UNIVERSITY OF NORTH CAROLINA CHARLOTTE

ITIS 6415/8415 and 4010 Cybersecurity Clinic | 3 credits

Faculty instructors, Email and Office: Available on Canvas

Course Description

This course provides students with hands-on, experiential learning through direct digital engagement with under-resourced regional nonprofits. Operating as a real-world consulting environment, students work in supervised teams to conduct baseline cybersecurity risk assessments. For Fall 2026, this regular clinic course serves as an implementation site for the Consortium of Cybersecurity Clinics' AI Risk Management (AIRM) Pilot and incorporates Faklaris and Ramesh's Participatory AI Risk Assessment (PARA) framework. Rather than isolating these into separate academic programs or external community workshops, we treat AI governance as an operational branch of standard organizational security. Student teams will execute a unified security audit that pairs traditional technical vulnerability scanning with socio-technical AI risk mapping -- evaluating issues like unvetted LLM data exposure, shadow AI usage, data provenance, and user agency within the client's existing workflow.

- **MS/BS/BA:** Can use as capstone course with permission of their program director.
- **Prerequisites:** Successful completion of 9 credit hours in the MS in Cybersecurity program, the BS or BA in Computer Science program, or equivalent work experience, or by permission of the instructor. The course can be repeated with permission.

How will this course be conducted?

This course is designed as a pro-bono Public Interest Technology Clinic where students, faculty, coaches, and community stakeholders work as partners in tech governance.

Classroom Seminar (Week 1-5): Materials and activities cover foundations of public-interest cybersecurity, threat modeling, operational security, and targeted training on the NIST AI Risk Management Framework (RMF 1.0) and PARA principles.

Clinical Service Model (Weeks 6-end): Live client engagement. Teams conduct technical scans, OSINT discovery, and internal ethnographic workflow interviews to map technical vulnerabilities alongside community data risks. Activities culminate in a comprehensive, accessible risk mitigation report delivered to the client.

Meeting Times

Mondays and Wednesdays, 5:30-6:45p, on Zoom and at client in-person locations.

Course Objectives:

- [CO1] Conduct integrated cybersecurity and AI risk assessments for resource-constrained organizations.
- [CO2] Operationalize the NIST AI RMF 1.0 and PARA frameworks alongside traditional cybersecurity standards (such as the MITRE ATT&CK framework).
- [CO3] Gather socio-technical data through staff interviews to catch operational AI vulnerabilities that automated network scans overlook.
- [CO4] Work collaboratively within a structured clinical hierarchy to solve complex security problems.
- [CO5] Translate complex technical vulnerabilities and algorithmic data risks into plain, actionable mitigation playbooks for non-technical leadership.

Required Materials

All materials are provided through Canvas. No purchases are required.

Grading Weights

Category	Percentage
Client Project Deliverables • Engagement Letter & Project Plan (10%) • Mid-point Progress Report (15%) • Final Risk Assessment Report (25%)	50%
Final Client Presentation	20%
Team Participation & Peer Evaluation	15%
Individual Professionalism & Class Attendance / Participation	15%
Total	100%

Grading Scale

Undergraduate:	Graduate:
A 100.0 % to 90.0%	A 100.0 % to 90.0%
B <90.0 % to 80.0%	B <90.0 % to 80.0%
C <80.0 % to 70.0%	C <80.0 % to 70.0%
D <70.0 % to 60.0%	U <70.0 %
F <60.0 %	

Other Course Policies

Syllabus Revision:
The standards and requirements set forth in this syllabus may be modified at any time by the course instructors. Notice of such changes will be by Canvas or email notice.

Non-Endorsement Statement:
This course engages diverse scholarly perspectives to develop critical thinking, analysis, and debate. Inclusion of a reading or other material does not imply endorsement.

Operational Security (OpSec) & Confidentiality:
Adherence to Clinic OpSec procedures is not just an academic requirement; it is for safeguarding the lives and livelihoods of our partners.

- **Code of Conduct:** All students must review and acknowledge the full Cybersecurity Clinic Code of Conduct before engaging with partners.
- **NDA:** All students must sign the Cybersecurity Clinic Confidentiality Agreement before engaging with partners.
- **Data Handling:** Secure communication and storage tools (as outlined in the Client Engagement Plan) must **always** be used with client data. Under no circumstances may any real client data, system logs, intake configurations, or interview transcripts be submitted to any public or third-party AI model without stakeholder consent.

Failure to follow these policies and others listed here may result in removal from the clinic and/or an F in the course, with disciplinary action per UNC Charlotte regulations.

Classroom & Clinic Expectations:
All participants are expected to facilitate a secure, safe, and inclusive working environment. This includes: (1) Respecting teammates and clients' perspectives and unique security contexts. (2) Encouraging accountability and honesty. (3) Reporting any security incidents or ethical dilemmas immediately to the clinic instructors.

We encourage your active participation in class. Each of us may have strongly differing opinions on the various topics of class discussions. The conflict of ideas is encouraged and welcome. The orderly questioning of the ideas of others, including the instructor's, is similarly welcome. However, we will exercise our responsibility to manage the discussions so that ideas and argument can proceed in an orderly fashion. If your conduct disrupts this, you will not be permitted to participate further.

Academic Integrity:
All students are required to read and abide by the Code of Student Academic Integrity. Violations of the Code of Student Academic Integrity, including plagiarism, will result in disciplinary action as provided in the Code. **Students are expected to submit their own work, either as individuals or contributors to a team assignment**. Definitions and examples of plagiarism and other violations are set forth in the Code. The Code is available from the Dean of Students or online at: http://legal.charlotte.edu/policies/up-407. Faculty may require students to demonstrate that graded assignments completed outside of class are their own work.

Policy on Use of AI and Other Productivity Tools:
In this course, students are encouraged to use tools for productivity such as Copilot, Gemini, and Notebook LM that are university-licensed and data-protected, per the "green" status on the AI Software Guidance webpage. No student may submit client data to ANY model or AI tool, regardless of license, without stakeholder consent. This is grounds for an automatic F and referral for disciplinary action. In their use of AI, students are responsible for identifying and removing any factual errors, biases, and/or fake references. Ask the instructors for advice about your use of non-UNC Charlotte tools.

The instructors use AI to assist with tasks such as generating ideas, checking grammar, drafting slide content and in-class activities, identifying research papers and news articles, and organizing materials. The purpose is to support efficiency, not to replace their judgment or expertise. All content has been reviewed and adapted to ensure it aligns with the objectives of this course. We disclose this so you understand that AI can be a helpful resource when used responsibly and critically.

Course Credit Workload:
Students should expect to work an average of 10 hours per week across class meetings, client and team interactions, and prep work and assignments. This workload may fluctuate based on client availability and project milestones. Professionalism in scheduling and communication is non-negotiable. Daily work is encouraged.

Attendance and Absences:
All students are expected to attend every class and clinic meeting and remain for the duration of the session. Regular class and clinic attendance is a student's obligation, as is a responsibility for all the work of class and clinic meetings. To learn more about excusable absences and student responsibilities to account for these and provide notice and documentation, see provost.charlotte.edu/policies-procedures/academic-policies-and-procedures/course-attendance-and-participation.

Failure to attend and/or provide reasonable notice and documentation of absences will result in a reduction of the Class Attendance/Participation grade at the least, and a grade of F at the most. Speak with an instructor if you foresee trouble with attendance.

Non-Discrimination:
All students and the instructors are expected to engage with each other respectfully. Any harassing or discriminatory conduct based on protected characteristics may constitute a violation of University Policy 501, Nondiscrimination. Any student suspected of engaging in such conduct will be referred to the Office of Civil Rights & Title IX.

Online Professionalism:
Video backgrounds must be neutral, and cameras should be positioned in a manner that prevents the introduction of distracting objects, messages, symbols, or other people. Alternatively, students may choose to blur their background, utilize one of the default Zoom virtual backgrounds, or use one of the UNC Charlotte branded backgrounds. Students who do not abide by this policy may be removed from the Zoom session and counted absent or otherwise receive a reduction in their participation grade. Repeated violations may be referred for potential disciplinary action under the Code of Student Responsibility.

All students are required to abide by the UNC Charlotte Title IX Grievance Policy, Sexual Harassment Policy, and the policy on Responsible Use of University Computing and Electronic Communication Resources. Sexual harassment is prohibited, even when carried out through computers or other electronic communications systems, including course-based chats, breakout rooms, or message boards.

Withdrawals:
Students are expected to complete all courses for which they are registered at the close of the add/drop period. If you are concerned about your ability to succeed in this course, it is important to make an appointment to speak with us as soon as possible. The University policy on withdrawal allows students only a limited number of opportunities available to withdraw from courses. It is important for you to understand the financial and academic consequences that may result from course withdrawal.

Incompletes:
The grade of I is assigned at the discretion of the instructor when a student who is otherwise passing has not, due to circumstances beyond his/her control, completed all the work in the course. The missing work must be completed by the deadline specified by the instructor, and no later than 12 months. If the I is not removed during the specified time, a grade of F, U, or N, as appropriate is automatically assigned. The grade of I cannot be removed by enrolling again in the same course, and students should not re-enroll in a course in which they have been assigned the grade of I.

Student Support

Academic Accommodations:
UNC Charlotte is committed to accessibility in education. If you have a disability and need academic accommodations, send the instructors your Accommodation Letter through the DS Portal as early as possible. We encourage you to meet with us to discuss the accommodations outlined in your letter. For more information about accommodations, contact the Office of Disability Services at 704-687-0040 or disability@charlotte.edu.

Mental Health Services:
It is common for college students to experience challenges that may interfere with academic success such as academic stress, sleep problems, juggling responsibilities, life events, substance misuse concerns, relationship concerns, or feelings of anxiety, hopelessness, or depression. If you or a friend is struggling, we strongly encourage you to seek support. Helpful, effective resources are available on campus at no additional cost.

- Visit the Counseling and Psychological Services (CAPS) website at caps.charlotte.edu for information about the broad range of confidential on-campus mental health services, health assessments, hours, and other information.
- Call CAPS at (704) 687-0311 if interested in scheduling an appointment with a counselor. After-hours crisis support is also available through this phone number.
- Contact the Center Wellness Promotion at (704) 687-7407, by email at wellness@charlotte.edu or visit the website at wellness.charlotte.edu for more information on how to develop healthy attitudes and behaviors as it relates to relationships, mental health, alcohol, tobacco, or other drugs and sexual health.

Help with Basic Needs:
UNC Charlotte defines “basic needs” as those needs that, when unmet, can hinder students’ ability to focus on and successfully complete their academic studies. Basic needs include food security, housing security, transportation, health and wellbeing, technology, and childcare. Any student who has trouble in one or more of these areas is urged to contact the Dean of Students Office for support with navigating campus and community resources. Students can also consult the Niner Needs website for a list of helpful resources designed to address student needs. (Niner Needs hyperlink: https://ninerneeds.charlotte.edu/)

Course Design Credits

Some aspects of the syllabus are inspired by those shared by the Consortium of Cybersecurity Clinics, its AI Risk Management Working Group, and member clinics, specifically those at University of California-Berkeley and Indiana University-Bloomington.

Weekly Schedule *(subject to change)*

Week	*Module / Unit*	*Preparation*	*Activities*
1	Introduction to Public-Interest Cybersecurity & AI Governance	**Read:** Pages 7 - 21 & 48 - 52 of *"An Introduction to Cybersecurity Ethics"* *(Shannon Vallor, The Markkula Center for Applied Ethics)* Intro to Public Interest Tech **Explore & use:** The National Cybersecurity Society, facts for small businesses, These fact sheets are tutorials on key cybersecurity topics of concern for your company and your employees NIST Cybersecurity and AI Risk Management Frameworks Participatory AI Risk Assessment principles and methods	Read and acknowledge Code of Conduct Read and sign Confidentiality Agreement Choose your current event topic and preferred presentation order Collect an initial survey data point from Padiyath et al. and from Faklaris et al. to establish baseline understanding Discuss the shifting threat landscape as resource-constrained nonprofits adopt unvetted automated tools
2	Ethics, Threat Diagnostics, and the AI Ecosystem	**Read**: Example articles for cybersecurity and AI risks in public interest contexts Arthur Turner. "Consulting Is More Than Giving Advice." Thomas Wedell-Wedellsborg. "Are You Solving the Right Problems?" Interviewing for Impact Ethical and social risks of harm from Language Models for an overview of the broader ethical and social considerations associated with LLMs. **Explore and use:** List of AI use cases provided in the PJMF AI Use Case Library and identify two use cases that stand out to you as particularly impactful ways to leverage AI tools for nonprofit organizations. OWASP Top 10 for Large Language Model Applications and the OWASP Top 10 for Agentic Applications - get a sense of the security risks associated with LLMs and agentic	Discuss prep materials Discuss "Rules of the Road": *Code of Conduct.* *Personal Risk.* *How to talk about the Clinic.* *Security Response Plan.* Current event briefing Review examples of how nonprofit organizations are leveraging AI tools to further their mission Demystify algorithmic dependencies. Learn how to map client data flows and trace third-party vendor APIs for implicit vulnerabilities.

Week	*Module / Unit*	*Preparation*	*Activities*
		AI, and identify one example from each resource that would be particularly harmful for nonprofit or community organizations.	
3	Threat Modeling & Socio-Technical Risk Taxonomies	**Read:** Electronic Frontier Foundation, "Surveillance Self-Defense: Your Security Plan" - **know** the definitions of underlined terms. Jorge Luis Sierra "Digital and Mobile Security for Mexican Journalists and Bloggers" Le Blond et al. "A look at targeted attacks through the lense of an NGO" SAFETAG Guide. **Skim** to Section 2.2, then **read Section 2.2 and Section 2.3**. **Explore and use:** About PESTLE (use an ad-blocker!) STRIDE and MITRE ATT&CK frameworks EU GPAI Code of Practice (safety and security chapter) Appendix 1 on systemic risks and other considerations MIT AI risk taxonomy	Problem Diagnosis and Reframing How to inventory cyber security and AI assets Current event briefing Discussion on AI's impact for cybersecurity and the breadth of AI risks Technical threat modeling meets socio-technical risk identification Share a template with student pairs to fill out a "Stakeholder Power Map," where they identify who has the most to lose from an AI error, prioritizing those voices as the primary "risk definers" for the rest of the clinic.
4	Information Gathering & The Art of Technical Translation	**Read:** Sarah Jeong, Charlie Warzel, Brianna Wu, Joan Donovan. New York Times. "Everything is GamerGate" - *Read all of the four essays.* Angela Chen. The Verge. "Moderating content doesn't have to be so traumatic" Sam Dubberley & Michele Grant. First Draft. "Journalism and Vicarious Trauma" **Explore:** The EFF's Security Education Companion The Moderator's Toolkit	Discussion of prep materials Current event briefing "Translation Sprint" – practice rewriting technical AI disclosures into plain-language client notifications "Inquiry Pairs" – mapping technical vulnerabilities to specific humanistic harms Discuss how qualitative data can help uncover how clinic AI systems or surrounding infrastructure impact their daily routines and decision-making power. Practice a mini risk /impact assessment for a hypothetical

Week	*Module / Unit*	*Preparation*	*Activities*
		Explore EU AI Act Tracker	client, identifying key stakeholders and degrees of expected impact, using the Australian Government AI Impact Assessment Tool (or the Consortium's own AI Risk Assessment) as the template.
5	AI Security Policy Gaps & Clinical Intake Prep	Review the NIST AI Risk Management Framework Read the U.S. State Department Risk Management Profile for Artificial Intelligence and Human Rights Review Figure 4 and related content from the paper, OECD Framework for the Classification of AI Systems Read chapter 6 on Policy and Governance of the 2025 AI Index Report (pg. 325-364) https://hai.stanford.edu/ai-index/2025-ai-index-report (NOTE this is an annual report and can be updated yearly for the most up to date information) Introduce tools such as Microsoft Counterfit and Credo AI Other materials on Canvas	Discussion of prep materials Current event briefing Discussion on responsibility throughout the value chain or tech stack e.g. when is an developer, deployer, or user held responsible? How are community or end-user or bystander perspectives incorporated, if at all? Conduct mock clinical intake interviews. Practice identifying organizational "red lines" for automated data processing using standard industry frameworks. Collect an intermediate assessment based on the Padiyath et al. (2025) scale and require a short reflection
6	Client Engagement & Boundary Setting	Review assigned Client dossier and initial intake notes. Review resources on interview and communication best practices Begin recruitment and scheduling for stakeholder interviews. Other materials to be assigned.	Discussion of prep materials Current event briefing Formal launch with assigned nonprofit clients. Establish communication protocols and secure data storage zones. Finalize and sign the Client Engagement Letter. Discuss and gather feedback on AI Risk Assessment framework from client staff
7	Digital Footprinting & OSINT Discovery	**Explore and use:** OSINT resources https://inteltechniques.com/links.html OSINT Techniques https://www.osinttechniques.com/osint-tools.html	Discussion of prep materials Current event briefing Research client's digital footprint and external AI dependencies Map out the client's public-facing footprint. Audit external cloud integrations and identify unvetted,

Week	*Module / Unit*	*Preparation*	*Activities*
		OSINT Framework Awesome OSINT **Try**: SECALERTS - Automated Security Audit Other materials to be assigned.	consumer-grade generative AI platforms utilized by staff (Shadow AI). Present initial Discovery Reports for critique
8	Internal Asset Mapping & Workflow Verification	Review client software manifests and privacy disclosures.	Conduct target interviews with staff members to trace data entry points, cloud storage use, and software configurations. Submit the Mid-point Progress Report. Submit a round of survey data to test retention of materials from AIRM modules.
Semester Recess			
10	System Configuration & Vulnerability Review	Map system states against the MITRE ATT&CK framework and OWASP LLM guidelines.	Analyze current software update policies, user credential management, and access controls for internal tools.
11	Technical Security Scans & Infrastructure Testing	Review all techniques in your assigned areas of the MITRE ATT&CK framework and list the ones that can impact our client Determine from that list what should be the top 3 priorities based on size on impact and the client's ability to take defensive action Conduct technical scans and ethnographic interviews	Perform supervised vulnerability scanning across designated client endpoints and configurations. Document critical network errors alongside unsafe data-sharing vectors. Present draft Vulnerability Reports for critique that pair technical issues with human impacts
12	Applying the PARA Framework Pillars	Review the Participatory AI Risk Assessment (PARA) guidelines developed by Faklaris and Ramesh focusing on three core pillars: Agency, Data Sovereignty, and Equity.	Evaluate the client's automated operations against PARA: Is client data used to train third-party models without consent? Are there clear human-in-the-loop overrides for algorithmic assessments? Present draft PARA guidelines report for critique
13	Drafting the Integrated Audit Report	Review the client's existing IT and data-use guidelines (if available) and other materials from prior weeks.	Synthesize technical network scan logs with qualitative workflow data. Build a unified risk ledger pairing systems flaws with operational privacy gaps.
13	Policy Revision and Participatory Design	Describe defensive countermeasures the client can take either directly, e.g. create a separate password protected wifi guest account, or indirectly, e.g.	Discuss with client staff their organizational red lines for AI use and for usable security and privacy

Week	*Module / Unit*	*Preparation*	*Activities*
		request a vendor of a critical SaaS application to share SLAs, RPOs and RTO, and negotiate according to defensive requirements Using the client's current draft of an AI policy (or a provided template if they have none), students will evaluate the policy against the PARA "Three Pillars of Responsibility": *Agency & Disclosure:* Does the policy mandate that the nonprofit's clients are informed when they are interacting with or being evaluated by an AI system? Is there a clear "human-in-the-loop" opt-out for critical services? *Data Sovereignty:* Does the policy protect the data of marginalized groups? For example, if a nonprofit uses a generative AI tool to summarize client intake notes, where does that data live, and is it used to train third-party models? *Equity & Redress:* Does the policy provide a clear pathway for a community member to contest an algorithmic decision (e.g., being denied a service based on a screening tool)?	Discuss the AI Risk Assessment framework with the client and revise as needed Use design probes, such as visual or tactile prompts (Tabassum et al. 2026, Yildirim et al. 2023) to help non-technical users describe how they want their data handled, that are known to elicit qualitative data on people's preferences Draft modular, low-overhead acceptable use policies for generative tools and core systems tailored for resource-strapped IT environments
14	Technical Mitigation Playbooks	Review prior materials.	Complete the step-by-step remediation plan for the client. Focus heavily on accessible configuration changes, password management, and turning off third-party training data toggles
14	Report Dual-Review & Presentation Rehearsal	Identify and incorporate feedback, critiques and "vetoes" from preceding activities and sessions.	Submit the draft report for formal dual-review oversight by instructors and industry mentors. Rehearse the translation presentation to ensure clarity for non-technical stakeholders.
15	Final Hand-off & Project Debrief	Finalize all artifacts; complete the Padiyath et al. Agency Index survey and other required surveys for the AIRM project.	Deliver the Final Risk Assessment Report and give the concluding presentation to client leadership. Complete the final project debriefing and hand off all secure resources to the client. Update your professional interest profile to indicate if this clinical experience has influenced your interest in careers such as Trust and Safety, AI Governance, or Public Interest Technology.

Curriculum Credits

This weekly schedule is adapted from the 2023-2025 University of North Carolina at Charlotte Cybersecurity Clinic blueprint, with modifications incorporating the 2025-26 Consortium of Cybersecurity Clinics AI Risk Management Working Group materials and Faklaris and Ramesh's 2026-28 Participatory AI Risk Assessment course curriculum project.