Machine-learning Classifiers for Security in Connected Medical Devices

Sida Gao
Computing and Software Systems
University of Washington - Bothell
Bothell, Washington
Email: gaosida@uw.edu

Geethapriya Thamilarasu
Computing and Software Systems
University of Washington - Bothell
Bothell, Washington
Email: geetha@uw.edu

***Abstract*—Medical devices equipped with wireless connectivity and remote monitoring features are increasingly becoming connected to each other, to an outside programmer and even to the Internet. While Internet of Things technology enables healthcare professionals to fine tune or modify medical device settings without invasive procedures, this also opens up large attack surfaces and introduces potential security vulnerabilities. Medical device hacks are slowly becoming a reality and it becomes more critical than ever to defend and protect these devices from security attacks. In this paper, we assess the feasibility of using machine learning models to efficiently determine attacks targeted on a medical device. Specifically, we develop feature sets to accurately profile a medical device and observe any deviation from its normal behavior. We test our method using different machine learning algorithms and provide a comparison analysis of the detection results.**

***Index Terms*—Internet of Health, medical device security, secure health, machine learning security**

I. Introduction

Internet of Things, the next wave of technology is making strong advances in healthcare. Connected medical devices greatly improve the quality and effectiveness of healthcare delivery and service, especially benefitting the elderly, patients with chronic conditions, and those requiring constant supervision. Wireless connectivity and integration of IoT features to medical devices enables remote monitoring and access to medical information for prevention, maintenance, diagnosis and treatment of patient conditions as well as reduces the costs needed to maintain and support these devices. However, this increased connectivity also introduces newer security vulnerabilities in the healthcare domain [1]. While no real-world incidents have been documented yet, medical device hacks, including hacking of an insulin pump and a pacemaker have been successfully demonstrated by researchers [2], [3]. Security attacks on wireless medical devices can disable the device and potentially cause life-threatening damage to the patients. For instance, resource depletion attacks can target resource constrained devices such as an implanted medical device jeopardizing the availability of these devices [4]. Resource depletion not only targets endpoint host devices but also network resources such as processing capability or memory consumption. Medical device hacks can affect patients physiology. For instance, patients could be flooded with drugs or hacking an implant to ignore treatments. Replay attacks can occur on these devices, where legitimate message to turn off a device might be recorded and replayed at a later time causing the attacker to turn off a therapy or quietly change the state of the device.

Recently, a number of solutions have been proposed to address the different security vulnerabilities in wireless medical devices. In this paper, we focus on applying machine learning techniques to identify and detect attacks on implanted medical devices(IMD). We develop a solution that does not necessitate any hardware or software modifications to the device. Our goal is to implement a few different machine learning algorithms for attack detection and evaluate their feasibility and performance. Our solution uses an external detection device that monitors the network and uses machine learning classifiers to detect anomalies.

The remainder of the paper is organized as follows. In Section II, we provide a brief review of related research in this domain. In Section III, we discuss our network and adversary model. In Section IV, we propose the machine learning based detection mechanism. In Section V, we discuss our simulations and provide a comparison analysis of different learning algorithms..

II. Related Work

Security solutions for wireless medical devices have received a great deal of attention in the recent years. Gollakota et al. [5] proposed a solution, where an external high powered device known as Shield, prevents unauthorized access to an implanted medical device by emitting a jamming signal whenever it detects an unauthorized wireless link between the IMD and a remote device. The drawback of this scheme is that it mandates a patient to carry an external device in their person at all times. The patient may also lose or forget to carry the shield. A novel authentication solution-using patient's unique encrypted heartbeat is proposed in [6] to prevent unauthorized access to patient data.

In [7], the authors proposed a biometric based authentication to deal with emergencies where a patient may be unconscious. The two-level security defense solution first tests for fingerprints and the weight of patient and then extracts patients iris information for authentication. IMDGuard is another solution,

where an external device mediates the communication between an IMD and an external programmer device [8]. This solution assumes a shared secret key between the medical device and the guardian. It also introduces a significant communication overhead. Hei et al. [9] proposed a mechanism to detect resource depletion attacks using the access pattern of the device. The limitation of this solution is that the attacker can force the medical device to respond to its malicious messages. Denning et al. [10] proposed an external device, called the cloaker that acts as a proxy authorizing all communication between the IMD and the programmer. In case of emergency scenarios, the external device can be easily removed allowing for open communication. RFID Guardian proposed in [11] is a portable electronic device that scans all devices in its range, manages RFID keys, authenticates nearby programmers that request access to the IMD, and blocks all unauthorized readers. While originally proposed for RFID systems, the authors propose to integrate the Guardian into a device that the patient always carries, such as a cellphone or wearable device.

In addition to encrypting medical data and authenticating access to medical devices, anomaly based detection approaches are gaining popularity to detect security attacks in these networks. Henry et al. proposed a scheme [12] that tracks acoustic bowel sounds to detect correlated physiological changes when an insulin bolus is administered. Medmon is an anomaly based detection mechanism, where an external entity acts as a security monitor that observes anomalous transmissions by examining physical and behavioral characteristics of the communication to and from the IMD [13]. While there is a growing interest in exploring anomaly based detection techniques for security, their research on wireless medical devices is still limited. Our work specifically aims to advance the research in this area. We contribute to the existing research by exploring and evaluating machine learning algorithms for detecting security attacks on medical devices.

III. Network and Attack Model

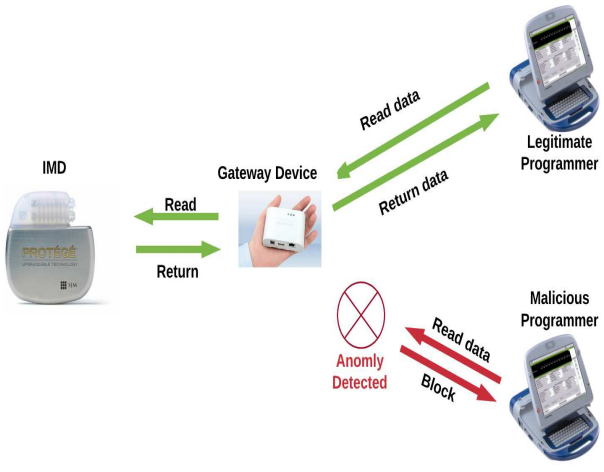

Fig. 1. IMD communication with network entities

Our network model consists of the IMD device, the programmer and the gateway node. As shown in Figure 1, the programmer is an external device responsible for controlling, monitoring and maintaining the IMD. The programmer uses radio frequency transmission through wireless channels to communicate with an IMD. The gateway node in the network is responsible for monitoring the wireless communication between IMD and other entities in the network. IMD regularly communicates with the gateway for granting access to external devices. Common communication scenarios in this network include communication between programmer and an IMD; communication between a reader and an IMD; or communication between an IMD and the gateway.

In our system, attacker is an entity that does not have authorized access to IMD. Our attack model mainly consists of a forced authentication attack that results in resource depletion of the device. Specifically, the malicious entity sends repetitive communication requests to the IMD for authentication. In addition to authentication requests, attacker can also flood the IMD by targeting specific IMD functions such as reading patient data and retrieving patient therapy history. When the medical device responds to large number of authentication requests or other functions, it results in serious depletion of its battery power resources, eventually causing a denial of service in the network. We also consider scenarios where an attacker is capable of sidestepping authentication and encryption protocols used in the network and where an attacker impersonates a programmer or a reader device to access IMD's data for malicious purposes.

We make the following assumptions in our system:

1) The attacker is an active adversary that is able to inject data over communication channel to read or modify patient data.
2) IMD, programmer, and gateway are honest entities of the network. These entities conform to communication protocol.
3) An attacker is a single and an external entity with respect to the network.
4) An attacker does not have physical access to an active IMD.
5) IMD, gateway and the attacker are in each other's range. The attacker has to be in IMD's range to launch an attack. The gateway is also within IMD's and attacker's range.
6) An attacker may either be a sophisticated dedicated device, or a simple device with limited resources.

IV. Our intrusion detection system

In this section, we propose machine learning based attack detection for securing wireless medical devices from malicious entities. We rely on an external gateway node to identify and detect attacks on the medical device. The gateway monitors all wireless communication between the IMD device and other entities in the network and detects any abnormal or irregular network behavior. An irregularity or an abnormal behavior of a network entity is defined as, any deviation from its normally

observed behavior. The deviation can occur in the physical characteristics of a signal transmitted from the device or the frequency of signal emissions from the device. When the gateway detects an attack, it initiates a passive response by sending a warning signal to the patient's personal device such as a mobile phone or a laptop. Moreover, the gateway ensures not to interfere with any ongoing communication between entities in the medical system.

We first propose a specification based intrusion detection system using decision tree algorithm to address the security attacks on the medical devices. We then compare the decision tree algorithm with other machine learning algorithms.

A. Decision-Tree Learning

Decision tree is a type of classification algorithm. The predictive model maps observations about an attribute to conclusions about the attribute's target value. Non-terminal nodes in the tree are used to represent attributes and terminal nodes are used to represent decision outcome of the algorithm. In this paper, nodes of the tree are monitored security attributes. Each branch represents outcome of monitored security attribute, whereas leaf nodes are the result of classification as normal or abnormal.

During the training phase, decision-tree learning is used to general a normal profile. As the model learns the behavior of the network, it modifies and updates the profile. In this paper, we develop a new set of features that include attributes specific to security as input to the learning model. A security attribute (X_i) is a transmission's physical or logical property. In this paper we focus on the following security attributes to build the decision tree:

- X_1= Type of action performed by the programmer
- X_2= Time when action is performed
- X_3= Number of times a same type of action occurs
- X_4= Elapsed time since last occurrence of same action
- X_5= Received signal strength indicator
- X_6= Type of Day
- X_7= Location

A collected measurement of all security attributes in our scheme is defined as a record: R = $\{X_1,X_2,X_3,X_4,X_5,X_6,X_7\}$.

Our detection scheme uses decision tree classification algorithm to decide if a particular record is benign or malicious. Developing the decision tree requires sufficient data to train the program. We define a set of security attributes as described below and monitor them for a substantial period of time to build a normal behavior profile. The features were selected based on the typical usage of a medical device.

- Type of action: This attribute specifies the type of action performed on an IMD. Depending on the type of the IMD that is accessed, these actions may vary in nature. For instance, in a pacemaker, examples of IMD actions include setting the device clock, changing its settings or accessing the medical data.
- Time of action: This attribute specifies the exact time at which a type of action is requested from an IMD by the programmer. The time of action can be any time during the day, unless the IMD is accessed in case of an emergency situation during the night.
- Number of action occurrences: This attribute specifies the number of occurrences of an action. The value varies depending on the type of IMD and is often limited to a definite value for a time period. For example, if a programmer is designed to access pacemaker data at a periodic interval during the day, deviation from this behavior can indicate an attack.
- Time interval since last occurrence of an action: This attribute specifies the time elapsed since the last occurrence of an action and helps to determine the frequency of that action. This feature can be used to evaluate if there is a change in frequency of a given action.
- Signal strength indicator: This attribute provides the signal strength or a received transmission. We use signal strength indicator to detect an attack in cases where, the signal transmitted by external device has abnormally low or high strength. Since the attacker has to be in the communication range of IMD device to perform an attack, the signal strength is expected to fall within certain range. If the received signal strength is significantly different from the pre-determined range, it could indicate an attack.
- Day: The type of day is one of the most significant attributes for detecting an attack on IMDs. Usually, the doctor or the patient or the patient's family/friends are the only users that access an IMD to record readings. This type of access has a stable frequency such that, the users access an IMD either on a weekday or a weekend or a holiday. Our scheme detects an attack if the attacker tries to access data from IMD on an odd day.
- Location: Location attribute denotes the location of a device that accesses an IMD. This attribute makes sure a certain type of actions is carried out in certain locations. For example, a doctor's clinic is a location where an action of changing functional settings of a pacemaker is performed. If the same action is performed from any other location apart from the doctor's clinic, the scheme detects an attack in action.

The proposed approach builds a decision tree based on the above interval ranges of monitored attributes (Normal Behavior). If an anomaly occurs, abnormal value record is detected and classified as normal or abnormal. For attributes that fall below or above a threshold, we check difference between predicated measurement and current measurement. If that difference is greater than the specified threshold, we detect abnormal record. If the difference is larger than the threshold, gateway raises a warning that an attack is happening.

V. Evaluation and Results

In this section, we discuss the simulations and experiments conducted to evaluate the performance of the proposed decision tree based intrusion detection system and provide a comparison with SVM and K-means algorithms. We used Castalia,

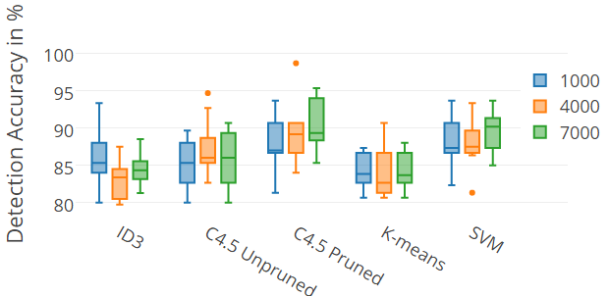

Fig. 2. Box and whisker diagram of detection Accuracy after 10-fold Cross Validation

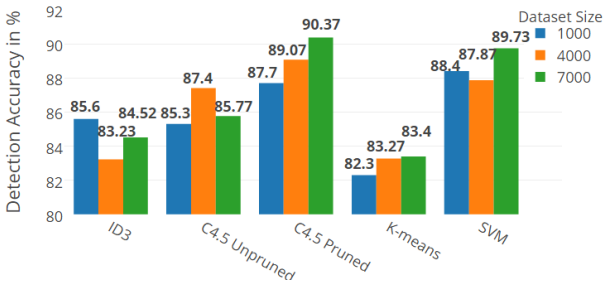

Fig. 3. Comparison of Average Detection Accuracy Against Resource Depletion Attacks for ID3 and C4.5

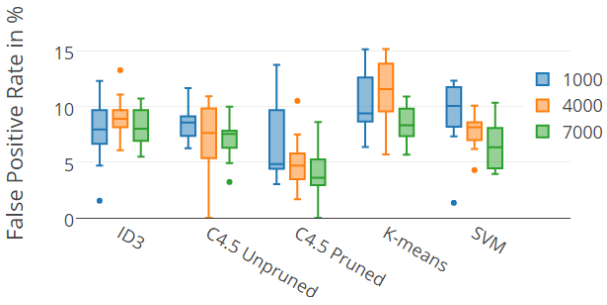

Fig. 4. False Positive Rate after 10-fold Cross Validation

the wireless body area network simulator for our experiments. Nodes in our simulation network includes the IMD device, a legitimate programmer and a malicious programmer. For our experiments, we first pre-processed the patient's access data. Recall that the patient's access data are denoted as a record: R = $\{X_1,X_2,X_3,X_4,X_5,X_6,X_7\}$, representing reader action type, time when action is performed, number of times a same type of action occurs, the time interval of the same action, signal strength, day, location, respectively. For X_1, we denote {Authenticate, Execute, Read, Write} as the four different types of action. For X_2, we label 24 values of the hours as {1,2,3,...,23,24}. For X_3, there are 3 categories, {1,1-5,¿5}. For X_4, we classify them to three categories as well {Less than one day,Less than three days, Less than a week}. For X_5, we categorize it into {Low, Normal, Strong}. For X_6, we have the 7 days in week {Mon, Tue, Wed, Thur, Fri, Sat, Sun}. For X_7, we have {Hospital1, Hospital2, Home, Office, Unknown}.

We used three different datasets with the total sample size of 1000, 4000 and 7000 respectively. For each dataset, 80% of the data were used as training data and the remaining 20% samples for testing. We tested the performance of our approach on decision tree, K-means and SVM algorithms. We used 10-fold cross-validation to determine the accuracy of the descision-tree based detection model. In k-fold cross-validation, the original sample is randomly partitioned into k equal size sub-samples. Of the k sub-samples, a single sub-sample is retained as the validation data for testing the model, and the remaining k-1 subsamples are used as training data. The cross-validation process is then repeated k times, with each of the k sub-samples used exactly once as the validation data. The advantage of this method is that all observations are used for both training and validation, and each observation is used for validation exactly once.

A. Detection Accuracy

Figure 2 shows the distribution of the results after the 10-fold cross validation. Figure 3 shows on average pruned C4.5 has the best permanence among other algorithms. ID3 achieves an average accuracy of 85.6%, 83.23% and 84.52% on different datasets. While C4.5 unpruned have 85.3%, 87.4% and 85.77% and C4.5 pruned has the best performance of 87.7%, 89.07% and 90.37% on our 3 datasets. We see from our results that K-means had weaker performance. We also note that the size of the dataset had no direct coorelation with the detection accuracy of the detection.

B. False Positive Rate

Figure 4 shows the distribution of the false positive rate after the 10-fold cross validation. Figure 5 illustrates the comparison of average false positive rate for all algorithms in consideration. ID3 achieves an average false positive rate of 7.95%, 9.1% and 8.11%. C4.5 unpruned did slightly better with 8.44%, 7.16% and 7.11% respectively. C4.5 pruned has a tremendously low false positive rate especially as the size of the dataset increases. It has 5.09% on the 4000 dataset and 4.09% on the 7000 dataset. K-means unsurprisingly has the highest false positive rate. For most algorithms in our experiment, a larger dataset resulted in lower false positive rate.

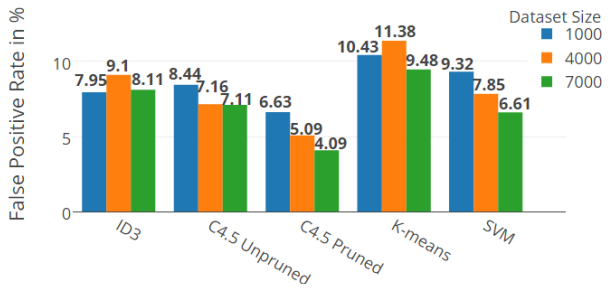

Fig. 5. Comparison of Average False Positive Rat for ID3 and C4.5

C. Training Speed and Prediction Speed

We also measured the training speed and prediction speed of the different detection algorithms. Figure 6 shows the algorithm overhead for different datasets. Both ID3 and SVM require some data transformation a real life scenario given their limitations. These algorithms also demonstrated slower training speed. C4.5 has a lower overhead, especially with pruning because prunning will reduce the complexity of tree. Figure 7, presents the comparison of prediction speed over different datasets. As shown in the results, ID3 and K-means have a relatively slower speed in terms of detection. The speed of K-means algorithm also depends on K, which in our case is only equal to 2. As observed from the results, despite the fact that SVM has a slow training speed, once the support vectors are generated, the prediction speed for SVM is the fastest among other algorithms.

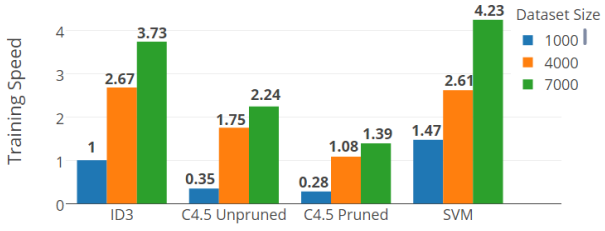

Fig. 6. Comparison of Training Speed

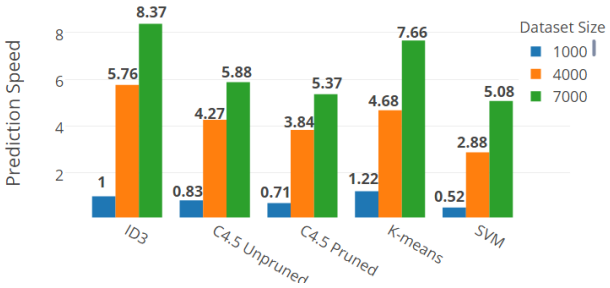

Fig. 7. Comparison of Prediction Speed

VI. Conclusion

In this paper, we evaluated the feasibility of using machine learning algorithms to detect security attacks in implantable wireless medical devices. We developed a feature set specific to the use of IMD devices and conducted experiments to test the performance of different learning algorithms including decision tree, SVM and K-means algorithms. Our results show that decision tree based algorithms achieve the highest detection accuracy, low false positive rate, fast training and prediction speed among all other algorithms. However our approach fails if the attacker is an insider and more familiar with the schedule and patterns of data access to/from the medical device. Future work includes expanding our feature set to address this limitation.

References

[1] W. H. Maisel, "Safety issues involving medical devices: implications of recent implantable cardioverter-defibrillator malfunctions," *JAMA*, vol. 294, no. 8, pp. 955–958, 2005.

[2] C. Li, A. Raghunathan, and N. K. Jha, "Hijacking an insulin pump: Security attacks and defenses for a diabetes therapy system," in *e-Health Networking Applications and Services (Healthcom), 2011 13th IEEE International Conference on*, pp. 150–156, IEEE, 2011.

[3] D. Halperin, T. S. Heydt-Benjamin, K. Fu, T. Kohno, and W. H. Maisel, "Security and privacy for implantable medical devices," *IEEE pervasive computing*, vol. 7, no. 1, pp. 30–39, 2008.

[4] D. R. Raymond and S. F. Midkiff, "Denial-of-service in wireless sensor networks: Attacks and defenses," *IEEE Pervasive Computing*, vol. 7, no. 1, pp. 74–81, 2008.

[5] S. Gollakota, H. Hassanieh, B. Ransford, D. Katabi, and K. Fu, "They can hear your heartbeats: non-invasive security for implantable medical devices," *ACM SIGCOMM Computer Communication Review*, vol. 41, no. 4, pp. 2–13, 2011.

[6] L. Berko, "We need to make implantable medical devices more secure," 2013.

[7] X. Hei and X. Du, "Biometric-based two-level secure access control for implantable medical devices during emergencies," in *INFOCOM, 2011 Proceedings IEEE*, pp. 346–350, IEEE, 2011.

[8] F. Xu, Z. Qin, C. C. Tan, B. Wang, and Q. Li, "Imdguard: Securing implantable medical devices with the external wearable guardian," in *INFOCOM, 2011 Proceedings IEEE*, pp. 1862–1870, IEEE, 2011.

[9] X. Hei, X. Du, J. Wu, and F. Hu, "Defending resource depletion attacks on implantable medical devices," in *Global Telecommunications Conference (GLOBECOM 2010), 2010 IEEE*, pp. 1–5, IEEE, 2010.

[10] T. Denning, K. Fu, and T. Kohno, "Absence makes the heart grow fonder: New directions for implantable medical device security.," in *HotSec*, 2008.

[11] M. R. Rieback, B. Crispo, and A. S. Tanenbaum, "Rfid guardian: A battery-powered mobile device for rfid privacy management," in *Australasian Conference on Information Security and Privacy*, pp. 184–194, Springer, 2005.

[12] N. Henry, N. Paul, and N. McFarlane, "Using bowel sounds to create a forensically-aware insulin pump system," in *Presented as part of the 2013 USENIX Workshop on Health Information Technologies*, 2013.

[13] M. Zhang, A. Raghunathan, and N. K. Jha, "Medmon: Securing medical devices through wireless monitoring and anomaly detection," *IEEE transactions on biomedical circuits and systems*, vol. 7, no. 6, pp. 871–881, 2013.