Robust and Energy Efficient Malware Detection for Robotic Cyber-Physical Systems

Upinder Kaur
kauru@purdue.edu
Purdue University
West Lafayette, Indiana, USA

Z. Berkay Celik
zcelik@purdue.edu
Purdue University
West Lafayette, Indiana, USA

Richard M. Voyles
rvoyles@purdue.edu
Purdue University
West Lafayette, Indiana, USA

2022 ACM/IEEE 13th International Conference on Cyber-Physical Systems (ICCPS) | 978-1-6654-0967-4/22/$31.00 ©2022 IEEE | DOI: 10.1109/ICCPS54341.2022.00048

Abstract

Cyber-Physical Systems (CPS) increasingly use multiple robots as edge devices to enhance their functionalities. However, this introduces new security vulnerabilities such as control channel attacks and false data injection that an adversary can exploit to put the users and environment at risk. In this paper, we build a robust malware detection system strengthened by carefully crafted adversarial samples. We generate adversarial samples within the bounds of domain constraints and integrate them into model training to improve the model's robustness. Additionally, we formulate an objective function to distribute the computation of malware detection to multiple edges, making optimal use of the robot mesh network to reduce power consumption. In the adjoining poster, we show the details of the dataset and the models, and illustrate the specifics of our contributions.

Keywords: security, cyber-physical systems, robotics

1 INTRODUCTION

The proliferation of CPS with robotic edges, such as mobile robot networks in agriculture and manufacturing, demands comprehensive scrutiny to safeguard users. Recent efforts have largely focused on strengthening the performance of physical and communication layers of robotic systems while its security against malware has been vastly understudied. With their increasing adoption, these systems become an attractive target for false data injection and control channel attacks. Such attacks are particularly dangerous since an adversary can completely alter the behavior of the robot and as robots operate near living beings, a deviation from its expected behavior can prove to be fatal [3].

The threat of physical harm demands a robust defense in robotic systems. However, their large-scale use such as in agriculture, coupled with the limited compute power on low-level embedded devices, makes the application of traditional security measures, such as remote attestation, cryptography and access control, widely impractical [1]. Hence, the recent advances in machine learning (ML) allows for a potential defense technique to detect such attacks as it can be scaled to embedded hardware and updated easily. In this work, we address the two critical aspects of deploying ML models for security in robotic systems:

Robustness. While ML models have been used to detect diverse malware [1], they are known to be susceptible to adversarial samples, which are carefully crafted feature values with the goal of confusing a model, resulting in misclassification. We craft adversarial samples that respect the constraints of the domain and integrate them into the model training to improve the robustness of the malware-detecting ML models.

Optimization for Energy Consumption. The interconnectivity afforded by the robot mesh network can be leveraged to reduce the computational load of the nodes in applying the ML models. This can be achieved by distributing the computation task and communicating the results over the mesh network. To do so, we formulate this goal as an optimization problem, which aims at minimizing the communication cost with respect to the overall energy consumption of robots.

2 PROPOSED FRAMEWORK

We propose a novel robust security framework that protects robotic systems from malware attacks while optimizing for power and computation. Recently, we introduced the RoboMal dataset, a collection of malicious and benign controller software for a mobile robots, along with a baseline LSTM detection model [1]. Our work enables ML-based malware detection in robotic systems with carefully crafted diverse attacks. In this work, we extend our previous framework to include adversarial examples into model training to defend against adversarial examples.

2.1 Malware Detection with Adversarial Training

We consider an adversary, which can manipulate the controller code and aims to make minimal changes to it to cause maximal damage in the system. We assume the adversary has a black-box or white-box knowledge of the model and aims to craft samples to evade detection. To encounter such attacks, we incorporate adversarial examples into the training dataset. However, unlike traditional adversarial training in other domains, such as in image recognition wherein images are perturbed arbitrarily and independently, adversarial examples in robotic systems must adhere to domain constraints [4] to represent legitimate samples. For example, a crafted malicious code should successfully compile without any errors to be effective. To address this issue, we use a

978-1-6654-0967-4/22/$31.00 ©2022 IEEE
DOI 10.1109/ICCPS54341.2022.00048

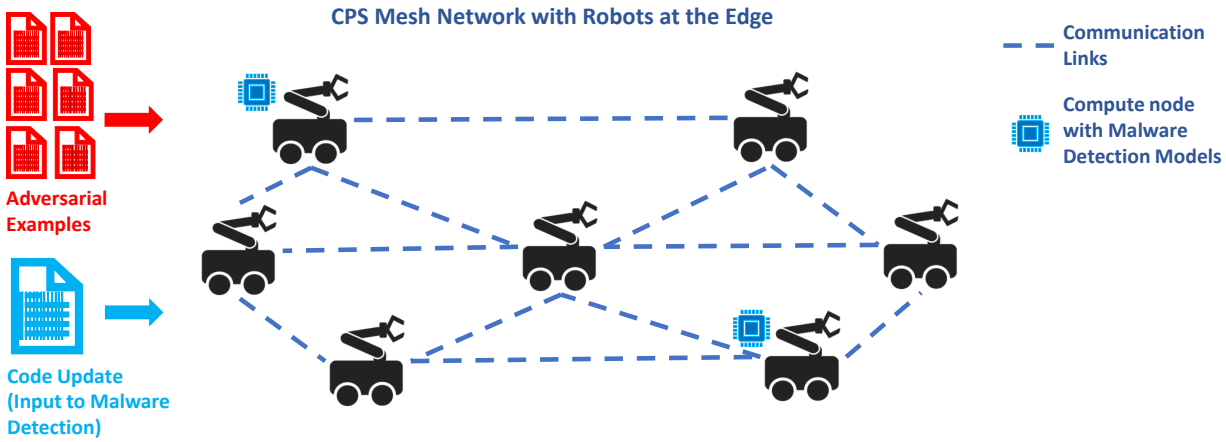

Figure 1. An example CPS mesh network with robot edges where the computation is distributed across nodes.

virtual robot simulator to verify the viability of adversarial examples, thereby identifying a valid subset that can be used as an input while training the model to improve its robustness and generalization. We also measure their similarity, using l_2 norm, to the original dataset to identify the adversarial samples.

2.2 Optimization Formulation

The robot mesh network affords the capability to distribute the computation across edges so that not all nodes have to expend energy doing the same computation for malware detection. To formulate an objective function, we consider a mesh network of homogeneous robot nodes, as shown in Fig. 1, built on the IEEE 802.15.5 standard protocol [2] with a tree structure rooted in network gateway. Our goal is to minimize the total energy consumption of each node by minimizing the maximum and overall costs of communication. To this aim, we define the following objective:

$$\mathtt{Cost} = \min_{x_i} \sum_{i \in N} \left(\sum_{o \in O} (d_{i,o}) + \max_{o \in O} (d_{i,o}) \right) x_i$$

where $x_i = 1$ if i is a computing node, otherwise 0. The routing path decides the transmit power, $d_{i,j}$. We define the set of constraints on the objective function:

$$\sum_{i \in O} \sum_{i \in N} \left(1_{P_{-i,o}}(n) E_r + 1_{P_{-i,-o}}(n) E_t \right) x_i + E_t |O| \leq E_{comp}(n)$$

$$E_{comp}(n) = E_T - E_{comm}(n) - E_{PHY}(n)$$

$$E_{comm}(n) = n_r * E_r(n) - n_t * E_t(n)$$

$$1 \leq i \leq N, x_i \in \{0, 1\}, O \in N$$

where, $1_{P_{-i,o}}$ returns 1 if $n \neq i$ belongs to the path from compute node to output, otherwise returns 0. Similarly, $1_{P_{-i,-o}}$ returns 1 if i is an intermediate and $o \in O$ output nodes, otherwise returns 0. For the $n \in N$ nodes, the total available energy is E_T with $E_t, E_r, E_{comm}, E_{comp}$ and E_{PHY} being the energy for transmission, receiving, total communication, computation, and for physical tasks, respectively. Number of received and transmitted messages are n_r and n_t, respectively.

3 SUMMARY AND FUTURE WORK

In this paper, we propose a robust distributed framework for malware detection for robotic Cyber-Physical Systems. We craft adversarial samples that respect the domain constraints and use them in model training to strengthen the ML models against adversarial samples. For future work, we will assess the efficacy of the crafted samples using the available baseline ML models. Further, leveraging the mesh network, we optimize the energy expenditure by distributing the computation of detection among different nodes. We plan to test the optimization on a CPS with mobile robots and analyze the impacts on power consumption and latency.

Acknowledgments

The authors acknowledge the support of USDA Grant 2018-67007-28439 in the fulfillment of this work.

References

[1] Upinder Kaur, Haozhe Zhou, Xiaxin Shen, Byung-Cheol Min, and Richard M Voyles. 2021. RoboMal: Malware Detection for Robot Network Systems. In *Proceedings of the Fifth IEEE International Conference on Robotic Computing*.

[2] Myung Lee, Rui Zhang, Jianliang Zheng, GS Ahn, Chunhui Zhu, Tae Park, Sung Cho, Chang Shin, and Jun Ryu. 2010. IEEE 802.15. 5 WPAN mesh standard-low rate part: Meshing the wireless sensor networks. *IEEE Journal on Selected Areas in Communications* 28 (2010), 973–983.

[3] Santiago Morante, Juan Victores, and Carlos Balaguer. 2015. Cryptobotics:Why robots need cyber-safety. *Frontiers in Robotics & AI* (2015).

[4] Ryan Sheatsley, Blaine Hoak, Eric Pauley, Yohan Beugin, Michael J Weisman, and Patrick McDaniel. 2021. On the robustness of domain constraints. In *Proceedings of the 2021 ACM SIGSAC Conference on Computer and Communications Security*. 495–515.