Guest Editors' Introduction

ÖZGÜR KAFALI, University of Kent
 NATALIA CRIADO, King's College London
 MARTIN REHAK, Cisco Systems
 JOSE M. SUCH, King's College London
 PINAR YOLUM, Utrecht University

ACM Reference format:

Özgür Kafalı, Natalia Criado, Martin Rehak, Jose M. Such, and Pinar Yolum. 2018. Guest Editors' Introduction. *ACM Trans. Internet Technol.* 18, 3, Article 26 (March 2018), 4 pages.
<https://doi.org/10.1145/3177884>

Cybersecurity underpins the lives of ordinary people—their safety, work, health, and entertainment. Yet despite its importance, cybersecurity is often approached in a reactive manner—taking corrective actions to “patch” vulnerabilities after they are detected or exploited. In the absence of fundamental improvements in formal processes and methods, the same or similar problems can recur. In contrast, a scientific vision for security and privacy should be *proactive*. The focus of this TOIT special section is to provide a forum to discuss and advance security and privacy by means of *Artificial Intelligence* (AI) approaches. Approaches that are intelligent and self-adaptive are crucial to deal with the complexities of effectively protecting sensitive assets in all security-critical domains. This is where research from the AI community can make a difference in security and privacy. Specifically, AI can help address a long-standing problem, namely, that security and privacy are attempted to be inserted as an afterthought, which is rarely adequate.

We identify the following key areas in AI, and review advancements in such areas that would help solve fundamental hard problems regarding security (CyBOK 2017) and privacy (Such 2017).

Normative models: Secure and privacy-aware governance of sociotechnical systems can only be achieved by bridging the divide between technical solutions to security and privacy such as access control systems and the human and social factors associated with the users of such systems. Development of unified computational models for the social and the technical tiers of sociotechnical systems as well as their verification are therefore of utmost importance. Previous work (Barth et al. 2006; Chopra et al. 2014; Criado and Such 2016; Kafalı et al. 2016a; Singh 2013) has shown that a formalization of norms captures sociotechnical system concepts by regulating user interactions. Norms serve as a computational method for establishing user accountability by capturing a precise

Authors' addresses: Ö. Kafalı, School of Computing, University of Kent, Canterbury, Kent, CT2 7NF, UK; email: R.O.Kafali@kent.ac.uk; N. Criado and J. M. Such, Bush House, King's College London, Strand Campus, 30 Aldwych, London, WC2B 4BG, UK; emails: {natalia.criado, jose.such}@kcl.ac.uk; M. Rehak, Cisco Systems, Karlovo namesti 10, 120 00 Prague 2, Czech Republic; email: marrehak@cisco.com; P. Yolum, Utrecht University, Princetonplein 5, 3584 CC, Utrecht, Netherlands; email: p.yolum@uu.nl.

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for third-party components of this work must be honored. For all other uses, contact the Owner/Author.

2018 Copyright is held by the owner/author(s).

1533-5399/2018/03-ART26

<https://doi.org/10.1145/3177884>

standard of interaction to which each user is held (Kafalı and Singh 2017), which no purely technical solution can establish. Enforcement of sanctions (Nardin et al. 2016) would enhance such accountability by providing compensation for norm violations and acting as deterrence against violating norms. Normative models can be helpful even in domains where norms are unknown *a priori*. For instance, in social media, in which roles and relationships between users are constantly evolving and norms tend to be implicit in collapsed contexts and communities, norm-based computational models have been shown to accurately learn social expectations about privacy following a bottom-up process (Criado and Such 2015).

Agent-based reasoning: Distributed AI techniques leveraging intelligent agents are increasingly applied to security and privacy problems. In the context of resolving privacy conflicts among multiple users, for example, game-theoretic mechanisms (Rajtmajer et al. 2017; Such and Rovatsos 2016) and argumentation frameworks (Fogues et al. 2017; Gao et al. 2016; Kökciyan et al. 2017) can help agents that act on behalf of their users to compute an optimal joint privacy policy based on the users' individual privacy preferences. This is especially important in cases where satisfying all user preferences is not possible due to competing needs. Formal reasoning on competing needs among liveness (usability) and safety (security and privacy) (Kafalı et al. 2017a) can help develop design patterns to transform sociotechnical system specifications by formally computing the tradeoffs among such needs. In practical domains, agents can automatically help users negotiate or set privacy-related permissions, such as which friends can access their data in social media (Misra and Such 2017) or which applications can access their smartphone data (Baarslag et al. 2017).

Human factors: Breaches are inevitable, no matter how secure a software system is or how well its users are trained (DoD 2015; Murphy 2015; Verizon 2017). This is even more critical in systems where the security or privacy of an individual depends on what others do, like in social media, where any of your friends can post information about you (Such et al. 2017), which is completely out of your control. Recent studies (Kafalı et al. 2017b) demonstrate the impact of human error in breaches, and formally analyze how well existing regulations account for such breaches via AI techniques such as semantic similarity and domain ontologies. Automating the extraction of information from valuable artifacts such as breach reports would help such analysis, since breach reports not only capture examples that the software and its users should avoid, but also describe actions to prevent, detect, and recover from future breaches.

Digital forensics: As the variety of breaches increases, such as those caused by a malicious or criminal attack, a system glitch, or human error (Partners 2016), so does the generated volume of system logs to help detect such breaches. Consequently, forensics analysis to sort through the logs and produce a diagnosis is becoming increasingly difficult. Mean times to identify and contain a data breach can reach 229 and 82 days, respectively (Institute 2016). During this invisible period, adversaries can exploit private data. Therefore, proactive approaches to detecting and diagnosing data breaches, by building forensics capabilities into software, are highly desirable. Typically, a digital forensics process (Sammons 2015) consists of three core phases: *Collection* of digital evidence; *examination* of evidence via monitoring and diagnosis; and *reporting* and validation of forensics hypotheses. Regarding collection of evidence, rather than blindly logging every user action, intelligent logging mechanisms can separate the behaviors of regular users from potential misusers (Kafalı et al. 2016b), anticipate their intentions (Geib and Goldman 2001; Sadri 2012), and determine the amount of detail to be logged accordingly.

Articles: We have included five articles in this TOIT special section, covering a diverse range of topics. These articles have been subject to a thorough and multiple-round peer review process with the reviews of at least three experts on the topic, and selected from 14 initial submissions.

In "Genetic Algorithms for Solving Problems of Access Control Design and Reconfiguration in Computer Networks," the authors propose a method based on genetic algorithms for the

problem of access control configuration and reconfiguration. The contribution is very important as access control configuration is an NP-hard problem and appropriate heuristics are needed, and it is a perfect example of how AI techniques can be used toward addressing long-standing security problems.

Coordinating agents' actions to achieve a common goal usually requires exchanging some information which can be of a private nature; for example, agents may prefer not to reveal information about their private processes and motivations. In "Quantifying Privacy Leakage in Multi-Agent Planning" the authors propose a general computational method to measure privacy leakage in multi-agent planning. This work presents the first approach for the formal treatment of privacy leakage in privacy preserving multi-agent planning. In the article, the authors illustrate how the approach can be applied to quantify leakage and demonstrate that, based on that, it is possible to establish comparisons between the information leaked by different multi-agent planner algorithms.

Security threat models, often represented as attack graphs or attack trees, are a formal means for enumerating all potential attacks regarding a system. In "Decision Networks for Security Risk Assessment," the authors use decision networks, an extension of Bayesian networks, for the probabilistic analysis of potential attacks and their consequences, thereby assessing risk in critical infrastructures. Their technique is experimented in a real SCADA system operating in a power grid as well as in a real-world scenario for a Border Gateway Protocol session.

In "SHARE: A Stackelberg Honey-Based Adversarial Reasoning Engine," the authors develop a foundation for an attacker-defender scenario, based on Stackelberg games, for formalizing how a "noisy rich" attacker (i.e., an adversary who learns from blocked attacks and keeps exploiting vulnerabilities) explores an enterprise network and attacks it. Experiments are performed on networks with varying sizes. To induce deception, the defender can modify the network by placing honey nodes and apparent vulnerabilities (gathered from National Institute of Standards and Technology's National Vulnerability Database).

Describing policies regarding data practices, for example, collection, use, and sharing of private data, in a formal manner is crucial for developing automated verification methods to enable secure and privacy-aware collaboration among organizations. In "Fine-Grained Access Control via Policy-Carrying Data," the authors describe a formal policy language based on first-order logic for specifying access control rules for managing sensitive data. The language is based on a normative model (permissions, obligations, and prohibitions) to associate policies with datasets. Using their formalization, one can monitor and verify whether certain data practices comply with the associated policy descriptions. The authors support their formalization with intuitive and realistic examples from security-critical domains such as healthcare, and provide an implementation of their policy language.

REFERENCES

- Tim Baarslag, Alper T. Alan, Richard Gomer, Muddasser Alam, Charith Perera, Enrico H. Gerding, and M. C. Schraefel. 2017. An automated negotiation agent for permission management. In *Proceedings of the 16th Conference on Autonomous Agents and MultiAgent Systems (AAMAS'17)*. 380–390.
- Adam Barth, Anupam Datta, John C. Mitchell, and Helen Nissenbaum. 2006. Privacy and contextual integrity: Framework and applications. In *Proceedings of the IEEE Symposium on Security and Privacy (SP'06)*. IEEE Computer Society, 184–198.
- Amit K. Chopra, Fabiano Dalpiaz, F. Başak Aydemir, Paolo Giorgini, John Mylopoulos, and Munindar P. Singh. 2014. Protos: Foundations for engineering innovative sociotechnical systems. In *Proceedings of the 18th IEEE International Requirements Engineering Conference (RE'14)*. IEEE Computer Society, 53–62.
- Natalia Criado and Jose M. Such. 2015. Implicit contextual integrity in online social networks. *Information Sciences* 325 (2015), 48–69.
- Natalia Criado and Jose M. Such. 2016. Selective norm monitoring. In *Proceedings of the 25th International Joint Conference on Artificial Intelligence (IJCAI'16)*. 208–214.
- CyBOK. 2017. The Cyber Security Body of Knowledge. Retrieved November 2017 from <https://www.cybok.org/>.

- DoD. 2015. The United States Department of Defense Cybersecurity Culture and Compliance Initiative. Retrieved September 2015 from <http://www.defense.gov/Portals/1/Documents/pubs/OSD011517-15-RES-Final.pdf>.
- Ricard L. Fogues, Pradeep K. Murukannaiah, Jose M. Such, and Munindar P. Singh. 2017. Sharing policies in multiuser privacy scenarios: Incorporating context, preferences, and arguments in decision making. *ACM Transactions on Computer-Human Interaction (TOCHI)* 24, 1 (2017), 5.
- Yang Gao, Francesca Toni, Hao Wang, and Fanjiang Xu. 2016. Argumentation-based multi-agent decision making with privacy preserved. In *Proceedings of the 15th Conference on Autonomous Agents and Multiagent Systems (AAMAS'16)*. 1153–1161.
- Christopher W. Geib and Robert P. Goldman. 2001. Plan recognition in intrusion detection systems. In *Proceedings of the DARPA Information Survivability Conference & Exposition II (DISCEX'01)*, Vol. 1. IEEE, 46–55.
- Ponemon Institute. 2016. Cost of Data Breach Study: Global Analysis. Retrieved November 2017 from <https://securityintelligence.com/cost-of-a-data-breach-2016/>.
- Özgür Kafalı, Nirav Ajmeri, and Munindar P. Singh. 2016a. Revani: Revising and verifying normative specifications for privacy. *IEEE Intelligent Systems* 31, 5 (Sept. 2016), 8–15.
- Özgür Kafalı, Nirav Ajmeri, and Munindar P. Singh. 2017a. Kont: Computing tradeoffs in normative multiagent systems. In *Proceedings of the 31st Conference on Artificial Intelligence (AAAI'17)*. 3006–3012.
- Özgür Kafalı, Jasmine Jones, Megan Petruso, Laurie Williams, and Munindar P. Singh. 2017b. How good is a security policy against real breaches? A HIPAA case study. In *Proceedings of the 39th International Conference on Software Engineering (ICSE'17)*. IEEE Computer Society, 530–540.
- Özgür Kafalı and Munindar P. Singh. 2017. Guest Editors' Introduction - Improving Cybersecurity: User Accountability and Sociotechnical Systems. (2017). IEEE Computing Now. <https://www.computer.org/web/computingnow/archive/improving-cybersecurity-april-2017-introduction>.
- Özgür Kafalı, Munindar P. Singh, and Laurie Williams. 2016b. Nane: Identifying misuse cases using temporal norm enactments. In *Proceedings of the 24th IEEE International Requirements Engineering Conference (RE'16)*. IEEE Computer Society, 136–145.
- Nadin Kökciyan, Nefise Yaglikci, and Pinar Yolum. 2017. An argumentation approach for resolving privacy disputes in online social networks. *ACM Transactions on Internet Technology (TOIT)* 17, 3 (June 2017), 27:1–27:22.
- Gaurav Misra and Jose M. Such. 2017. PACMAN: Personal agent for access control in social media. *IEEE Internet Computing* 21, 6 (2017), 18–26.
- Sean Murphy. 2015. Is cybersecurity possible in healthcare? *National Cybersecurity Institute Journal* 1, 3 (March 2015), 49–63.
- Luis G. Nardin, Tina Balke-Visser, Nirav Ajmeri, Anup K. Kalia, Jaime S. Sichman, and Munindar P. Singh. 2016. Classifying sanctions and designing a conceptual sanctioning process model for socio-technical systems. *The Knowledge Engineering Review* 31, 2 (March 2016), 142–166.
- Crowd Research Partners. 2016. Insider Threat Spotlight Report. Retrieved November 2017 from <http://crowdresearchpartners.com/portfolio/insider-threat-report/>.
- Sarah Rajtmajer, Anna Squicciarini, Jose M. Such, Justin Semonsen, and Andrew Belmonte. 2017. An ultimatum game model for the evolution of privacy in jointly managed content. In *International Conference on Decision and Game Theory for Security*. Springer, 112–130.
- Fariba Sadri. 2012. Intention recognition in agents for ambient intelligence: Logic-based approaches. In *Agents and Ambient Intelligence. Ambient Intelligence and Smart Environments*, Vol. 12. IOS Press, 197–236.
- John Sammons. 2015. *Digital Forensics: Threatscape and Best Practices*. Syngress.
- Munindar P. Singh. 2013. Norms as a basis for governing sociotechnical systems. *ACM Transactions on Intelligent Systems and Technology (TIST)* 5, 1 (Dec. 2013), 21:1–21:23.
- Jose M. Such. 2017. Privacy and autonomous systems. In *Proceedings of the 26th International Joint Conference on Artificial Intelligence (IJCAI'17)*. 4761–4767.
- Jose M. Such, Joel Porter, Sören Preibusch, and Adam Joinson. 2017. Photo privacy conflicts in social media: A large-scale empirical study. In *Proceedings of the 2017 CHI Conference on Human Factors in Computing Systems*. ACM, 3821–3832.
- Jose M. Such and Michael Rovatsos. 2016. Privacy policy negotiation in social media. *ACM Transactions on Autonomous and Adaptive Systems (TAAS)* 11, 1 (2016), 4.
- Verizon. 2017. Data Breach Digest. Retrieved November 2017 from <http://www.verizonenterprise.com/verizon-insights-lab/data-breach-digest/2017/>.

Received November 2017; revised November 2017; accepted November 2017