

Dr. Blaine Nelson

Background

- 1995-1999: Student at Franklin County High School, Frankfort, KY
- 1999-2003: Undergraduate studies (Computer Science) at the University of South Carolina, Columbia
- 2003-2010: Research assistant (Computer Science) at the University of California, Berkeley
- Since March 2011: Post-doc at the Center for Bioinformatics (ZBIT), University of Tübingen

Research Interests

- Computer Security
- Machine Learning
- Secure/Robust Machine Learning

Recent Activities

- Co-organizer: [AISec 2012](#) Workshop
- Coordinator: [2012 Perspectives Workshop - Machine Learning Methods for Computer Security \(Schloss Dagstuhl, Germany\)](#)
- Member: [AISec 2011](#), [PSDML 2010](#) Workshops

Current Projects

- Investigating robustness of classifiers against adversarial noise.
- Designing new classifier techniques to be more robust.

Publications

1. Ling Huang, Anthony D. Joseph, Blaine Nelson, Benjamin I. P. Rubinstein, and J. D. Tygar. **Adversarial Machine Learning**, In *Proceedings of the 4th ACM Workshop on Artificial Intelligence and Security*, ACM, 21 October 2011.
2. Blaine Nelson, Benjamin I. P. Rubinstein, Ling Huang, Anthony D. Joseph, and J. D. Tygar. **Classifier Evasion: Models and Open Problems**, In *Privacy and Security Issues in Data Mining and Machine Learning*, volume 6549 of Lecture Notes in Computer Science, 2011, pages 92-98.
3. Marco Barreno, Blaine Nelson, Anthony D. Joseph, and J. D. Tygar. **The Security of Machine Learning**, In *Machine Learning Journal*, 81(2), 2010, pp. 121-148. (Technical Report: EECS Department, University of California, Berkeley, Tech. Rep. UCB/EECS-2008-43, 2008.)
4. Blaine Nelson, Benjamin I. P. Rubinstein, Ling Huang, Anthony D. Joseph, Shing-hon Lau, Steven Lee, Satish Rao, Anthony Tran and J. D. Tygar, **Near-Optimal Evasion of Convex-Inducing Classifiers**, In the *Proceedings of the Thirteenth International Conference on Artificial Intelligence and Statistics (AISTATS 2010)*, 2010.
5. Benjamin I. P. Rubinstein, Blaine Nelson, Ling Huang, Anthony D. Joseph, Shing-hon Lau, Satish Rao, Nina Taft, and J. D. Tygar, **ANTIDOTE: Understanding and Defending against Poisoning of Anomaly Detectors**, In the *Proceedings of the Internet Measurement Conference (IMC 2009)*, 2009.
6. Benjamin I. P. Rubinstein, Blaine Nelson, Ling Huang, Anthony D. Joseph, Shing-hon Lau, Satish Rao, Nina Taft, and J. D. Tygar, **Stealthy Poisoning Attacks on PCA-based Anomaly Detectors**, In *ACM SIGMETRICS Performance Evaluation Review*, 2009.
7. Blaine Nelson, Marco Barreno, Fuching Jack Chi, Anthony D. Joseph, Benjamin I. P.

- Rubinstein, Udam Saini, Charles Sutton, J. D. Tygar, and Kai Xia. **Misleading learners: Co-opting your spam filter**, Book chapter in Jeffrey J. P. Tsai and Philip S. Yu (eds.) *Machine Learning in Cyber Trust: Security, Privacy, and Reliability*, pg. 17-51, 2009. [[pdf](#)]
8. Marco Barreno, Peter L. Bartlett, Fuching Jack Chi, Anthony D. Joseph, Blaine Nelson, Benjamin I. P. Rubinstein, Udam Saini, and J. D. Tygar, **Open Problems in the Security of Learning**, *In the Proceedings of the First ACM Workshop on AISEc*, pg. 19-26, 2008.
 9. Benjamin I. P. Rubinstein, Blaine Nelson, Ling Huang, Anthony D. Joseph, Shing-hon Lau, Nina Taft, and J. D. Tygar, **Evading Anomaly Detection through Variance Injection Attacks on PCA (Extended Abstract)**, *In the 11th International Symposium on Recent Advances in Intrusion Detection (RAID 2008)*, pg. 394-395, 2008. **Winner of the RAID08 Best Poster Award.**
 10. Benjamin I. P. Rubinstein, Blaine Nelson, Ling Huang, Anthony D. Joseph, Shing-hon Lau, Nina Taft and J. D. Tygar, **Compromising PCA-based Anomaly Detectors for Network-Wide Traffic**, EECS Department, University of California, Berkeley, Tech. Rep. UCB/EECS-2008-73, 2008.
 11. Blaine Nelson, Marco Barreno, Fuching Jack Chi, Anthony D. Joseph, Benjamin I.P. Rubinstein, Udam Saini, Charles Sutton, J. D. Tygar, and Kai Xia, **Exploiting Machine Learning to Subvert Your Spam Filter**, *In the Proceedings of the First USENIX Workshop on Large-Scale Exploits and Emergent Threats (LEET'08)*, San Francisco, CA, April 15, 2008.
 12. Blaine Nelson, and Anthony D. Joseph, **Bounding an Attack's Complexity for a Simple Learning Model**, *In the Proceedings of the First Workshop on Tackling Computer Systems Problems with Machine Learning Techniques (SysML)*, Saint-Malo, France, June, 2006.
 13. Marco Barreno, Blaine Nelson, Russell Sears, Anthony D. Joseph, and J. D. Tygar, **Can Machine Learning Be Secure? (Invited paper)**, *In the Proceedings of the ACM Symposium on InformAtion, Computer and Communications Security (ASIACCS'06)*, Taipei, Taiwan, March, 2006.

Theses

1. Blaine Nelson, **Behavior of Machine Learning Algorithms in Adversarial Environments** (PhD dissertation). University of California, Berkeley, Department of EECS technical report UCB/EECS-2010-140. November 23 2010.
2. Blaine Nelson, **Designing, Implementing, and Analyzing a System for Virus Detection** (Master's dissertation) University of California, Berkeley, Department of EECS technical report UCB/EECS-2006-27, March 19 2006.

Address, Phone, Fax, Email

Eberhard-Karls-Universität Tübingen
Wilhelm-Schickard-Institut für Informatik
Lehrstuhl Kognitive Systeme
Sand 1
D - 72076 Tübingen

Germany

Email: [blaine<DOT>nelson<AT>gmail<DOT>com](mailto:blaine@nelson@gmail.com)