traffic. Accordingly, the perturbed feature values of the malicious network traffic are within the region of benign network traffic. As such, a disparate feature set is essential to differentiate a malicious and benign sample.

Our contributions in this work are as follows:

- We propose a novel defense-in-depth approach for network security classifiers using a hierarchical ensemble of classifiers, each using a disparate feature set.
- Subsequently we present an evaluation of the hierarchical ensemble to defend an existing network security classifier against an AML attack.
- Additionally, we discover a novel set of features to detect network scanning activity.

The outline of this paper is as follows. We present related work in the defense of AML in section 2. We then present background in section 3 and our novel AML defense in section 4. An evaluation of our AML defense follows in section 4. Lastly, we present a conclusion and future work in section 5.

2 Related Work

There is a limited breadth of research on the defense of network security classifiers against AML. Conversely, the work in a defense against AML in image recognition is extensive. However, many of the defensive techniques in image recognition do not effectively transfer to a network security context.

A comprehensive survey of AML defensive techniques specifically applied within image recognition is provided by Biggio and Roli in [5]. Their survey discusses defensive techniques grouped into the broad categories of Adversarial Training, Detection and Rejection of samples far from the training data, and classifier ensembles. The following further illustrates each of these categories of defensives and evaluates the techniques for transferability to a network security context.

The work in Adversarial Training is primarily within the context of image recognition to protect a classifier from an AML attack. Adversarial training employs AML to optimize the adversarial objective to create modified examples via perturbation of attack examples to cause the classifier to change its prediction from the original sample. Thus, the discovered adversarial samples are correctly labeled in the data set used during the training phase of the classifier. The process of adversarial training is analogous to the repetitive cycle of a traditional network security signature developed by a human security analyst in response to a discovered attack.

Adversarial Training as a defense against AML in the context of image recognition is extensively studied in [6,7,11,13]. The use of adversarial training is effective at defending image recognition classifiers against AML since the AML perturbations are minimal and do not mimic the features of the intended class. However, within a network security setting, the perturbations do not need to be minimal. Additionally, during a mimicry AML attack, the perturbed attack sample feature values are within the range of benign samples. Thus, the goal of an attacker is to evade (cause misclassification) an intrusion detection system (IDS) classifier by perturbing the attack network traffic feature values to mimic those of a benign (i.e., web browsing) session. Therefore, an adversarial sample produced by a mimicry AML attack would be indistinguishable from the actual benign samples.

The next category of defenses is Detection and Rejection of samples far from the training data. This category of defenses is studied within the context of image recognition, malware detection, and spam detection in [2,12,17]. Conceptually these defenses rely on the ability to reject samples, which are anomalous in comparison to the in-class samples from within the training dataset. In their work the defenses prove to be effective at defending against AML. However, these techniques are ineffective at defending a network security classifier against an AML mimicry attack. Recall the objective of an AML mimicry attack is to perturb the malicious samples to match the range of feature values of a benign sample. Therefore, these perturbed samples are not anomalous in comparison to the in-class training dataset.

Another category of defense is the use of ensemble techniques. Ensemble techniques are employed as a defense within the image recognition, span, and malware detection in [2,4,15,17]. The concept of the use of an ensemble for the defense against an AML attack is transferable to a network security context. The use of multiple classifiers employed as an ensemble is recognized to be beneficial to achieve an optimal performance. Researchers in [18] propose an ensemble of classifiers using a diverse set of features to describe the payload of a packet. Their approach shows an improved performance and avoids the curse of dimensionality compared to a single classifier containing all features.

Another example of an ensemble technique to defend against evasion is demonstrated by researchers in [19,20] for information fusion to enable a resilient biometric indicator. Their work shows the use of an ensemble of classifiers for biometric identification. Similarly, an example of the use of an ensemble of classifiers to evaluate a sample from disparate perspectives is proposed to defend against image recognition AML attacks in autonomous vehicles [16].

Subsequently, a combination of classifiers using a disparate feature set is demonstrated in [14] to perform better than a single classifier containing all features. Thus, in [3] researchers suggest that an ensemble of classifiers trained using loosely correlated disparate feature sets is more effective than using a single classifier containing all features. Additionally, their ensemble randomization approach is a random sampling of the training data set. Furthermore, their work suggests that a traditional IDS makes a final decision on maliciousness by fusing the scores produced by several independent modules, each based on a subset of the whole feature space.

We believe an ensemble approach is effective at defending against AML. While the related work ensemble approaches use an ensemble, they are not hierarchical and do not focus on the use of disparate feature sets. A traditional IDS compares network traffic with numerous known malicious signatures in a hierarchical fashion to produce the output label of malicious or benign. Our proposed approach for an AML defense uses a hierarchical ensemble combined with disparate feature sets.

2 Background

We propose a novel method to defend a network security classifier using machine learning against AML attacks. Our approach is composed of a hierarchical ensemble of heterogeneous classifiers using disparate feature sets. The assumptions of the threat environment and the theoretical foundations of our approach follows.

2.1 Assumptions

We define the assumptions of our novel AML defense within the perception of the adversary's knowledge of a machine learning based network security classifier. Our approach assumes the adversary does not have access to the trained model or the training dataset (e.g. grey-box). Additionally, we presume the adversary is aware of the feature space of the target network security classifier. It is conceivable that the attacker would have knowledge of the feature space, as many classifiers are either open source or release a whitepaper describing the features.

2.2 Defense-in-Depth

Our novel AML defense builds on a security concept known as "defense-in-depth". The defense-in-depth strategy leverages a combination of defensive layers, such that some layers will strengthen and mitigate other layers' weaknesses [21,23]. Each layer of the defense is disparate and has the property of "independent vulnerabilities", which would require an adversary to have both expertise and time to evade all defensive layers [10,21]. Thus, the defense-in-depth methodology increases the cost to the adversary in terms of time, skillset, and money.

Therefore, in terminology of the defense-in-depth model, each independent classifier of the AML defense hierarchical ensemble is a layer. Thus, each independent classifier has an "independent vulnerability". Furthermore, an independent classifier mitigates another's vulnerability. Within the context of AML defense of a network security classifier, a vulnerability is the malicious network traffic evasion (e.g. malicious network traffic misclassified as benign).

3 Defense of AML

In our proposed AML defense, each successive classifier must be resilient to an adversarial attack against the proceeding classifier in the ensemble. Each independent classifier in our AML defense is composed of a disparate feature set. Thus, there exist alternative methods relying on disparate feature sets to detect the same type of attack.

Furthermore, our approach encompasses a set of rules to combine classification decisions from each respective classifier within the ensemble to predict the label of a sample.
A typical ensemble classifier combines the predictions of each independent classifiers by a majority voting process. However, the ensemble in our AML defense approach employs a method for prediction combination like a stacked ensemble. Specifically, our AML defense is a hierarchical ensemble.

The authors of [10] demonstrate a hierarchical rule-based ensemble architecture to combine predictions of the independent classifiers with the intent of adaptive learning. They achieved adaptive learning using the hierarchical ensemble to enable the integration of future developed classifiers. The researchers in [10], trained a new classifier to be added to the ensemble, using only a data set containing the newly discovered attacks. Accordingly, proceeding models in the ensemble do not train with the data set containing the latest discovered attacks. Additionally, each of the classifiers in the ensemble depend on the same feature set. Furthermore, the researchers enable the use of three classes with a label of attack, normal, or anomaly.

Our proposed AML defense approach leverages the concept of a hierarchical ensemble with a rule set for prediction combination similar to [10], but with novel differences. The objective of our proposed hierarchical ensemble is to defend against an AML mimicry attack as opposed to enabling adaptive learning. Furthermore, our hierarchical ensemble uses only two classes with a label of attack or benign. Moreover, our hierarchical ensemble contains independent classifiers using a disparate feature set. Thus, each successive classifier in the hierarchy mitigates an AML vulnerability of its predecessor. Additionally, our AML defense enables the integration of new independent classifiers to counter AML attacks.

Figure 1 depicts our novel AML defense ensemble, with two independent classifiers H_1 and H_2 respectively. Both x^1 and x^2 are samples composed of the disparate feature set of H_1 and H_2 respectively. The training of each classifier uses the malicious and benign samples composed of the respective features. Algorithm 1 presents the ruleset for the hierarchical AML defense ensemble shown in figure 1. For simplicity, we only present an ensemble of two classifiers. However, the integration of additional classifiers would operate iteratively according to Algorithm 1.

The addition of classifiers to our AML defense ensemble, results in a further cost to the adversary. Recall, the objective of a defense in depth model is to increase the cost to the adversary by adding disparate layers which mitigates another's vulnerabilities. Our AML defense approach operates in a similar fashion with the use of a hierarchical ensemble composed of independent classifiers, each with a disparate feature set.

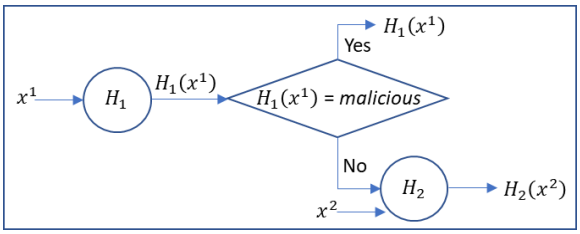

Figure 1: AML Defense Ensemble

Algorithm 1: AML Defense Ensemble ruleset

1: $if\ (H_1(x_1)\ = malicious)\ then$
2: $\quad output \leftarrow H_1(x_1)$
3: $else\ output \leftarrow H_2(x_2)$

Thus, the total cost $C(t)$ to the adversary to evade each classifier H_i in a hierarchical ensemble of size n, where t_i is the time required to implement an attack on classifier H_i is as follows:

$$C(t) = \sum_{i=1}^{n} t_i \quad (1)$$

Therefore, a greater number of layers increases the cost to the adversary to successfully evade detection by the ensemble classifier. Each independent classifier requires the adversary to expend time in perturbation of the disparate features and attack implementation of the respective classifier layer. Consequently, our AML defense approach also comes at a cost to the defender to implement and maintain.

3.1 Defense of AML Example Scenario

We provide a further understanding of our proposed hierarchical ensemble for AML defense through a botnet detection scenario. The objective of the classifier in this scenario is to detect the presence of a botnet within a network. Again, for simplicity we limit the number of classifiers in this scenario to two.

The target classifier H_1 in figure 1 performs detection of botnet presence by using the frequency of packet sizes. The adversary counters the target classifier by perturbing the malicious packet size frequencies to mimic benign network traffic. Thus, to increase the cost to the adversary and counter their attack, a defender integrates an additional classifier using a disparate feature set to form the hierarchical ensemble. An alternate detection method using a disparate feature set of domain name features exists. Thus, the second classifier H_2 in figure 1 uses the disparate feature set based on the domain name.

Consequently, the cost of a successful attack by the adversary increases. As, the adversary must counter and implement attacks, or two classifiers based on different feature sets. Assume the time for an adversary to defeat classifiers H_1 and H_2, respectively is 336 and 672 hours. Accordingly, to equation 1 the total cost for a successful attack by the adversary is 1,008 hours. Thus, the addition of the classifier H_2 increases the cost of a successful attack by 672 hours.

3.2 Defense of AML Wireless Networks

Our AML defense extends to both wired and wireless networks. The features in the examples and evaluation of our AML defense is at the networking and application layer. These protocols are the same in both wired and wireless networks. Additionally, network scanning is also a challenge in wireless networks such as Internet of Things (IoT). Moreover, the features of the independent classifiers within our AML defense can be characteristics within a wireless context.

4 AML Defense Evaluation

We evaluate our hierarchical ensemble AML defense to defend an existing network scanning classifier from an AML attack. The AML attack on an existing network scanning classifier is further discussed in our previous work [9]. Further evaluation of the AML defense and alternative classifiers are offered in our work [8]. We present the existing network scanning classifier and dataset, followed by an evaluation of our AML defense.

4.1 Target Classifier and Dataset

The objective of the existing network security classifier is to detect the presence of network scanning activity in a network. The dataset and classifier are derived from [22]. The data set consists of packet captures collected for one hour from benign (no scanning activity) and malicious (scanning activity) hosts. The three features of the dataset are as follows:

- percent_tcp_unsuc: Percentage of TCP flows with unsuccessful connections
- percent_flow_udp: Percentage of network flows that are UDP
- percent_flow_icmp: Percentage of network flows that are ICMP

The dataset split is 80% and 20% for training and testing, respectively. The test dataset consists of 57 scanning and 51 benign samples. Before introducing the AML attack, the baseline accuracy of the network scanning classifier is 100%. The adversary perturbs all 57 scanning samples reduces the accuracy to 47%. The confusion matrix in figure 2 indicates the 57 perturbed scanning samples (malicious) misclassified as benign.

	Predicted Benign	Predicted Malicious
Actual Benign	51	0
Actual Malicious	57	0

Figure 2: Confusion matrix after AML attack

4.2 Evaluation

We leverage our proposed AML defense architecture to defend the existing network scanning classifier discussed. Thus, we motivate a second random forest classifier using disparate features to defend the existing network scanning classifier. For simplicity, we only use two classifiers in the hierarchical ensemble.

Within the context of our AML defense hierarchical ensemble of figure 1, H_1 and H_2 are respectively the target classifier and proposed classifier based on a disparate feature set. The features of the second classifier H_2 are motivated based on the expected behavior of TCP resets and destination ports visited by a host. The disparate feature set of H_2 is as follows:

- n_dst_port: Number of unique destination ports
- entropy_dst_port: Entropy of the destination ports
- n_dst_tcp_unsuc_port: Number of unique destination ports with unsuccessful TCP connections

Ideally, TCP resets should occur in rare conditions of network and application errors or resource exhaustion [1]. However, in practice offending applications and servers also end TCP connections with a reset rather than a proper TCP close with a "Fin" flag [1]. Consequently, a small

subset of unique destination ports occurs due to recurrent TCP reset behaviors in an ideal network.

The presence of network scanning on a host will increase the number of unsuccessful TCP flows (i.e. ended with a TCP reset) due to ports not being open or conducting half-open (i.e. prematurely ending the TCP handshake with a reset) scans. Additionally, a network scanner inherently increases the number of destination ports as the objective is to identify open ports on a host. In contrast a benign host typically visits a small subset of destination ports. Thus, a host containing a network scanner increases the entropy of the destination ports due to an increase of unique ports visited with a lower frequency. Moreover, a network scanner increases the number of unsuccessful TCP connections to destination ports due to ports not being open or conducting half-open scans.

We evaluate the effectiveness of our proposed hierarchical ensemble, composed of classifier H_1 and H_2, from an AML attack. The AML attack does not change the number of ports scanned. Thus, our hierarchical ensemble counters the AML attack and restores the classifier accuracy to 100% detection of the scanning hosts. The confusion matrix of figure 3, shows all 57 of the perturbed scanning flows correctly classified as scanning.

	Predicted Benign	Predicted Malicious
Actual Benign	51	0
Actual Malicious	0	57

Figure 3: Confusion matrix for AML Defense

Therefore, the attacker would need to not only perturb the H_1 classifier features, but also the number of destination ports scanned. Thus, the cost to the adversary for a successful attack would increase by the time it takes the adversary to defeat classifier H_2. Assume the adversary expends 336 and 672 hours to defeat classifier H_1 and H_2 respectively. Accordingly, to equation 1, the cost of successful attack to the adversary is 1,008 hours.

A successful attack against both classifiers in the ensemble, necessitates the addition of classifiers (layers) to counter these attacks. Ideally, the features within successive layers of the ensemble are dependent on features of preceding layers. Thus, perturbation of features in preceding layers would induce an unwanted change in features of successive layers. These unwanted changes in features of successive layers enables detection of the attack.

5 Conclusion and Future Work

The increasing reliance on machine learning within network security classifiers consequently escalates the potential for adversarial use of AML to evade detection. Thus, cyber defenders must counter the adversary's use of AML to secure network security classifiers. Thus, countering AML attacks enables persistent attack detection.

To counter AML within a network security context, we proposed a novel AML defense using a hierarchical ensemble and motivated by a defense-in-depth methodology. Each independent classifier of the ensemble relies on a disparate feature set. Consequently, the use of multiple classifiers using a disparate feature set increases the cost to the adversary. We demonstrated the effectiveness of our proposed AML defense to protect an existing network scanning classifier. Additionally, we also motivated a novel feature set based on the destination ports visited, to detect network scanning.

We propose to enhance our AML defense approach in future work. A shortcoming of our approach is the increased cost to the defender for implementation of each independent classifier. Therefore, we propose combining our AML defense with a moving target defense approach. In this approach we envision a random selection of a subset of classifiers to use for detection. We also propose to evaluate methods for computing the number of layers required for a successful defense. Additionally, we propose to evaluate our AML defense with a variety of wired and wireless datasets and classifiers which have a greater number of features. Moreover, we propose to evaluate the effectiveness of decomposing a classifier with many features into multiple classifiers, each containing a small subset of the features.

REFERENCES

[1] Martin Arlitt and Carey Williamson. 2005. An analysis of TCP reset behaviour on the internet. *SIGCOMM Comput. Commun. Rev.* 35, 1 (January 2005), 37–44. DOI:https://doi.org/10.1145/1052812.1052823

[2] Battista Biggio, Igino Corona, Zhi-Min He, Patrick PK Chan, Giorgio Giacinto, Daniel S. Yeung, and Fabio Roli. 2015. One-and-a-half-class multiple classifier systems for secure learning against evasion attacks at test time. In *International Workshop on Multiple Classifier Systems*, Springer, 168–180.

[3] Battista Biggio, Giorgio Fumera, and Fabio Roli. 2008. Adversarial pattern classification using multiple classifiers and randomisation. In *Joint IAPR International Workshops on Statistical Techniques in Pattern Recognition (SPR) and Structural and Syntactic Pattern Recognition (SSPR)*, Springer, 500–509.

[4] Battista Biggio, Giorgio Fumera, and Fabio Roli. 2010. Multiple classifier systems for robust classifier design in adversarial environments. *International Journal of Machine Learning and Cybernetics* 1, 1–4 (December 2010), 27–41. DOI:https://doi.org/10.1007/s13042-010-0007-7

[5] Battista Biggio and Fabio Roli. 2018. Wild patterns: Ten years after the rise of adversarial machine learning. *Pattern Recognition* 84, (December 2018), 317–331. DOI:https://doi.org/10.1016/j.patcog.2018.07.023

[6] Joan Bruna, Christian Szegedy, Ilya Sutskever, Ian Goodfellow, Wojciech Zaremba, Rob Fergus, and Dumitru Erhan. 2013. Intriguing properties of neural networks. *arXiv preprint arXiv:1312.6199* (2013).

[7] Nilesh Dalvi, Pedro Domingos, Sumit Sanghai, and Deepak Verma. 2004. Adversarial classification. In *Proceedings of the tenth ACM SIGKDD international conference on Knowledge discovery and data mining*, 99–108.

[8] Michael J. De Lucia,. 2020. Machine Learning Enhanced Network Security. Doctoral Dissertation. University of Delaware, Newark, DE.

[9] Michael J. De Lucia and Chase Cotton. 2019. Adversarial machine learning for cyber security. *Journal of Information Systems Applied Research* 12, 1 (April 2019), 26.

[10] Wei. Fan and Salvatore J. Stolfo. 2002. Ensemble-based adaptive intrusion detection. In *Proceedings of the 2002 SIAM International Conference on Data Mining*. Society for Industrial and Applied Mathematics, 41–58. DOI:https://doi.org/10.1137/1.9781611972726.3

[11] Ian J. Goodfellow, Jonathon Shlens, and Christian Szegedy. 2014. Explaining and harnessing adversarial examples. *arXiv:1412.6572 [cs, stat]* (December 2014). Retrieved November 18, 2018 from http://arxiv.org/abs/1412.6572

[12] Roberto Jordaney, Kumar Sharad, Santanu K. Dash, Zhi Wang, Davide Papini, Ilia Nouretdinov, and Lorenzo Cavallaro. 2017. Transcend: Detecting concept drift in malware classification models. In *26th USENIX Security Symposium (Security 17)*, 625–642.

[13] Alex Kantchelian, J. Doug Tygar, and Anthony Joseph. 2016. Evasion and hardening of tree ensemble classifiers. In *International Conference on Machine Learning*, 2387–2396.

[14] Josef Kittler, Mohamad Hatef, Robert PW Duin, and Jiri Matas. 1998. On combining classifiers. *IEEE transactions on pattern analysis and machine intelligence* 20, 3 (1998), 226–239.

[15] Aleksander Kołcz and Choon Hui Teo. 2009. Feature weighting for improved classifier robustness. In *CEAS'09: sixth conference on email and anti-spam*.

[16] James Rundle and John McCormick. 2020. Bosch deploys ai to prevent attacks on cars' electronic systems. *Wall Street Journal*. Retrieved January 13, 2020 from https://www.wsj.com/articles/bosch-deploys-ai-to-prevent-attacks-on-cars-electronic-systems-11578306600

[17] Dongyu Meng and Hao Chen. 2017. Magnet: a two-pronged defense against adversarial examples. In *Proceedings of the 2017 ACM SIGSAC Conference on Computer and Communications Security*, 135–147.

[18] Roberto Perdisci, Guofei Gu, and Wenke Lee. 2006. Using an ensemble of one-class svm classifiers to harden payload-based anomaly detection systems. In *Sixth International Conference on Data Mining (ICDM'06)*, IEEE, 488–498.

[19] Arun A. Ross, Anil K. Jain, and Karthik Nandakumar. 2006. Information fusion in biometrics. *Handbook of Multibiometrics* (2006), 37–58.

[20] Arun Ross and Anil Jain. 2003. Information fusion in biometrics. *Pattern recognition letters* 24, 13 (2003), 2115–2125.

[21] W. Tirenin and D. Faatz. 1999. A concept for strategic cyber defense. In *MILCOM 1999. IEEE Military Communications. Conference Proceedings (Cat. No.99CH36341)*, IEEE, Atlantic City, NJ, USA, 458–463. DOI:https://doi.org/10.1109/MILCOM.1999.822725

[22] Sridhar Venkatesan, Shridatt Sugrim, Rauf Izmailov, Cho-Yu J. Chiang, Ritu Chadha, Bharat Doshi, Blaine Hoffman, E. Allison Newcomb, and Norbou Buchler. 2018. On detecting manifestation of adversary characteristics. In *MILCOM 2018 - 2018 IEEE Military Communications Conference (MILCOM)*, 431–437. DOI:https://doi.org/10.1109/MILCOM.2018.8599754

[23] 2019. Defense in depth. Retrieved November 21, 2019 from https://apps.nsa.gov/iaarchive/library/ia-guidance/archive/defense-in-depth.cfm