Acknowledgments

We gratefully acknowledge the contributions and assistance of our colleagues in making this book possible, who include but are not limited to Sadia Afroz, Scott Alfeld, Tansu Alpcan, Rekha Bachwani, Marco Barreno, Adam Barth, Peter Bartlett, Battista Biggio, Chris Cai, Fuching Jack Chi, David Fifield, Laurent El Ghaoui, Barbara Goto, Rachel Greenstadt, Yi Han, Ling Huang, Michael Jordan, Alex Kantchelian, Hideaki Kawabata, Marius Kloft, Pavel Laskov, Shing-hon Lau, Chris Leckie, Steven Lee, Justin Ma, Steve Martin, Brad Miller, Satish Rao, Fabio Roli, Udam Saini, Tobias Scheffer, Russell Sears, Anil Sewani, Arunesh Sinha, Dawn Song, Nedim Šrndić, Charles Sutton, Nina Taft, Anthony Tran, Michael Tschantz, Kai Xai, Takumi Yamamoto, and Qi Zhong. We additionally thank Matthias Bussas and Marius Kloft for their careful proofreading of Chapter 4 and the staff at Cambridge University Press including Heather Brolly and Julie Lancashire for their help in preparing this manuscript. We would also like to thank the many colleagues with whom we have had fruitful discussions at the Dagstuhl Perspectives Workshop on Machine Learning Methods for Computer Security (Joseph, Laskov, Roli, Tygar, & Nelson 2013) and at the ACM Workshop on Artificial Intelligence and Security (AISec) and other workshops and conferences.

The authors are currently at the University of California, Berkeley, the University of Melbourne, and Google. We thank these institutions. While we were writing this book, some of the authors were at Universität Tübingen, Universität Potsdam, Università di Cagliari, IBM Research, and Microsoft Research, and we also thank those institutions. We offer special thanks to our support staff, including Angie Abbatecola, Kattt Atchley, Carlyn Chinen, Barbara Goto, Damon Hinson, Michaela Iglesia, Shane Knapp, Jey Kottalam, Jon Kuroda, Lena Lau-Stewart, Christian Legg, and Boban Zarkovich.

We are grateful for the financial sponsors of this research. We received U.S. government funding from the Air Force Office of Scientific Research, Homeland Security Advanced Research Projects Agency, National Science Foundation, and State Department DRL, and in some cases through UC Berkeley laboratories (DETERlab and TRUST). Some authors received additional support from the Alexander von Humboldt Foundation, the Australian Research Council (DE160100584), the Center for Long-Term Cybersecurity, the Future of Life Institute, Oak Ridge National Laboratory, and the Open Technology Fund. The opinions expressed in this book are solely those of the authors and do not necessarily reflect the views of any funder.

The authors could not have written this book without the support, encouragement, and patience of their friends and families.