Prior Cyber Research Summer School Projects

2019-2022

LA-UR-22-28918

<cyber.toaster>

https://cyberfire.training/school/

SSNzkSNARK: Secure Neural Network Verification System Using zkSNARKs (2019)

Student: Zachary DeStefano
Mentor(s): Michael Dixon

- zk-SNARKs are computationally efficient zero-knowledge proof systems that can verify that a computation was performed correctly.
- PCD (Proof Carrying Data) is a construction that allows for a proof to be constructed that attests to the entire history (sequential and parallel) of the execution of an arithmetic circuit multiple times recursively over its previous outputs.
- This project involves constructing efficient PCD zk-SNARKs for verifiable neural network execution.
- We utilize training using recursive proof composition and additional lower level optimizations.
- Potential applications of this research includes nuclear treaty verification, data integrity, and supply chain security.

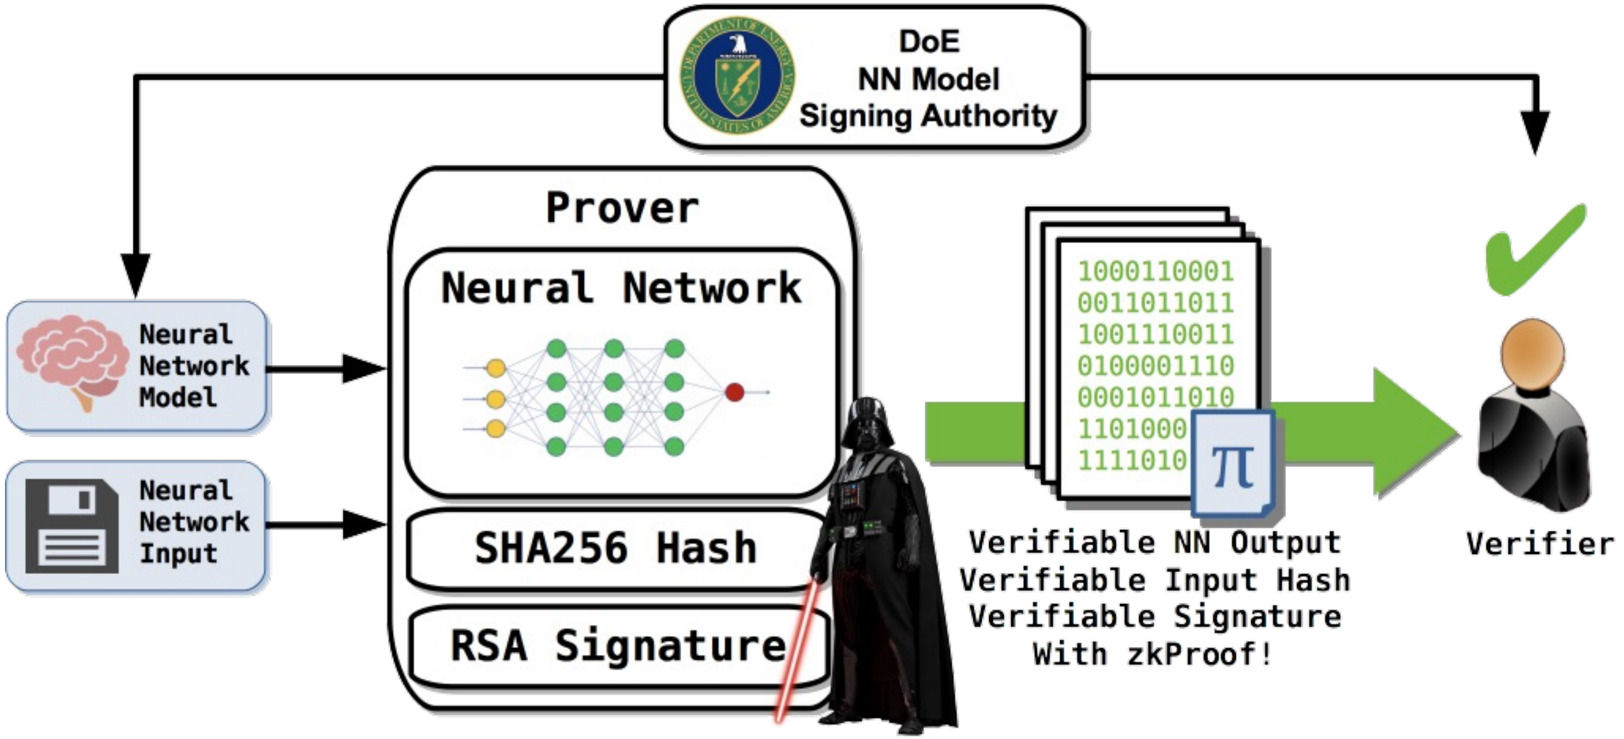

Anomalous Event Detection using Non-negative Poisson Tensor Factorization (2020)

Student: Maksim E. Eren
Mentor(s): Juston S. Moore, Boian S. Alexandrov

- Detecting malicious anomalous network activities and distinguishing them from unusual but benign events is a fundamental challenge for cyber defenders.
- Non-negative tensor factorization, a powerful unsupervised machine learning method, that can naturally model multi-dimensional data to capture the complex and multi-faceted details of behavior profiles.
- Our method generalize to unseen types of attacks by detecting deviations from normal behavior, without knowledge of specific attack signatures.
- Our new unsupervised statistical anomaly detection methodology matches or surpasses state-of-the-art supervised learning baselines across several challenging and diverse cyber application areas with extreme class imbalance, including detection of compromised user credentials, botnets, spam e-mails, and fraudulent credit card transactions.
- We provide a publicly available Python library, named pyCP_APR.
- pyCP_APR is part of the SmartTensors software package that won the R&D 100 2021 award.

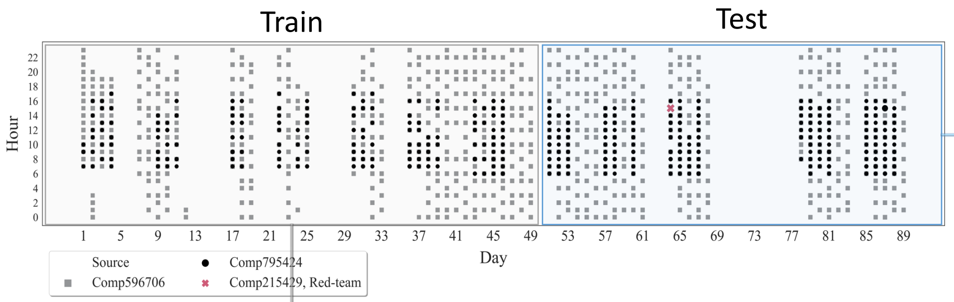

1 Model the normal behavior

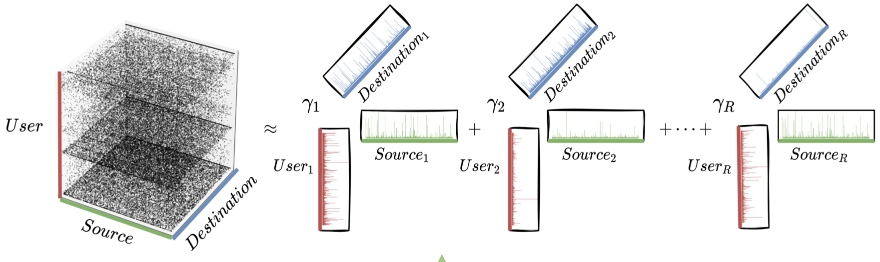

2 Get the tensor coordinates

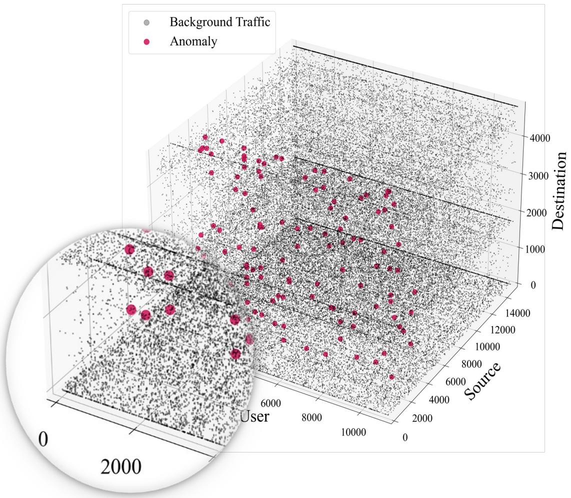

3 Obtain the anomaly scores

$$\mathfrak{X}_{i_1,i_2,i_3} \sim \text{Poisson}(\lambda_{i_1,i_2,i_3})$$

$$\lambda_{i_1,i_2,i_3} = \sum_{r=1}^{R} \gamma_r \prod_{d=1}^{3} \theta_{r,i_d}^{(d)}$$

Secure System Composition and Type Checking using Cryptographic Proofs (2021)

Student: Dani Barrack
Mentor(s): Michael Dixon, Boris Gelfand

- Formally verifying the correctness of systems of systems involves verification of their compatibility.
- Conventional approaches require the exhaustive checking of an entire system's state space and the undesirable exposure of data.
- We overcome this limitation by including zero-knowledge proofs (ZKPs) with system outputs that assert desired system properties without revealing sensitive information.
- This approach allows us to ensure system integrity without checking every computational path, extends our trusted computing base well beyond our own system, and grants us fine-grained control over which bits of information to keep secret.

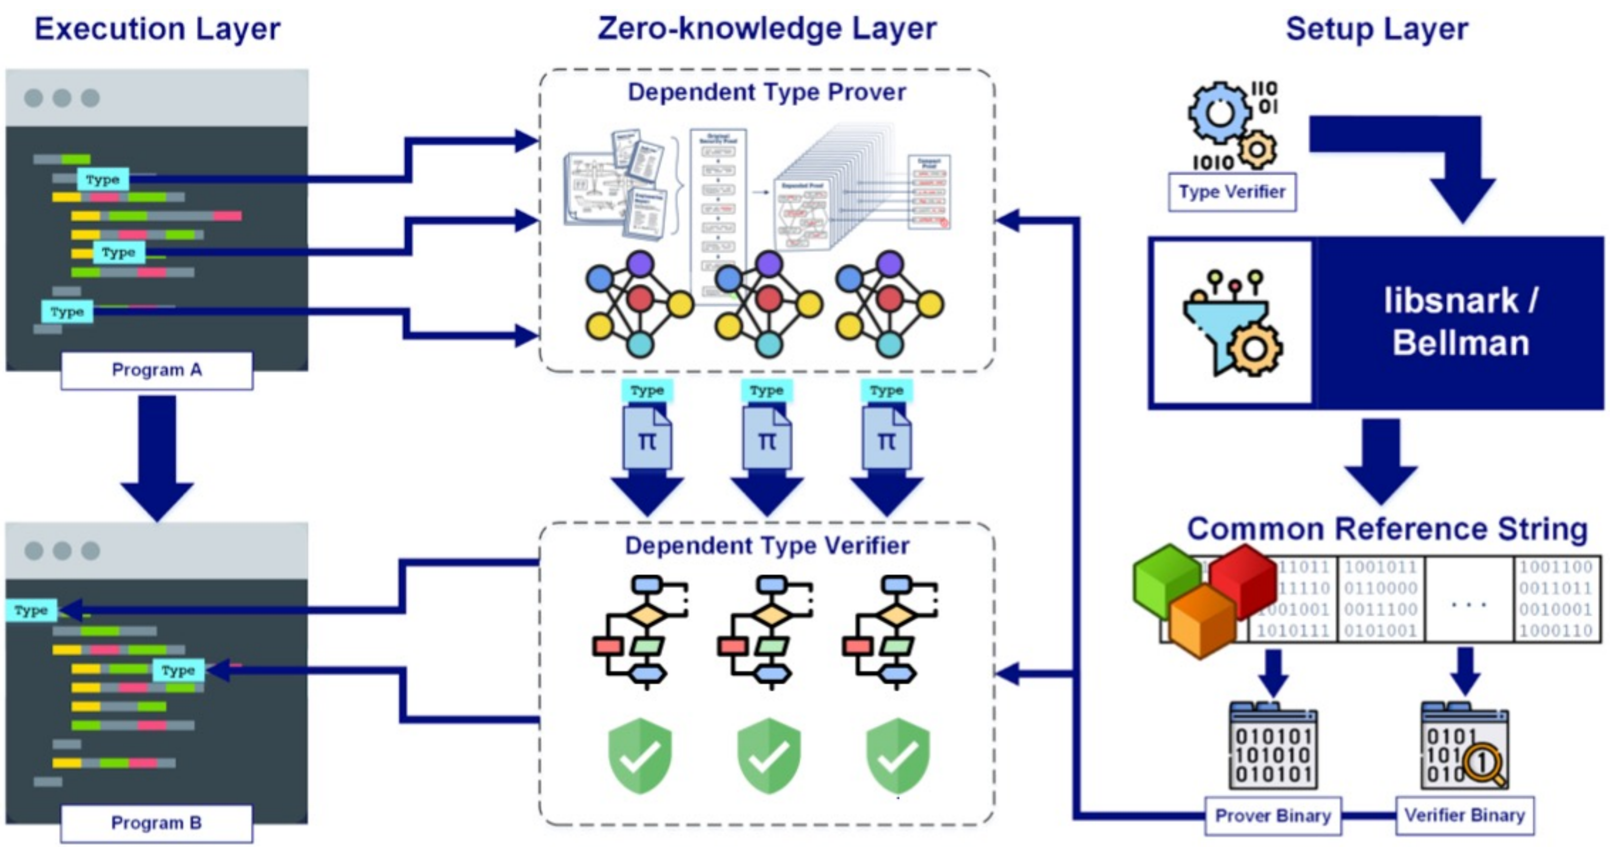

Detecting Electrical Anomalies via Overlapping Measurements (2021)

Similarity Timeseries → **Anomaly Detection** → **Compare and Analyze**

- **Similarity Timeseries**
  - Dynamic Time Warping
- **Anomaly Detection**
  - Autoregression, Level Shift, Rolling Average
- **Compare and Analyze**
  - Individual vs Merged Data

Student: Sina Sontowski
Mentor(s): Nigel Lawrence, Deepjyoti Deka

- As cyber-attacks against critical infrastructure become more frequent, it is increasingly important to be able to rapidly identify and respond to these threats.
- We are investigating methods for using multiple independent systems with overlapping electrical measurements to rapidly identify anomalies.
- Prior research has explored the benefits of fusing measurements.
- Overlapping measurements from an existing electrical system has not been investigated.
- We explore the potential benefits of combining overlapping measurements both to improve the speed/accuracy of anomaly detection and to provide additional validation of collected measurements.

Zero Knowledge Proofs of Certified Robustness (2022)

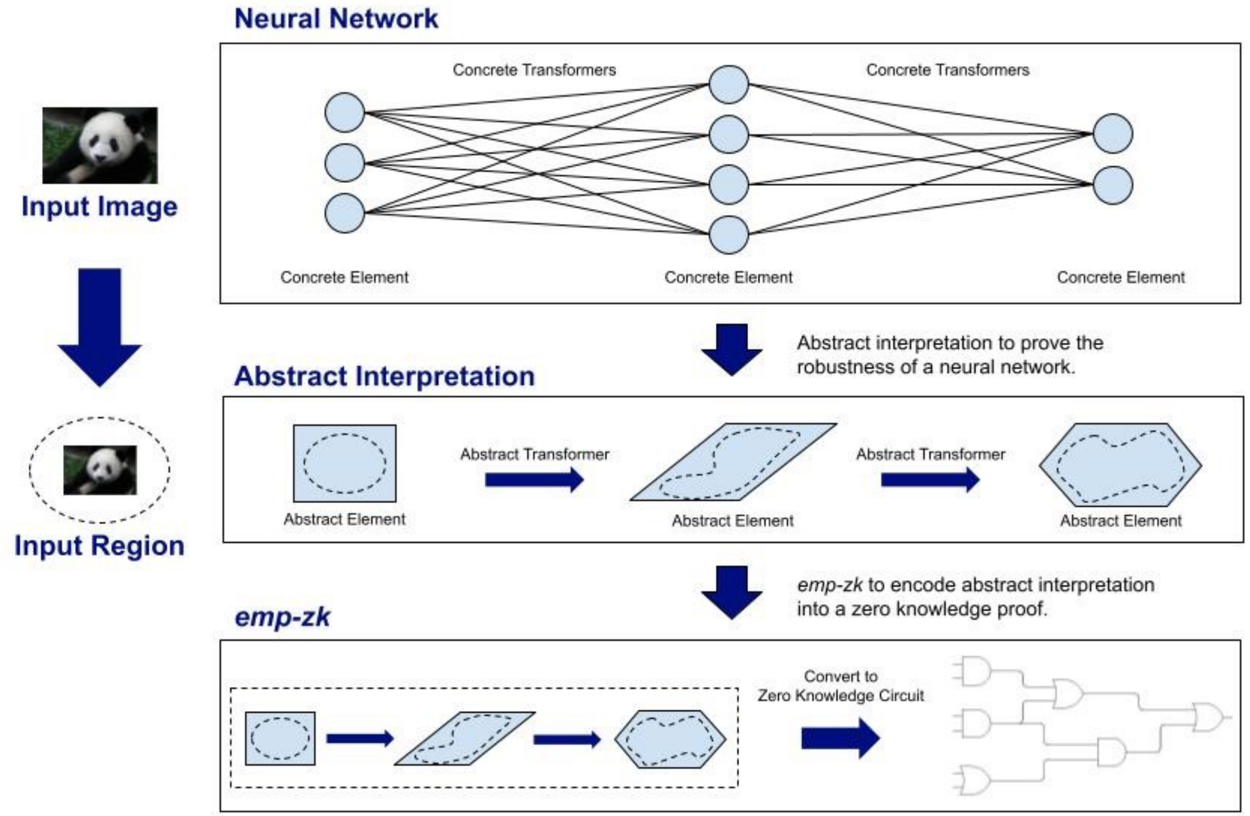

Student: Jack Cheng
Mentor(s): Michael Dixon, Zachary DeStefano

- The threat of adversarial examples that cause neural networks to misclassify inputs has led to the need to certify that models are robust against these adversarial attacks.
- Prior research efforts have developed means of successfully proving model robustness; however, the processes require intimate knowledge and handling of the model or access to the training data itself. i.e proving model robustness to untrusted third parties would require:
    - Sharing sensitive information contained in model weights
    - Network architecture
    - Input training data
- As neural networks become more widely used for mission critical applications involving sensitive data, the ability to attest to model robustness to a third party without revealing sensitive information will become paramount.
- We demonstrate a method of solving this problem by using abstract interpretation to certify robustness and capturing a proof of that fact in a zero-knowledge proof.
- We turn to an approach to prove robustness using abstract interpretation by:
    - Over-approximating the robustness region
    - Running the over-approximation through the neural network abstractly
    - Proving that all points within the resulting over-approximation are classified the same

Cryptographic Structure of Physical Unclonable Functions (2022)

Student: Apollo Albright
Mentor(s): Boris Gelfand, Michael Dixon

- Authentication of systems is an essential feature of secure communication between parties.
- Traditional authentication systems with an ID stored in non-volatile memory are susceptible to being spoofed by copying the ID to a malicious machine.
- One solution to this problem is to store the ID using a physical unclonable function (PUF).
- We prove that a class of linear optical PUFs can be learned to arbitrary precision with arbitrarily high probability, even in the presence of noise, given access to polynomially many challenge-response pairs and polynomially bounded computational power, under mild assumptions about the distributions of the noise and challenge vectors.
- We use a matrix Chernoff bound to formulate polynomial bounds for the required number of samples and the computational complexity of a linear regression algorithm, based on size parameters of the PUF, the distributions of the challenge and noise vectors, and the probability and accuracy of the regression algorithm.