

## **Inaugural Workshop on Boosting Resilience through Artificial Intelligence and Decision Support (BRAIDS)**

**Background.** Stakeholders of all types are interested in reliably characterizing the computer networks relevant to their mission space. For instance, planners and architects want to know if the network has sufficient capacity to handle expected traffic surges, or the resilience to recover from unforeseen events. Defenders want to establish a robust baseline for network behavior, in order to identify unexpected events or unusual behavior. Recently, significant attention has been directed at developing Artificial Intelligence entities (AIs) for decision support to enhance the security and resilience of computer networks. The raw data enabling these efforts are frequently time series; time series which are used predictively, rather than forensically.<sup>1</sup>

**The Problem.** Data scientists studying network behavior sometimes pull common statistical tools “out of the box” to use in predictive analysis. However, many available signal processing, machine learning, and AI libraries implicitly assume certain statistical properties of the stochastic observables. When applied to data that violate the implicit assumptions, the tools will perform poorly or generate high false alarm rates. Unfortunately, network time series are often non-Gaussian, non-stationary, and non-ergodic, making the application of many common tools inappropriate without first transforming the data. For instance, an ML classifier with complex non-stationary input data could become increasingly inappropriate as its input evolves away from the conditions that prevailed at training. It is hardly surprising that analyses built on inappropriate foundations fail to hold up.

Substantial work has been done on non-stationary and non-ergodic time series in other fields, notably finance, hydrology, and geophysics, where rich rewards have been gained by quantitative analysis. However, this extensive armory of econometric tools has not yet been deployed in the cyber domain.

**Our Proposal.** We propose a systematic empirical study of cyber-network observables to

- Characterize statistical properties of the observables.
- Identify cyber variables (or transformations thereof) that exhibit convenient statistical properties.
- Develop analytics to exploit emergent behavior, especially for trend projection and anomaly detection.
- Determine how to use evolution of statistical properties to indicate that AI retraining may be required.
- Investigate whether changes in the statistical properties of time series can expose adversarial learning.

This research will enable robust predictive cyber analytics for effective decision support. As a first step, we announce a two-day Workshop, with the purpose of assembling a community of researchers interested in pursuing these goals and/or with expertise in potentially useful tools and techniques from other fields. Day 1 of the Workshop will consist of presentations of current research results in cyber analytics as well as potentially useful techniques from other fields. Day 2 will be devoted to forming subgroups of researchers interested in particular lines of research. These will be charged with developing a research program for the upcoming year.

The Workshop will be held as a Special Event within the 2019 IEEE High Performance Extreme Computing Conference (HPEC '19), held 24 - 26 September 2019 at the Westin Hotel, Waltham, MA USA. Accepted papers will appear in the Conference Proceedings.

DISTRIBUTION STATEMENT A. Approved for public release. Distribution is unlimited. This material is based upon work supported by the United States Air Force under Air Force Contract No. FA8702-15-D-0001. Any opinions, findings, conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of the United States Air Force.

---

<sup>1</sup> “Left of boom” rather than “right of boom,” in the colorful phrase of the IED world.