Lecture Notes in Computer Science   **11281**

More information about this series at

Vinod Ganapathy · Trent Jaeger
R. K. Shyamasundar (Eds.)

Information Systems Security

14th International Conference, ICISS 2018
Bangalore, India, December 17–19, 2018
Proceedings

Editors
Vinod Ganapathy
Indian Institute of Science
Bangalore, India

R. K. Shyamasundar
Indian Institute of Technology Bombay
Mumbai, India

Trent Jaeger
Pennsylvania State University
University Park, PA, USA

Preface

This volume contains the papers presented at the 14th International Conference on Information Systems Security (ICISS 2018), held at the Indian Institute of Science, Bengaluru, Karnataka, India, during December 17–19, 2018. In response to the call for papers, 53 submissions were received. One submission was withdrawn by the authors while a second one was desk-rejected, leaving a total of 51 papers to be evaluated by the Program Committee. All submissions were evaluated on the basis of the novelty and significance of their scientific content. The Program Committee, comprising 47 members, reviewed all submissions via a single-blind review process. Each paper was reviewed by three or more reviewers and discussed. After discussions, the committee selected a total of 23 papers for presentation at the conference. The papers covered a wide range of topics as is reflected in the list of papers in this volume.

In addition to the peer-reviewer papers, we were also fortunate to have four eminent speakers delivering keynote presentations at the conference: Venkatramanan Siva Subrahmanian (Dartmouth University), Atul Prakash (University of Michigan–Ann Arbor), Prateek Saxena (National University of Singapore), and Sriram Rajamani (Microsoft Research India). The program also consisted of two tutorials: one by Nishanth Chandran and Divya Gupta (Microsoft Research India), and a second one by Somesh Jha (University of Wisconsin–Madison). We are really thankful to these speakers for taking time off from their busy schedules and contributing to ICISS 2018.

Many individuals contributed to making ICISS 2018 a success, and it is a pleasure to thank them. We are thankful to all members of the Program Committee and all external reviewers for their efforts in reviewing the papers and participating in the discussion and selection process. We thank the Steering Committee, our publicity chair (Anu Mary Chacko of NIT Calicut), and the Web team (Rounak Agarwal, Aditya Shukla and Kripa Shanker, of IISc Bangalore). We thank Kushael and Shankar from the Department of Computer Science and Automation, IISc Bangalore, for providing administrative support and taking care of many logistical details. Needless to say, we are thankful to all the authors who submitted their work to ICISS 2018, and our keynote speech and tutorial speakers who accepted our invitation to present at the conference.

We were fortunate to receive financial support for the conference from Sonata Software, IISc Bangalore, and Microsoft Research India. We are thankful to Omprakash Subbarao (Sonata Software), Y. Narahari (IISc Bangalore), and the team of Chiranjib Bhattacharyya (IISc Bangalore), Sriram Rajamani, Satish Sangameswaran, and Christina Gould-Sandhu (Microsoft Research India) for providing financial sponsorship. We also appreciate the support of Springer, in particular Alfred Hofmann and Anna Kramer, in publishing the proceedings as well as the monetary support for the conference. We would also like to acknowledge EasyChair for their conference management system, which was freely used to manage the process of paper submissions and reviews.

We hope that you find these proceedings interesting and useful in your own research.

October 2018

Vinod Ganapathy
Trent Jaeger
R. K. Shyamasundar

Organization

Program Committee

Vijay Atluri	Rutgers University, USA
Gogul Balakrishnan	Google, USA
Arati Baliga	Persistent Systems, India
Mridul Sankar Barik	Jadavpur University, India
Lorenzo Cavallaro	Royal Holloway, London, UK
Sambuddho Chakravarty	IIIT-Delhi, India
Frederic Cuppens	IMT Atlantique, France
Ashok Kumar Das	IIIT-Hyderabad, India
Lorenzo DeCarli	Worcester Polytechnic Institute, USA
Rinku Dewri	University of Denver, USA
Mohan Dhawan	IBM Research India
Adam Doupe	Arizona State University, USA
Earlence Fernandes	University of Washington, USA
Vinod Ganapathy	IISc Bangalore, India
Vijay Ganesh	University of Waterloo, Canada
Siddharth Garg	New York University, USA
Kanchi Gopinath	IISc Bangalore, India
Trent Jaeger	Pennsylvania State University, USA
Sushil Jajodia	George Mason University, USA
Suman Jana	Columbia University, USA
Aniket Kate	Purdue University, USA
Ram Krishnan	University of Texas-San Antonio, USA
Subhomoy Maitra	ISI-Kolkata, India
Pratyusa Manadhata	HP Labs, USA
Debdeep Mukhopadhyay	IIT-Kharagpur, India
Divya Muthukumaran	Imperial College London, UK
Adwait Nadkarni	College of William and Mary, USA
Eiji Okamoto	University of Tsukuba, Japan
Biswabandan Panda	IIT-Kanpur, India
Arpita Patra	IISc Bangalore, India
Goutam Paul	ISI-Kolkata, India
Phu Phung	University of Dayton, USA
Atul Prakash	University of Michigan–Ann Arbor, USA
Indrakshi Ray	Colorado State University, USA
Bimal Roy	ISI-Kolkata, India
Diptikalyan Saha	IBM-Research India
R. Sekar	StonyBrook University, USA
Sandeep Shukla	IIT-Kanpur, India

Anoop Singhal	National Institute of Standards and Technology, USA
Arunesh Sinha	University of Michigan–Ann Arbor, USA
Pramod Subramanyan	IIT-Kanpur, India
Laszlo Szekeres	Google, USA
Mohit Tiwari	University of Texas at Austin, USA
Mahesh Tripunitara	University of Waterloo, Canada
Venkatakrishnan V. N.	University of Illinois, Chicago, USA
Hayawardh Vijayakumar	Samsung Research, USA
Stijn Volckaert	University of California, Irvine, USA
Vinod Yegneswaran	SRI International, USA

Additional Reviewers

Bruhadeshwar Bezawada
Ayantika Chatterjee
Warren Connell
Sabrina De Capitani di Vimercati
Akshar Kaul
Manish Kesarwani
Haining Wang

Keynote Abstracts

Bots, Socks, and Vandals: An Overview
of Malicious Actors on the Web

V. S. Subrahmanian

Department of Computer Science and Institute for Security, Technology,
and Society, Dartmouth College, Hanover NH 03755, USA
vs@dartmouth.edu

Abstract. In this paper, we discuss four types of malicious actors on social platforms and online markets: bots, sockpuppets, individuals committing review fraud, and vandals.

Online social networks and e-commerce platforms are increasingly targeted by malicious actors with a wide variety of goals. Bots on Twitter may seek to illicitly influence opinion. Sock-puppet accounts on online discussion forums (e.g. discussion threads on online news articles) may help push certain points of view. Vandals on Wikipedia may seek to inject false material into otherwise legitimate pages. Review fraud in online forums may illicitly promote a product or destroy a competing product's reputation.

The bulk of this talk will focus on identifying review fraud in online e-commerce platforms such as Amazon, eBay and Flipkart. Because an increase of 1 star in a product rating can, on average, lead to a 5–9% increase in revenues, vendors have strong incentives to generate fake reviews. We will present both an unsupervised model as well as a supervised model to identify users who generate fake reviews. We show that our framework, called REV2 [3], produces high performance in real world experiments. In addition, a report of 150 review fraud accounts on Flipkart was independently evaluated by Flipkart's anti-fraud team who reported that 127 of the predictions were correct.

Sockpuppet accounts – multiple accounts operated by a single individual or corporate "puppetmaster" – are also a popular mechanism used to inappropriately sway opinion in online platforms. For instance, social "botnets" [4] commonly use multiple "sock" accounts to implement coordinated bots. Sockpuppet accounts are also commonly used by trolls. I will report on recent work [2] on the characteristics and properties of sockpuppet accounts through a study that involves data from the Disqus platform. Disqus powers discussion threads and forums on a host of news and other websites. Sockpuppets are often used in such contexts to artificially boost an opinion or artificially generate controversy. I will also briefly discuss the use of bots in real world influence campaigns along with methods to detect them [1, 4].

Third, I will discuss the problem of vandals on Wikipedia. Though research has been done previously on automated methods to detect acts of vandalism on Wikipedia, we describe VEWS, a Wikipedia Vandal Early Warning System that seeks to detect vandals as early as possible and preferably before they commit any acts of vandalism.

We show that VEWS outperforms prior work – but that when combined with prior work, it predicts vandals with very high accuracy.

The talk will conclude with a discussion of different types of malicious actors on the web.

References

1. Dickerson, J.P., Kagan, V., Subrahmanian, V.: Using sentiment to detect bots on twitter: are humans more opinionated than bots? In: Proceedings of the 2014 IEEE/ACM International Conference on Advances in Social Networks Analysis and Mining, pp. 620–627. IEEE Press (2014)
2. Kumar, S., Cheng, J., Leskovec, J., Subrahmanian, V.: An army of me: sockpuppets in online discussion communities. In: Proceedings of the 26th International Conference on World Wide Web, pp. 857–866. International World Wide Web Conferences Steering Committee (2017)
3. Kumar, S., Hooi, B., Makhija, D., Kumar, M., Faloutsos, C., Subrahmanian, V.: Rev2: fraudulent user prediction in rating platforms. In: Proceedings of the Eleventh ACM International Conference on Web Search and Data Mining, pp. 333–341. ACM (2018)
4. Subrahmanian, V.S., et al.: The darpa twitter bot challenge. Computer **49**(6), 38–46 (2016). https://doi.org/10.1109/MC.2016.183

Robust Physical-World Attacks on Deep Learning Visual Classifiers and Detectors (Invited Talk)

Atul Prakash

Computer Science Division, University of Michigan,
Ann Arbor, MI 48109, USA
aprakash@umich.edu

Abstract. Recent studies show that the state-of-the-art deep neural networks (DNNs) are vulnerable to adversarial examples, resulting from small-magnitude perturbations added to the input [1, 5, 6, 8, 10, 12]. Given that that emerging physical systems are using DNNs in safety-critical situations such as autonomous driving, adversarial examples could mislead these systems and cause dangerous situations. It was however unclear if these attacks could be effective in practice with real-world objects [7], with some researchers finding that the attacks fail to translate to physical world in practice [9]. We report on some of our findings [2, 3] for generating such adversarial examples that can be physically realized using techniques such as stickers placed on real-world traffic signs. With a perturbation in the form of only black and white stickers, we modified real stop signs, causing targeted misclassification in over 80% of video frames obtained on a moving vehicle (field test) for state-of-the-art image classifiers, LISA-CNN and GTSRB-CNN. Our recent results [4] suggest that object detectors, such as YOLO [11], are also susceptible to physical perturbation attacks. I discuss some of the implications of the work on the design of robust classifiers and detectors for safety-critical applications.

Keywords: Adversarial machine learning · Input perturbation
Physical attacks · Deep learning · Security · Robust classifiers · Robust detectors

References

1. Carlini, N., Wagner, D.: Towards evaluating the robustness of neural networks. In: 2017 IEEE Symposium on Security and Privacy (SP), pp. 39–57. IEEE (2017)
2. Eykholt, K., et al.: Robust physical-world attacks on deep learning visual classification. In: Proceedings of Computer Vision and Pattern Recognition Conference (CVPR 2018). IEEE, June 2018. (Supersedes arXiv preprint arXiv:1707.08945, August 2017)
3. Eykholt, K., et al.: GitHub Repo on Robust Physical Perturbations Code. https://github.com/evtimovi/robust_physical_perturbations
4. Eykholt, K., et al.: Attacking object detectors with adversarial stickers. In: Proceedings of 12th Usenix Workshop on Offensive Technologies (WOOT), Baltimore, MD, (arXiv:1807.07769) (supersedes arXiv:1712.08062), August 2018

5. Goodfellow, I.J., Shlens, J., Szegedy, C.: Explaining and harnessing adversarial examples. arXiv preprint arXiv:1412.6572 (2014)
6. Kos, J., Fischer, I., Song, D.: Adversarial examples for generative models. arXiv preprint arXiv:1702.06832 (2017)
7. Kurakin, A., Goodfellow, I., Bengio, S.: Adversarial examples in the physical world. arXiv preprint arXiv:1607.02533 (2016)
8. Liu, Y., Chen, X., Liu, C., Song, D.: Delving into transferable adversarial examples and black-box attacks. arXiv preprint arXiv:1611.02770 (2016)
9. Lu, J., Sibai, H., Fabry, E., Forsyth, D.: No need to worry about adversarial examples in object detection in autonomous vehicles. arXiv preprint arXiv:1707.03501 (2017)
10. Papernot, N., McDaniel, P., Jha, S., Fredrikson, M., Celik, Z.B., Swami, A.: The limitations of deep learning in adversarial settings. In: 2016 IEEE European Symposium on Security and Privacy (EuroS&P), pp. 372–387. IEEE (2016)
11. Redmon, J., Farhadi, A.: YOLO9000: better, faster, stronger. CoRR, abs/1612.08242 (2016)
12. Sabour, S., Cao, Y., Faghri, F., Fleet, D.J.: Adversarial manipulation of deep representations. arXiv preprint arXiv:1511.05122 (2015)

Specifying and Checking Data Use Policies

Sriram K. Rajamani

Microsoft Research, Bengaluru, India
`sriram@microsoft.com`

Cloud computing has changed the goals of security and privacy research. The primary concerns have shifted to protecting data in terms of not only who gets to access data, but also how they use it. While the former can be specified using access control logics, the latter is relatively a new topic and relatively unexplored.

We describe a language called Legalese, which we designed to specify data use policies in cloud services. Legalese [1] uses propositional logic together with type-state to specify constraints on data use, retention and combination of data. Next, we describe a notion called Information Release Confinement (IRC) [2], which can be used to specify that data does not leave a region except through specific channels such as API calls. IRC has been used to specify and verify confidentiality of cloud services that use Intel SGX enclaves. Finally, we speculate on combining these two approaches to specify and check stateful policies on data use in cloud services.

References

1. Sen, S., Guha, S., Datta, A., Rajamani, S.K., Tsai, J.Y., Wing, J.M.: Bootstrapping Privacy Compliance in Big Data Systems. In: IEEE Symposium on Security and Privacy, pp. 327–34 (2014)
2. Sinha, R., et al.: A design and verification methodology for secure isolated regions. In: PLDI, pp. 665–681 (2016)

Contents

Privacy

Client Security and Authentication

Invited Keynote