CSEC-559/659 Generative AI in Cybersecurity Course Syllabus

NOTE: The information presented in this syllabus is subject to expansion, change, or modification during the semester.

Course Description

Generative AI and large language models (LLMs) such as ChatGPT have taken the cybersecurity world by storm, with Microsoft introducing Security CoPilot and Google launching the Google Cloud Security AI Workbench. In this project-based course, we invite students who are already proficient in cybersecurity to delve into the application of Generative AI through real-world case studies.

With an emphasis on hands-on learning, we will give students the freedom to explore, apply, and critique the use of Generative AI in various cybersecurity tasks inspired by their experiences from prior class work and internships. They will also gain experience using non-cloud models that would be needed for proprietary information.

Course Learning Outcomes

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

Prerequisites

CSEC-559: CSEC-380 or Permission of Instructor
CSEC-659: CSEC-742 or CSEC-380 or Permission of Instructor

Course Materials

RIT MyCourses will be used to access course content, submit assessments (e.g., assignments), complete quizzes, and view grades.

Required Materials

No textbook or materials are required. Materials will be provided on MyCourses.

Expectations

This is a 3-credit semester course. You should plan to spend several hours each week outside of the classroom to complete the work required for this course, in addition to the time you attend and participate in class.

Important RIT Deadlines

For term dates, including add/drop and withdrawal deadlines refer to the RIT Calendar.
Use the arrows at the top and bottom of the page to review additional syllabus information.

Instructor Information

Instructor (of record): Dr. Matthew Wright

- Email:
- Office Hours: My office hours will be conducted in person and via Zoom at

Instructor: Dr. Christopher Schwartz

- Email:

Instructor: Nate Mathews

- Email: nate.mathews@rit.edu

We will answer emails and Slack messages within 24 hours between Monday and Friday. Responses to emails and messages received on Saturdays, Sundays, breaks, or holidays may be delayed. We will respond to those emails and messages within 48 hours of the next academic day.

Discord

This course uses Discord for communication between students and the instructor. Students will be invited to the server ( ) and will be expected to check it regularly.

Expectations

All course content, topics, concepts, skills, techniques and tools within this course are for educational and research purposes only. Your use of them must be legal, ethical, moral, responsible, and professional. Contact your instructor with any questions and/or concerns.

Students are expected to adhere to the following. Failure to adhere to these will result in failing the course and being considered for expulsion from RIT. Degree revocation may also be considered. If you do not agree with these expectations or cannot or will not adhere to them, you should drop the course during the add/drop period or withdraw from the course.

- Act and be legal, ethical, moral, responsible, and professional.
- Do not break any laws.
- Do not access any protected (e.g., password protected, etc.) systems or material for which you do not have explicit written permission.
- Only use the course content, topics, concepts, skills, techniques and tools and your knowledge, skills, and abilities for good.
- Do not use the course content, topics, concepts, skills, techniques, and tools or your knowledge, skills, and abilities for any illegal, immoral, unethical, and/or malicious activity.
- Ask the instructor any questions and discuss with the instructor any concerns prior to taking any action.

Warnings

Some course content may include additional acceptable and unacceptable actions, uses, and/or guidelines. The instructor is not a lawyer. No legal advice is provided within this course. Seek legal advice from your legal counsel.

Liability Release

The instructor assumes no liability and will not be held responsible or liable for your actions with the course content, topics, concepts, skills, techniques and tools or your use of your knowledge, skills, and abilities. In addition, the instructor assumes no liability or responsibility for errors or omissions and no liability for damages related to your use of the course content, topics, concepts, skills, techniques and tools. Therefore, the instructor is released from all liability.

Course Schedule

Refer to the myCourses Content and Calendar tabs for details on availability and due dates. All times listed in the course are US Eastern time, unless otherwise noted.

Topics may be added, removed, or modified as time permits.

- 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: Retrieval-Augmented Generation
- Week 7: Feb. 27 & Feb. 29: Fine-tuning
    - Project 3 Assigned
- Week 8: March 5 & March 7: Present Project 2
- SPRING BREAK: March 12 & March 14
- Week 9: March 19 & March 21: Guest Speakers
- Week 10: March 26 & March 28: Present Project 3 Idea
- Week 11: April 2 & April 4: Ethics and Law 2
- Week 12: April 9 & April 11: Guest Speakers
- Week 13: April 16 & April 18: Research & Trends in GenAI
- Week 14: April 23 & April 25: Present Project 3

Components

In this course, we will use *ungrading*, in which you will be given feedback on your work but no grades. You will provide the instructors with self-assessments after each project and assign yourself a final grade. The instructors reserve the right to change your self-assigned grade, but will only do so if it (a) helps your grade, or (b) your self-evaluation is far outside of what the instructors have observed throughout the semester.

As suggested by the students on the first day of class, we believe that the following elements should be considered in your self-evaluation:

- Effort
- Completing assignments on time
- Active engagement, both in class and with the assignments
- Learning
- Making assigned deliverables look presentable/professional

Please feel free to discuss your grade at any time with the instructors.

Late Submissions/Due Date Extensions

See MyCourses for specific due dates. If you believe you will miss a due date, please contact us as soon as possible. Although we are not assigning grades, late submissions impact our ability to provide timely feedback to all students and demonstrate a lack of respect for our learning community. Failing to appear for a presentation is particularly bad. Please hold yourself to a high standard of professionalism when it comes to deadlines.

When you do your self-evaluations, please regard timely delivery as a significant component of your grade. To provide everyone with flexibility for real-world challenges and constraints, we as a class have agreed upon this late submission policy: *If you can foresee problems that prevent you from submitting on time, alert the instructors to the issue ASAP, and at least 24 hours before the deadline. Exceptions will be made for emergencies. Work with the instructors on a new deadline.*

See below for information on RIT's policy on excused absences.

RIT Policy D04.0 Attendance

Provost Statement on RIT Attendance Policy

Emailed Submissions

Submissions must be submitted to myCourses by the due date. Submissions via email or any other means to the instructors will not be accepted.

Extra Credit

Extra credit assignments may be offered during the course. For your self-evaluations, you can treat these much like you treat extra credit in other courses, except that you should assign the value of the extra credit in rough proportion to the effort and attention required, and whether you believe you did a good job.

News Presentations

Students will create and deliver a 5-minute presentation on news related to generative AI and ideally Gen AI in cybersecurity.

Projects

Students will work in teams to engage with Gen AI tools for cybersecurity tasks. The challenge is 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. While each task is likely to generate some documents and possibly code, the main submission will be an individual reflection report from each student discussing what they did, what they learned, what was successful, and what was unsuccessful. Teams will also present their task and their findings to the class.

There will be three projects:

- Project 1: Apply prompt engineering methods to product a document
- Project 2: Use Gen AI to help you develop a program
- Project 3: Apply fine-tuning and/or RAG to create a Gen AI model for a specific use case.

Academic Integrity

As an institution of higher learning, RIT expects students to behave honestly and ethically at all times, especially when submitting work for evaluation in conjunction with any course or degree requirement. RIT requires all students to become familiar with the RIT Honor Code and with RIT's Academic Integrity Policy.

Academic Adjustments

RIT is committed to providing academic adjustments to students with disabilities. If you would like to request academic adjustments such as testing modifications due to a disability, please contact the Disability Services Office. Contact information for the DSO and information about how to request adjustments can be found at the Disability Services Office. After you receive academic adjustment approval, it is imperative that you contact me as early as possible so that we can work out whatever arrangement is necessary.

Title IX

Title IX violations are taken very seriously at RIT. RIT is committed to investigate complaints of sexual discrimination, sexual harassment, sexual assault and other sexual misconduct, and to ensure that appropriate action is taken to stop the behavior, prevent its recurrence and remedy its effects. Additional information and policies are available at:

https://www.rit.edu/fa/compliance/title-ix

https://www.rit.edu /fa/compliance/title-ix-policies-and-resources.

RIT is committed to providing a safe learning environment, free of harassment and discrimination as articulated in our university policies located on our governance website. RIT's policies require faculty to share information about incidents of gender-based discrimination and harassment with RIT's Title IX coordinator or deputy coordinators when incidents are stated to them directly. The information you provide to a non-confidential resource which includes faculty will be relayed only as necessary for the Title IX Coordinator to investigate and/or seek resolution. Even RIT Offices and employees who cannot guarantee confidentiality will maintain your privacy to the greatest extent possible.

If an individual discloses information during a public awareness event, a protest, during a class project, or advocacy event, RIT is not obligated to investigate based on this public disclosure. RIT may however use this information to further educate faculty, staff, and students about prevention efforts and available resources.

If you would like to report an incident of gender-based discrimination or harassment directly, you may do so by using the online Sexual Harassment, Discrimination, and Sexual Misconduct Reporting or anonymously by using the Compliance and Ethics Hotline.

If you have a concern related to gender-based discrimination and/or harassment and prefer to have a confidential discussion, assistance is available from any of RIT's confidential resources (listed below).

RIT Counseling and Psychological Services: 585-475-2261 (V), 585-475-6897 (TTY)
www.rit.edu/counseling

RIT Student Health Center: 585-475-2255 (V) www.rit.edu/studentaffairs/studenthealth

RIT Ombuds Office: 585-475-7357; 585-475-6424; 585-286-4677 (VP)
www.rit.edu/ombuds/contact-us

NTID Counseling and Academic Advising: 585-475-6400
www.rit.edu/ntid/caas

Center for Religious Life: 585-475-2137 www.rit.edu/studentaffairs/religion

Attendance

As stated in the RIT Attendance Policy (https://www.rit.edu/academicaffairs/policiesmanual/d040), "It is the responsibility of all students to attend their scheduled classes regularly and punctually in order to promote their progress and to maintain conditions conducive to effective learning. Absences, for whatever reason, do not relieve students of their responsibility for fulfilling normal requirements in any course. In particular, it is the student's responsibility to make individual arrangements in advance of missing class due to personal obligations such as religious holidays, job interviews, athletic contests, etc., in order that [they] may meet [their] obligations without penalty for missing class."

Please notify your instructors in advance if you expect to be absent from class so that we can determine the best way to address the absence.

Starfish

This course uses the RIT Starfish communication tool in order to promote student success. If I am concerned about a student's academic performance, I may issue an academic alert, which will notify the student and the student's advisors. If you receive an academic alert from Starfish for this course, please communicate with me, your advisor, or other appropriate contact as soon as possible to discuss the issue, its potential impact on your success in this course, and resources to assist you.

RIT Code of Conduct for Computer Use

The RIT Code of Conduct for Computer Use guides the use of computer and network resources at RIT. This is a summary of the RIT Code of Conduct. The full text is available here. The computing, network and information resources of the Rochester Institute of Technology are

intended to support the mission of teaching, scholarly activity and service for the Institute's students, faculty, and staff. Appropriate use of the computer and networking facilities by members of RIT's academic community should always reflect academic honesty and good judgment in the utilization of shared resources and observe the ethical and legal guidelines of our society.

Sharing Protected Information on the Internet

When sharing copyrighted content on the Internet with your classmates, please make sure that you link to a legal source. Repeated access to illegal sources may cause you or your classmates to receive warnings through the Copyright Alert System, as well as possible downgrades in Internet service and possible disabling of your RIT account. View more about the RIT Copyright Infringement policy.

Use of Copyrighted Material

Certain materials used in this course are protected by copyright and may not be copied or distributed by students. You can find more information at https://www.rit.edu/academicaffairs/policiesmanual/c032.

Students may not copy, publish, distribute, display, modify or create derivative works of the course materials in any medium for any purpose. Additionally, students may not sell, rent, lease, trade, or otherwise transfer the course materials.

Use of Artificial Intelligence (AI)

Unlike any other course you may have had, *the expectation in this class is that you use Gen AI tools as much as possible*. “Cheating,” in this class, may constitute not using enough Gen AI in your work. Please be sure to include your own actual observations and reflections in your self-reflection submissions, though you are also encouraged to use Gen AI to enhance those as well.

Emergencies

In the event of a University-wide emergency, course requirements, classes, deadlines and grading schemes are subject to changes that may include alternative delivery methods, alternative methods of interaction with the instructor, class materials, and/or classmates, a revised attendance policy, a revised semester calendar, and/or a revised grading scheme. Should the campus be impacted by weather or other issues, emergency notifications should be available on a banner at the main RIT website (https://www.rit.edu) and may also be posted in other official communications.

Student Support Availability

Please feel free to reach out about any difficulty you may be having that may impact your performance in this course as soon as it occurs. In addition to your academic advisor, we strongly encourage you to contact the many other support services on campus that stand ready to assist you. These include the Academic Success Center, College Restoration Program, Disability Services, English Language Center, Higher Education Opportunity Program, Spectrum Support program, RIT Ombuds, and TRiO Support Services. Students can find out about specific services and programs at www.rit.edu/studentaffairs.

Communication

RIT email is the official communication mechanism. Email will be used to make official announcements. Students are to use email for notification of class absence, grades discussion, and requests. Students may also use email to discuss and ask questions about course topics and assignments. All academic-related communications must be sent from students' RIT email addresses.

Discord will be used for more real-time communication between students and me. Students may only use Discord to discuss and ask questions about course topics and assignments.

Technology in the Classroom

Place all electronics in silent mode to limit classroom disruptions.