

Gen AI & Cybersecurity

Chris Schwartz
Nate Mathews
Matt Wright

Apr. 2024

Funding

- **Project:** “Creating a Course on LLMs in Cybersecurity”
- **Funding:** NCAE-C
- **Source:** National Security / Towson University

Purpose of the Course

- **Gen AI changing cybersecurity in practice**
 - Microsoft's Security CoPilot, Google Cloud Security AI Workbench
- **Project-based course**
 - Hands-on exploration, application, & critique of Gen AI
 - Real-world case studies from prior class work, internships
 - Using non-cloud models for proprietary information

Course specs

- **Prerequisites**

- CSEC-559: CSEC-380 (*Web Security*) or Permission of Instructor
- CSEC-659: CSEC-742 or 380 or Permission of Instructor

- **Use of AI in class policy**

- Use Gen AI tools extensively
- Avoid 'cheating' by not utilizing enough Gen AI
- Exception: personal observations and reflections
- Encourage Gen AI use to enhance submissions

Learning Outcomes

1. **Use Gen AI for cybersecurity tasks**
2. **Use some additional tech: API calls, RAG, fine-tuning**
3. **Critically evaluate #1**
4. **Some understanding of ethical and legal issues**

Course Schedule

- Week 1: Jan. 16 & Jan. 18: Intro to GenAI and Cybersecurity 
 - News Presentations Assigned
- Week 2: Jan. 23 & Jan. 25: Prompt Engineering Basics 
 - Project 1 Assigned
- Week 3: Jan. 30 & Feb. 1: Ethics and Law 1 
- Week 4: Feb. 6 & Feb. 8: Present Project 1 
- Week 5: Feb. 13 & Feb. 15: Code and tools 
 - Project 2 Assigned 
- Week 6: Feb. 20 & Feb. 22: Adam Shostack & LLM Basics 
- Week 7: Feb. 27 & Feb. 29: Retrieval-Augmented Generation 
 - Project 3 Assigned
- Week 8: March 5 & March 7: Present Project 2 
- SPRING BREAK: March 12 & March 14
- Week 9: March 19 & March 21: Virginia Dignum & Fine-tuning 
- Week 10: March 26 & March 28: Discuss and Develop Project 3 Ideas
- Week 11: April 2 & April 4: Ethics and Law 2 
- Week 12: April 9 & April 11: Jason Ross & Mike Pinch 
- Week 13: April 16 & April 18: Jason Leto & Trends in GenAI
- Week 14: April 23 & April 25: Present Project 3 

Projects

Students worked in teams to engage with Gen AI tools for cybersecurity tasks. The challenge was to identify a relevant cybersecurity task, gather or generate any necessary source materials, and then leverage Gen AI to manage as much of the task as possible.


- Project 1: Apply prompt engineering methods to produce a document
- Project 2: Use Gen AI to develop a programming (coding)
- Project 3: Use fine-tuning and/or RAG to create a Gen AI model for a specific use case.

Besides the project results, a key component was individual reflection reports from each student discussing what they did, what they learned, and what was successful, and what was unsuccessful.

- Most respondents felt the course **significantly enhanced their knowledge** of the technical aspects, cybersecurity impacts, and ethical/legal implications of Generative AI.
- Common ethical considerations discussed included **deepfakes**, **algorithmic bias** against certain groups, **data privacy** concerns, and issues around **consent and intellectual property** with AI-generated content.
- **Hands-on projects** and access to tools like **ChatGPT-4** were frequently cited as particularly **useful** and **interesting** aspects of the course. Prompt engineering skills were valuable learnings.
- **Confidence in abilities** to use Generative AI for various tasks generally increased from the **1-3 range before** the course to the **3-5 range after**.

- **Likelihood to use Generative AI** in future cybersecurity work also typically increased from **low/moderate before** the course to **high afterwards**.
- Generative AI was seen to have a **positive impact on brainstorming** and a **mixed to negative impact on coordination** and **workload distribution** in team projects. **Debugging** AI-generated code and **consistently getting desired outputs** were challenges.
- A few respondents wished for more coverage of deploying models in real-world settings, analyzing malware/networks with AI, and comparing capabilities of different models.

The class enhanced your knowledge?




Survey Findings

Before: Confidence in your ability to use Gen AI to ...	Average
... accelerate the speed of document creation?	2.83
... improve the quality of document creation?	2.61
... to increase your productivity in code development?	2.83
... extend your coding skills and capabilities?	2.89
... address cybersecurity problems?	2.11
... develop cybersecurity tools with new functionalities?	2.39


Survey Findings

After: Confidence in your ability to use Gen AI to ...	Average
... accelerate the speed of document creation?	4.39
... improve the quality of document creation?	4.17
... to increase your productivity in code development?	4.17
... extend your coding skills and capabilities?	4.11
... address cybersecurity problems?	3.50
... develop cybersecurity tools with new functionalities?	3.78

Likelihood to use Gen AI in your future work	Average
Before	3.06
After	4.17




Based on your experience in the projects, how would you assess your experience of Generative AI's impact on teamwork? For each component of teamwork, check either the positive or negative column. 17 ⓘ



How would you rate the effectiveness of the hands-on projects in the following areas? Use a scale of 1-5, with 1 signifying no effectiveness and 5 signifying maximum effectiveness. 18 ⓘ

	Average	Minimum	Maximum	Count
Increasing your practical knowledge of Gen AI	4.17	2.00	5.00	18
Increasing your practical knowledge of Gen AI specifically for cybersecurit...	4.06	2.00	5.00	18
Increasing your ethical considerations about applying Gen AI	3.50	1.00	5.00	18



Generative AI is a technology that's still rapidly developing. It's designed to create many types of content like essays, pictures, music, and more. It works by using machine learning models, which are complex algorithms that can figure out patterns and make decisions. They're usually made up of neural networks and are trained to understand different forms of data so that they can generate the content they're designed to generate.

I would define "Generative AI" as a tool that can be used to generate accurate and useful responses to one's contextual input, based on the chatbot's knowledge base.

A probabilistic ML model to use generate the most likely "proper" creation in response to a prompt. Prompts are often text based and become interpreted as vectors which act the raw data to be processed by a trained model to generate an image, most likely the "proper" one. The most common example is ChatGPT.

AI software that generates content from prompts given by users. Example is ChatGPT.

Predict what's the best action is for the next timestep.

...

Generative AI is a probabilistic AI model that that creates content based off of the most likely responses to a given prompt. They are often trined on large dates of creative data, and often have to have some form of language processing to break down the prompt into tokens.

...

Generate something from learned information and data.

...

Generative AI is any program that uses historical patterns to predictively create something new. Examples may include text generative LLMs or image generating models.

...

A cutting-edge tool that is trained to offer information based on submitting prompt.

...

Generative AI is some sort of model built on a set of training data that can produce its own content derived from that training. An example is ChatGPT. ...

Generative AI is a form of AI that can combine existing knowledge in new ways usually by predicting what goes next (most text generation) or basing the generation on existing data (most image generation). One classic example is ChatGPT. ...

AI that is able to spit out predictive answers based on the LLMs it has been trained on, like fancy predictive text. ...

Generative AI is the subject relating to learning about different AI models and also learning how to generate responses that accurately pertain to and answer the original inquiry. I think that an example of this would be like learning about prompt engineering. ...

Generative AI is a type of tool that uses a specifically trained language model that is designed to function in specific and general fields. An example of such tool is ChatGPT which is used by a general population for various different tasks. The point of the tool is to generate information from its trained set that is reliable, factual, and constructive.

Generative AI is the usage of a trained model to produce an output (can be text, slides, pictures, media...) that appears like it as been created by a human.

Generative AI is an artificial system that creates new input based off of training data. This training data usually consists of a collections of samples of what it would like as potential outputs. A good example of this is OpenAI or ChatGPT or Midjourney.

Generative AI generates content that fits a users prompt based on a corpus of data.

I'd define it as the creation of a product or idea through the assistance of AI

Generative AI refers to any AI model that takes some kind of prompt as an input and outputs content of some kind based in a way that mimics human behavior.

Generative AI is where artificial intelligence takes in datasets (documents, images, etc...) and generates something a user desires given that background knowledge.

Generative AI refers to deep learning algorithms which generate content of some kind when given tokens. For example, ChatGPT is a generative AI model which takes in text and responds with text.

Using other people's work to train an AI is an ethical consideration to take into account when training these LLMs

...

no ethics sorry

...

I thin generating false information is to be a major ethical consideration going forward.

...

Discussions around fake-porn made me puke :)

...

I thought that the ethical discussion about the AI sentience was interesting. I had not really considered the argument that it was sentient before.

...

Algorithmic bias can lead to certain groups being targeted by artificial intelligence.

...

As a cybersecurity student (and general enthusiast) I found the considerations for data privacy particularly interesting.

...

It is almost impossible to verify that the data has or has not been poisoned, there is a lot of opportunity and this will continue to increase, for malicious actors to take advantage of AI in these early stages especially.

...

One ethical consideration that we talked about was about artwork. I think that this conversation about AI generated art and taking away from actual artists really impacted my view on how AI impacts other fields that are not cybersecurity.

...

Algorithmic bias can lead to certain groups being targeted by artificial intelligence.

...

As a cybersecurity student (and general enthusiast) I found the considerations for data privacy particularly interesting.

...

It is almost impossible to verify that the data has or has not been poisoned, there is a lot of opportunity and this will continue to increase, for malicious actors to take advantage of AI in these early stages especially.

...

One ethical consideration that we talked about was about artwork. I think that this conversation about AI generated art and taking away from actual artists really impacted my view on how AI impacts other fields that are not cybersecurity.

...

The course discussed the issue of Generative AI in the field of deepfakes. In this case, it is possible to use AI to generate life like images^{...} and videos of a person or persons doing activities. These activities are usually things that person being faked has not done or has done but the primary issue is that these fakes are constructed without the person's consent. As a result, the ethics behind generative AI and its use with people pose an issue with manipulating information related to a person.

GenAI definitely being used to level the playing field was really eye opening, since artists can spend thousands of hours learning about^{...} art, where an average Joe can create a picture with a same level of artistic quality

There were very few times this was discussed. One I can remember was the ethics of the AI's "bias" when it came to race, creed, or other^{...} harmful biases that generative AI can have when portraying certain people, things, ideas, etcetera.

Deepfakes & model biases. Discussing how they can "harm" in a non-physical way was a very interesting discussion and put emphasis^{...} on why we need more laws & security in Gen AI. Rather than seeing them as a fun or useful tool it can be used as a weapon as well.

Who should be consider liable when it comes to improper prompting for malicious actions. Along with this I've also thought about the replacement of humans in core positions and if that should be allowed completely. ...

We talked a lot about how AI can unintentionally be trained in a way that causes it go generate content that uses stereotypes and other potentially questionable foundations. ...

AI being used in deepfake adult content. While it does not involve direct harm to an actual person, it still can affect their reputation and their mental wellbeing. ...

Course Syllabus

<https://tinyurl.com/24jm34fz>



Project Examples

<https://tinyurl.com/ydk9nczh>



Learning Outcome

Assessment Method(s)

Explain how generative AI models work and discuss their capabilities and limitations in cybersecurity contexts.

News Presentation,
Project Presentations
and Reports

Identify and critically evaluate potential applications of generative AI in areas like incident response, threat intelligence, and penetration testing.

Project Reports and
Presentations

Demonstrate technical proficiency in making API calls to generative AI and using open source tools

Final Project Report
and Presentation

Develop, implement, and report on case studies applying generative AI to real-world cybersecurity problems.

Project Reports and
Presentations

Discuss ethical implications of using generative AI in cybersecurity, including biases, accountability, transparency, and potential misuse.

Project Reports and
Presentations

Analyze how use of generative AI intersects with relevant laws, regulations, and compliance standards like data protection.

Project Reports and
Presentations

Evaluate when generative AI can be effectively utilized in cybersecurity workflows and when their limitations necessitate human judgment.

Project Reports and
Presentations

Survey Takeaways

- **Project-based work is a good way to let students explore**
- **GenAI can be hard with teamwork**
- **Students overall feel**
 - More positive about potentially using Gen AI for work
 - More confident in their ability to use Gen AI effectively