ISTI Rapid Response Data Challenge: Resilient Machine Learning Algorithms

Synopsis

Verification and validation techniques for data-driven models are an area of increasing importance. Academic researchers are producing a widening array of attacks against state-of-the-art machine learning (ML) models. Government agencies and civilians are relying heavily on these ML platforms. Therefore to make data analytics safe for real world deployment, we need to ensure that algorithms are understood by analysts, adaptable to changing situations, robust in new domains, and reliable in the presence of an adversary. In order to familiarize the workforce with resilient ML techniques the Information Science & Technology Institute (ISTI) solicits teams to participate in a data challenge.

Format

Teams of data scientists will compete to generate adversarial examples and robust classifiers on the provided datasets. Teams will train and submit a classifier on the provided training datasets. Classifiers will be ranked on team-submitted adversarial example datasets (for robustness). Teams will also compete to generate the most successful adversarial example (i.e. the example that tricks the most classifiers) and the most "human readable" adversarial example (i.e. the example that tricks the most human judges). The leaderboard for each category (best adversarial example, most robust classifier, and most human readable example) will be posted throughout the contest and the final winners for each category will be announced at the final presentation day.

Deliverables (required for teams receiving funding):

All teams will be asked to present their work to a general lab-wide audience at the end of the challenge. In addition, all teams will contribute to a written report to be shared openly.

Datasets (publicly available)

1.) RadioML RF modulation classification dataset (2016.10A).

https://www.deepsig.io/datasets

2.) Real and fake images of celebrities' faces (CelebA).
https://github.com/tkarras/progressive_growing_of_gans

Register your team by submitting the team name and members to mturcotte@lanl.gov with subject line "ISTI Adversarial ML Data Challenge". In addition, if any members of your team require funding please submit .25 - 1.0 page, 11pt font document as an attachment that includes the following:

- ● Brief description of proposed approach to the challenge;
- ● Brief bio of team members with qualifications and how participating will benefit you;
- ● Requested amount in FTE hours per individual requesting funding (do not exceed more than 200

  hours per team).

Deadlines

1.) Team registration and requests for funding due - July 08th
2.) Data Release + Kick-off presentation + Funding decisions - July 10th
3.) Adversarial examples for the human readable category - September 10th 4.) Final presentations + Winners announced - September 24th

For questions regarding this call, contact, email rr-istiml@lanl.gov.
Organizers: Juston Moore, Ben Migliori, Diane Oyen, Garrett Kenyon, Laura Monroe, Nga Nguyen