Signal Processing for Adversarial Machine Learning

General Chairs:

Pin-Yu Chen, *IBM Research AI*
Sijia Liu, *IBM Research AI*
Bo Li, *UIUC*

Technical Program Chairs:

Jinfeng Yi, *JD.com*
Cho-Jui Hsieh, *UC Davis*

Recent studies have highlighted the lack of robustness in state-of-the-art machine learning models. For instance, carefully crafted adversarial perturbations to natural images can easily cause modern classifiers trained by deep convolutional neural networks to yield incorrect predictions, while these adversarial examples can be made visually similar to the natural images, resulting in critical safety and security concerns of services and applications supported by machine learning models.

Signal processing and black-box optimization techniques, such as manifold analysis, data transformation and zeroth-order optimization, are becoming the core components in the research of adversarial machine learning. They are widely used to generate powerful adversarial examples to deceive target machine learning models and evade detection, as well as to provide robust and effective machinery against adversarial examples. This symposium aims to bring together researchers and practitioners from both academia and industry to report novel advances and to publish high-quality papers, in order to foster the field of signal processing for adversarial machine learning.

Topics include but are not limited to:

- Methods of generating adversarial examples to deceive state-of-the-art machine learning models
- Methods of mitigating adversarial perturbations towards robust machine learning
- Demonstrations of adversarial examples attacking real-world machine learning services in different domains, including but are not limited to images, videos, speeches and texts
- Explanation and interpretation towards the origins and transferability of adversarial examples
- Geometric analysis (e.g., manifold and subspace) of data distribution for adversarial machine learning
- Data transformation techniques for adversarial machine learning
- Zeroth-order (gradient-free) optimization techniques applied to adversarial machine learning or other machine learning aspects
- Security and privacy implications including but are not limited to data poisoning, model stealing and data privacy

Paper Submission: Prospective authors are invited to submit full-length papers (up to 4 pages for technical content including figures and possible references, and with one additional optional 5th page containing only references) and extended abstracts (up to 2 pages, for paper-less industry presentations and Ongoing Work presentations) via the GlobalSIP 2018 conference website. Manuscripts should be original (not submitted/published anywhere else) and written in accordance with the standard IEEE double-column paper template. The accepted abstracts will not be indexed in IEEE Xplore, however the abstracts and/or the presentations will be included in the IEEE SPS SigPort. Accepted papers and abstracts will be scheduled in lecture and poster sessions.

Important Dates:

- **June 17, 2018**: Paper submission due
- **Aug. 7, 2018**: Notification of Acceptance
- **Aug. 22, 2018**: Camera-ready paper due.

For inquiries please contact: Pin-Yu Chen (pin-yu.chen@ibm.com), Sijia Liu (sijia.liu@ibm.com), Bo Li (lxbosky@gmail.com), Jinfeng Yi (jinfengyi.ustc@gmail.com), Cho-Jui Hsieh (chohsieh@ucdavis.edu)