



Workshop Overview and Generative Machine Learning for Cybersecurity



Assefaw Gebremedhin and James Halvorsen
School of Electrical Engineering and Computer Science
Washington State University
VICEROY CySER Summer Workshop
May 22, 2023



Two Parts



Overview of CySER and
Summer Workshop 2023

Assefaw
Gebremedhin



Generative Machine
Learning for
Cybersecurity

James
Halvorsen

What is CySER?

- An Institute funded by the Department of Defense Air Force Command through the VICEROY program
 - VICEROY = Virtual Institutes for Cyber and Electromagnetic Spectrum Research and Employ
 - VICEROY Institutes are managed by the Griffiss Institute
- Directly responds to the VICEROY call
 - Training ROTC and DoD-aligned civilians in cybersecurity at the undergraduate and graduate level, with primary emphasis on undergraduate
- Builds a strong consortium in the Pacific Northwest for cybersecurity education and research
 - CySER brings together 5 institutions with complementary strengths and diversity of populations served
- Seeks to position WSU to attain Center of Academic Excellence in Cyber Operations (CAE-CO) designation
 - WSU is launching a new BS in Cybersecurity with emphasis on cyber operations starting from Fall 2023

CySER: Institutions and People

Washington State University (WSU)

- Bernard Van Wie (VSCBE; Lead PI)
- Assefaw Gebremedhin (EECS; Co-PI; Research & Curriculum Lead)
- Noel Schulz (EECS; Co-PI; Industry Lead)
- Venera Arnaoudova (EECS; CS Curriculum POC)
- Olusola Adesope (Education; Evaluator)
- Partha Pande (EECS; SP)
- Haipeng Cai (EECS; SP)
- Robert Crossler (MISE; SP)
- Jana Doppa (EECS; SP)
- Arda Gozen (MME; SP)
- Larry Holder (EECS; SP)
- Chris Hundhausen (former EECS; SP)
- John Miller (EECS; SP)
- Gabriel Nketah (Project Coordinator)
- James Crabb (Project Coordinator)



WSU/UI ROTC

- LTC Nicholas Jeffers (former)
- Major Paul Hyde (former)
- Major John Ford (current)

Montana State University (MSU)

- Clemente Izurieta (MSU Site Lead)
- LTC Lance Ratterman (former)
- LTC Christopher L'Heureux (current)



University of Idaho (UI)

- James Alves-Foss (UI Site Lead)
- Terence Soule



Columbia Basin College (CBC)

- Mathew Boehnke (former CBC Site Lead)
- T. Lee Williams (current CBC Site Lead)

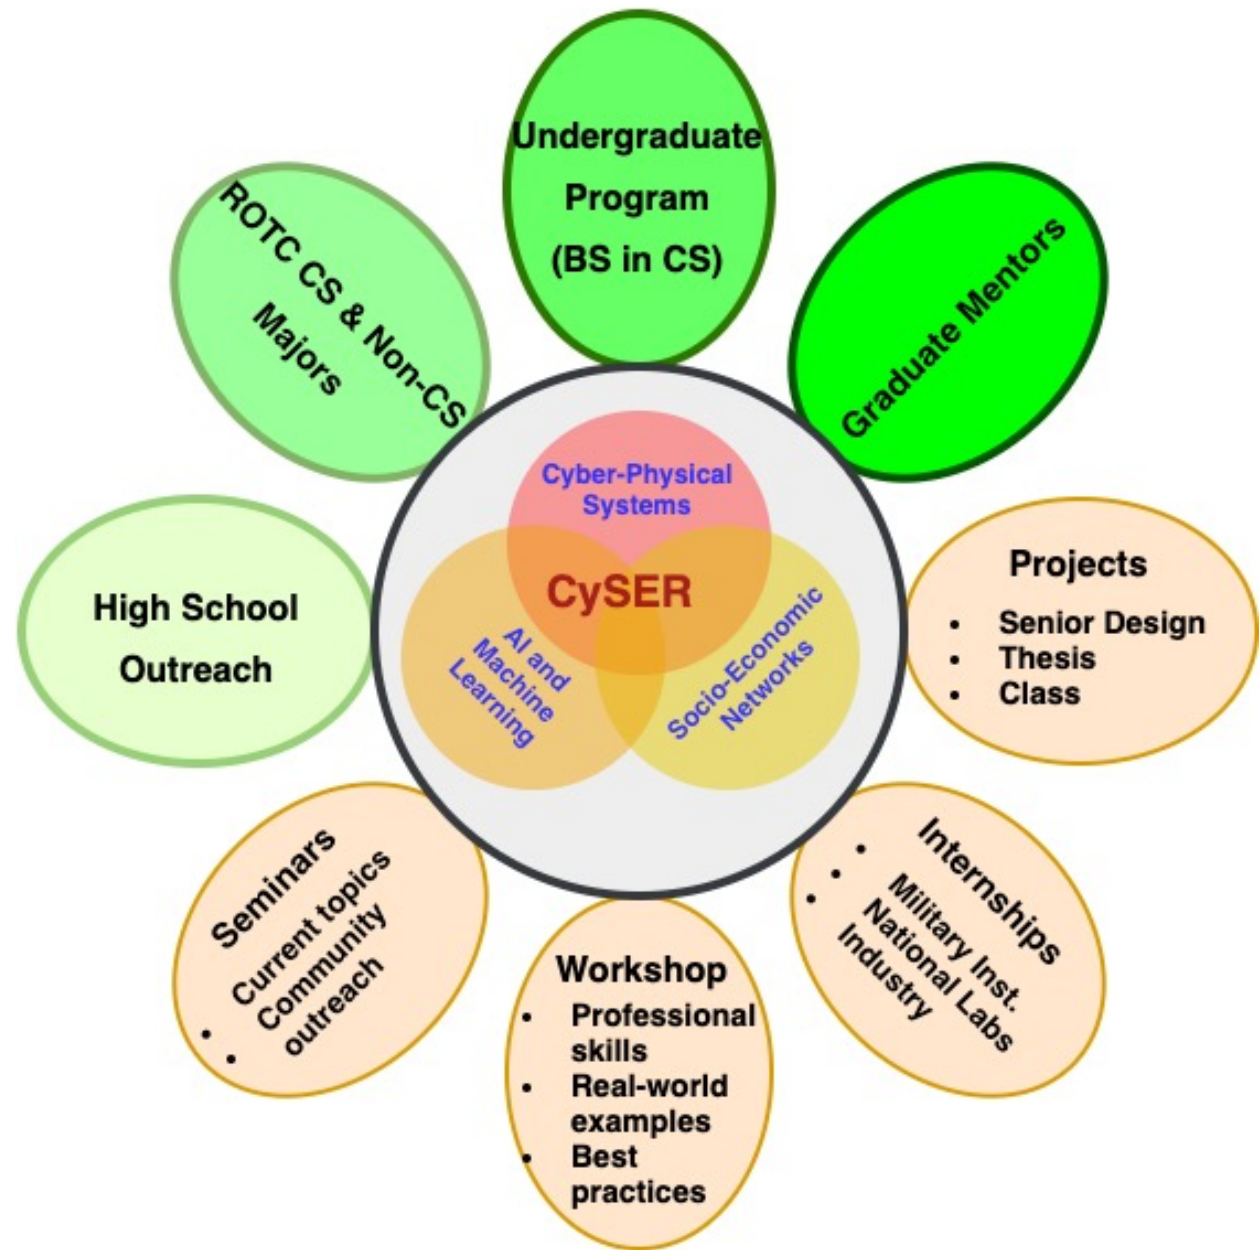


Central Washington University

- Capt. JonCharles Tenbusch (current CWU Site Lead)
- Andrew Van Den Hoek (former CWU Site Lead)
- David Douglas (ITAM)



CySER Program Elements



CySER Curriculum: Encapsulated via 3 Certificate Offerings at WSU



CySER CAE-CO Fundamentals

BS in Computer Science, Software Engineering

Lead by EECS

Approved by Faculty Senate in Fall 2022



CySER Basics

For non-CS majors

Affiliated with the MISE program in the College of Business

Approved by Faculty Senate in Fall 2022



CySER CAE-CO Advanced

MS/PhD students in CS, CE, EE, ME, ChE, MISE or similar field

Mentor CySER undergraduates on research projects

Lead by EECS

Approved by Faculty Senate in Spring 2023

2023 Summer Workshop Overview



Presentations and
Tutorials

Todd Hall 203



Poster Session

Tue May 23, 1pm – 2:15pm
Spark G10



Certificate Ceremony &
Celebration

Tue May 23, 2:30pm – 3:15pm
Spark G10



Team Building and
Leadership Session

Wed May 24, 10:30am – 12pm
Spark G10

Maj. Paul Hyde, WSU; LTC Matthew
Sheftic, WSU; Capt. JonCharles
Tenbusch, CWU



Field Trips

PNNL (May 25)
SEL (May 26)
Fairchild AFB (June 1)

Presentations: Technical

Machine learning and AI

- Generative ML for security – James Halvorsen, WSU (this session)
- Anomaly detection – Jana Doppa, WSU (Mon May 22, 1—2pm)

Networks and Information Security

- Legacy systems management – Julia Stachofsky, WSU (Mon May 22, 10:45—11:45am)
- Behavioral threats – Rob Crossler, WSU (Tue May 23, 8:30—9:30am)
- Risk assessment – Rob Crossler, WSU (Tue May 23, 9:30—10:30am)

Cyber-Physical Systems

- Power systems and cybersecurity – Noel Schulz, WSU (Tue May 30, 8:30—9am)
- Cybersecurity hot topics and ICS – Tim Schulz, Scythe (Tue May 30, 9—10:30am)
- Cybersecurity in ICS – Nathan Kipp, SEL (Tue May 30, 10:45am—12pm)
- Simulating cyber-attacks to biological systems -- Brenden Fraser Hevlin et al, WSU (Wed May 31, 1:30—2:30pm)

Software Security and Quality Assurance

- Hardware and malware detection and recovery using FPGAs – Chris Major, MSU (Fri May 26, 2:05—3pm)
- Securing the supply chain in a company -- Slater Weinstock, Casaba Security (Fri May 26, 3:15—4:15pm)
- Sustainability defenses against evolving mobile malware – Haipeng Cai, WSU (Wed May 30, 2:30—3:30pm)

Tutorials

User interface
for
interpretation of
code
vulnerabilities

John Miller, WSU-Tricities

Mon May 22, 2:15—4pm

Deep fakes and
ChatGPT

Deborah Wells, CWU

Wed May 24, 1—2:50pm

Graph mining for
insider threat
detection

Larry Holder, WSU

Wed May 24, 3:05—5pm

Digital forensics

Andrew Fallin, MSU

Tue May 30, 1—5pm

Linux tools for
binary reverse
engineering

Jim Alves-Foss, UI

Wed May 31, 8:30—10:30am

Presentations: Collaborations and Opportunities

Cyber overview and
opportunities at
AFRL

Sonja Glumich, AFRL/RI
(Tue May 23, 10:45—11:45am)

DoD collaborations
and opportunities
at Keyport NUWC

Aaron Darnton, NUWC
(Tue May 23, 3:30—4:30pm)

Internship
opportunities at the
Griffiss Institute

Jennifer McCullough, GI
(Wed May 24, 9:30—10:15am)

Vehicle
cybersecurity

Rita Barrios, UDM VICEROY
(Fri May 26, 1—2pm)

Army Cyber
Command

Matt Boehnke, CBC
(Sat May 27, 8:30—10am)

National Cyber
League

Matt Boehnke, CBC
(Sat May 27, 10:15—11:15am)

Cyber threat
intelligence (table
discussion)

Alexander Salazar and Daniel Brown, CISA
(Wed May 31, 10:45am—12:30pm)



Presentations: professional skills

- Day-in-the-life of a cybersecurity professional
 - Slater Weinstock, Casaba Security and Chris Hundhausen, OSU
 - Wed May 24, 8:30—9:30am
- Life-long learning: Getting the most out of your internship
 - Sola Adesope, WSU
 - Wed May 31, 2:30--3:30pm

Part II



Overview of CySER and
Summer Workshop 2023

Assefaw
Gebremedhin



Generative Machine
Learning for
Cybersecurity

James
Halvorsen



Machine Learning Review




- An Informal Definition of ML:
 - “A collection of techniques that enable a machine to learn the definition of a function”
- Approaches:
 - Supervised (have labels, learn mapping from feature vector to labels)
 - Unsupervised (no labels, learn ways to group feature vectors)
 - Reinforcement (learn a policy for an intelligent agent to act upon)
- ML Tasks:
 - Classification
 - Regression

Review: Applications of ML to Cyber Security



Intrusion Detection Systems (IDS)

Monitor activity on host system and/or network
Determine if/when intrusion occurs based on monitoring data



Incident Response

Analyze security alerts to determine the cause
Decide on an appropriate response to deal with the threat



Automated Red Teaming

Test the strength of cyber defenses
Treats compromising a system as an AI planning problem.

IDSs in Depth

Implementation details vary

- Host-based (HIDS) vs Network-based (NIDS)
- Signature-based vs Anomaly-based
- Rule-based (expert system) vs Statistical Approach

All approaches require data

- Host-based: process and filesystem data
- Network-based: packet captures, netflow
- Approaches using supervised learning also require labels

Difficulties with Machine Learning in IDS



Data Problems

Availability of labeled data
Need for effective pre-processing



False Positives

Can lead to incorrect responses (denies availability).

May cause alerts to be ignored, IDS usage abandoned.



Zero Day Attacks

Significant problems for anything Signature-Based.

Anomaly-Based approaches still imperfect.

If attacker has classifier model, can create attacks that evade detection.


An abstract digital cityscape rendered in shades of blue. The scene is composed of numerous rectangular blocks and cubes of varying sizes, some of which are hollow or have glowing interiors. The surfaces of these blocks are covered in a dense pattern of binary code (0s and 1s). Several bright, glowing points of light in blue, green, and red are scattered throughout the scene, some appearing to be part of the structure and others floating in the air. The overall effect is a sense of a complex, interconnected digital environment.

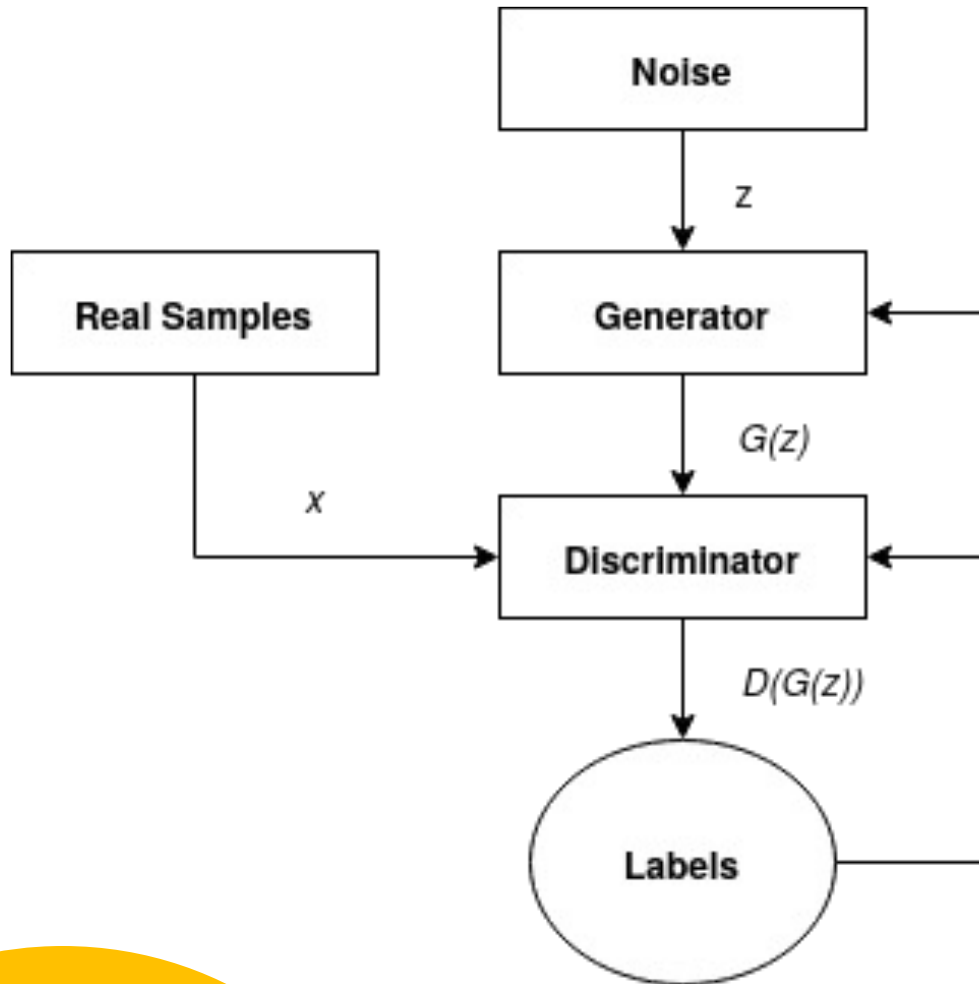
Generative Machine Learning Overview

- Lots of problems related to data in cyber security
- Task: create more/better data
 - Train a generative model on some smaller dataset
 - Learn statistical distribution of dataset
 - Create model that turns random noise into data
 - Essentially classification in reverse (for some label, generate feature vector)

Common Generative ML Algorithms

- Generative Adversarial Networks (GANs)
- Variational Autoencoders (VAEs)
- Diffusion Models
- Autoregressive Models






GANs in Depth


- Trains two ML models in adversarial process
- Generator (G) produces synthetic samples
- Discriminator (D) classifies samples as real or fake
- Performance of D used as loss function for G.



GAN Shortcomings

- 
- The adversarial process used for training GANs can make optimization difficult. Challenges include:
 - **Vanishing Gradient:** Discriminator performs too well; Generator cannot effectively learn from it
 - **Mode Collapse:** Generator learns to produce a few good samples to fool the discriminator, and nothing more
 - **Failure to Converge:** Generator and Discriminator do not reach equilibrium; resulting performance is poor
 - Can mitigate some of these with loss functions and/or regularization (example: Wasserstein GAN)
 - Alternatively, can consider other generative models

VAEs in Depth



$$\text{loss} = ||x - \hat{x}||^2 + \text{KL}[N(\mu_x, \sigma_x), N(0, I)] = ||x - d(z)||^2 + \text{KL}[N(\mu_x, \sigma_x), N(0, I)]$$

A diagram depicting VAE training. Source:

<https://towardsdatascience.com/understanding-variational-autoencoders-vaes-f70510919f73>

Introduced around the same time as GANs (2013)

Architecturally similar to existing work on autoencoders

Unlike traditional autoencoders, adds an extra step to decoding from latent representation

Avoids some of the shortcomings GANs face, particularly mode collapse

Encoder neural network converts data into latent representation

Decoder neural network reconstructs data from latent representation

VAE Shortcomings

Compared to GANs on image generation, known to produce blurrier outputs. Not clear how this shows up for other types of data

Less published research – not as easy to build off prior work

- Using Google Scholar, around 64.7k results for “Variational Autoencoder”
- By comparison, “Generative Adversarial Network” has 285k results.
- Side note: these numbers were 34k and 124k when this slide was used in a 2021 presentation!

Can be more complex to implement than GANs or other autoencoders



a. Samples generated by GAN



b. Samples generated by VAE

A comparison of GAN vs VAE in generating samples from the MNIST dataset. Taken from the paper “Deep Generative Models for Image Generation: A Practical Comparison Between Variational Autoencoders and Generative Adversarial Networks”

Diffusion Models



Relatively new method for image generation



Process involves “denoising”

Progressively add noise to training data
Learn to reverse process



Major successes already in AI art (Stable Diffusion)



Few applications outside of image generation – so far.




Gradual denoising applied to CIFAR-10 image dataset.

From the paper “Denoising Diffusion Probabilistic Models” by Ho et al, 2020

Autoregressive Models

- A feed-forward neural network for generating sequential data
- Uses sequence of previous inputs to determine next output
- Used in the “Generative Pre-trained Transformer” AI models (i.e. ChatGPT)



Animation from Google DeepMind’s WaveNet demonstrating feed-forward generation used in Autoregressive models.

Found at <https://www.georgeho.org/deep-autoregressive-models/>

Applications of Generative ML in Cyber Security

Improving dataset quality
and diversity

Evaluating performance of
cyber defenses

Generative models as IDSs

Improving Datasets with Generative ML



Want to train/test IDS or other ML application on security data



Problem: Data is often unbalanced (i.e. 95% benign, 5% various attacks)



Solution: Create data that is statistically similar to minority classes in data



Limitations: Can't create something from nothing; need *some* initial data to work with



Benefits: Can significantly decrease workload in creating security datasets

Evaluating Performance of Cyber Defenses

- Suppose we have a *good* IDS
 - How does it perform against novel attacks?
 - Can an attacker create minor perturbations to evade it?
- Generative models can create attacks IDSs haven't seen before
- May not need to be as complicated as automated red teaming (planning)
- Possible concern for future: what if generative AI is used for offensive purposes?

Generative ML Models as IDSs

- GANs and VAEs train two neural networks together
 - GAN: Use discriminator as a classifier, improved by presence of novel attack data
 - VAE: Create a generator for benign data. Use decoding loss as an anomaly detector
- Current research suggests generative models perform well as IDSs
- Very limited research on using models other than GANs/VAEs

Other work – TOMATO



- Threat **O**bservability and **M**onitoring **A**ssessment **T**Ool.
- Problem: Many different types of cyber security data
 - Host data – process events, filesystem events, syscalls, etc...
 - Network data – netflow, packet capture
- Depending on what data is captured, attacks may be confused with benign traffic
- Need to measure relevance of data to determine best tools to use for monitoring
- TOMATO provides a metric for this relevance, called “observability”

