
CSCE 6280: Advanced Topics in AI (spring 2027)

Lawful AI: Privacy, Security, and Safety

Instructor: Supreeth Shastri

Timing: Tu/Th 1:50 - 3:10 PM

Location: NTDP F285

Course description

Why does an AI company need a constitution? What does it mean for an AI system to be lawful, and who gets to decide? This PhD seminar takes the view that making AI lawful should not be left entirely to lawyers, and that computer scientists can design and evaluate lawful AI systems without needing a law degree. We will examine **AI privacy** through the GDPR, **AI security** through the Cyber Resilience Act (CRA), and **AI safety** through the AI Act. We take a first principles approach to explore how legal requirements translate—or fail to translate—into technical properties, system behavior, evidence, and mechanisms. *If questions about AI are keeping you up at night, this course may at least give you better questions to lose sleep over.*

Course structure

The course includes instructor-led lectures, student-led discussion of research papers, flipped analysis of legal materials, and guest lectures from two prominent legal scholars working on AI. We will also work with artifacts and incidents that do not fit neatly into a traditional paper-reading course. For example, Anthropic's Claude Constitution, AI safety and governance frameworks, and investigations of AI incidents such as the OpenAI–Hugging Face attack.

Student teams will undertake a semester-long research project at the intersection of AI and law. Projects may involve software development, empirical experimentation, AI incident investigation, legal analysis, or a combination of these approaches. Promising teams will receive additional mentoring toward developing their work into papers for AI/ML, security/privacy, or interdisciplinary research venues.

Who should take this course?

The course is intended for graduate students interested in exploring the intersection of AI with privacy, security, and safety through the lens of law. No legal background is expected; however, students are required to have completed a course in AI, ML, NLP or a closely related area. This course is ideal for students interested in exploring how legal obligations can be expressed, implemented, and evaluated as properties of AI systems.

Need more info?

Reach out to the instructor via email (supreeth.shastri@unt.edu) or explore our research in the Lawful Computing Lab (<https://lawfulcomputing.org>)
