Introduction to Privacy Engineering

Spring 2020

Time: TBD

Instructor: Tamara Bonaci (tbonaci@)

Office hours: By appointment

Teaching Assistant: TBD

Course website: TBD

Course assignments and dropbox: TBD

Course discussion board: TBD

Course gradebook: TBD

Course mailing list: TBD

Course Overview:

You are an owner of a successful start-up, providing a service where your system continuously monitors various health and activity information about your customer, in order to better manage the customer's chronic health condition. Your business relies on cloud computing technologies, mobile devices, and recent developments in artificial intelligence.

As a part of your business, you collect throws of data about your customers. The data that you collect, mine and guard is valuable – without it, your business doesn't exist. However, the data is also highly sensitive (e.g., it consists of patients' biomedical, activity, location information, collected over an extended period of time)! So, even though you are only planning to use that data for legitimate purposes (and you are never planning to abuse it), you need to understand how might such data be misused, and by whom. More importantly, you need to understand what your responsibilities are, and what are technical approaches that you can implement in order to use the data only as intended.

This course surveys recent privacy approaches, applicable to various technical disciplines. In doing so, the course puts an emphasis on biomedical technologies, and explores inference threat arising from recent advancements in artificial intelligence. We will briefly discuss the history of privacy, and compare major legal frameworks for privacy. We will then explore the phenomena of unravelling of privacy and inverse privacy. Next, we will survey several privacy design frameworks. We will then explore approaches such as anonymization, differential privacy, and polymorphic encryption.

Course Goals:

By taking this course, the students will:

- Gain an appreciation for the importance of privacy in engineering.
- Learn about current research and tech policy efforts in the areas of security and privacy, privacy engineering and usable privacy.
- Learn how to collect, analyze and reconcile the diverse technological, business and tech policy aspects impacting how and why information is collected, managed, used and shared in some system.
- Evaluate system designs based on privacy principles, and privacy requirements.

- Learn how to integrate privacy into the engineering lifecycle phases.
- Work as a member of an interdisciplinary team, to address critical system requirements in a privacy-sensitive ecosystem.

Course Outcomes:

Upon taking this course, students will:

- Be familiar with the different privacy aspects, applicable to diverse engineering systems.
- Develop critical thinking about the strengths and weaknesses of various privacy approaches and frameworks.
- Be able to implement different privacy paradigms, and embed them into various systems during both design and implementation phases.

Course Progression:

The following is the class progression covering the 10 weeks for the course.

Week 1: Course overview. Introduction to information security, privacy and policy.

Week 2: Foundations of privacy. Introduction to law of computer technology.

Week 3: Introduction to anonymization

Week 4: Static data anonymization – multidimensional data and time-series data.

Week 5: Introduction to differential privacy.

Week 6: Differential privacy II.

Week 7: Introduction to privacy-preserving cryptographic methods, and their applications.

Week 8: Polymorphic encryption and pseudomization. Conceptual hashes.

Week 9: Emerging topics: privacy engineering and biomedical applications.

Week 10: Emerging topics: privacy engineering and AI.

Finals week: Project presentations

About the Course:

The course will consist of ***readings and discussion, classroom presentations, privacy review assignment, and a project.***

Readings and Discussion:
Readings and discussions of assigned papers are an important part of this course. The goal of this exercise is to work together on understanding broader implications of the seminal security and privacy research. Additionally, we will also try to keep up with the state-of-the-art security and privacy research, with the focus on the recent privacy engineering work.

Prior to every lecture, a single research paper will be assigned, and you will be expected to post to the class discussion board about it. Your post should contain something original beyond what others have already posted. You may consider posting:

- A summary of the paper,
- Evaluation of the paper's strengths and weaknesses,
- Open research question related to the topic of the paper, or
- Question you would like to discuss in class.

All posts are due by **4pm on a lecture day,** and they will be graded on the scale of 0-2, where:

- 0 means missed or irrelevant post,
- 1 means a relevant post submitted, and
- 2 means a good and interesting post.

Post submitted after 4pm on the day of the class will receive no credit. There will be a total of nine reading assignments, and we will take eight best scores when determining your grade.

Classroom presentation:

Classroom presentations will be aligned with the general topic presented and discussed in class in a particular week, but with an intention of exploring either:

- Emerging and/or more advanced parts of the considered topic,
- Particular application areas for the considered topic,
- An intersection of the consider topic and another related discipline

As such, classroom presentations will follow a **flipped-classroom format**, and your classmates and your instructor will be learning from you.
Topics covered by classroom presentations will be announced during the first week of classes. You will be expected to sign up for ***one presentation***. Each topic will be covered by ***at most two persons***, who will be expected to work together in preparing a 30-minutes long presentation. The presentation may consist of slides, a poster, a demo, or any combination thereof, and that material will be due by **5:30pm on the day of your presentation**.

You (as a group) will be graded based upon the quality of your presentation.

Privacy Review:
This course aims to sharpen our *privacy-conscious mindsets*, and to get us all to thinking about the world in a different way. In the light of that, this exercise is designed to get you to thinking about privacy (and security) in a context you might not normally do.

With the privacy review, your goals will be to evaluate the potential privacy (and security) issues of a new and/or emerging technology, and to discuss what could be done to address those potential threats.
You will choose a technology to analyze (you might get an idea for a technology from various technology news sources), and in a short (2-3 pages) write-up, you will present:

- A short summary of the evaluated technology,
- An analysis of potential privacy goals and weaknesses of the technology,
- An analysis of possible attackers and threats, and
- A short discussion about possible defense strategies.

The privacy review will be due in the middle of the quarter

Project:
The final component of this course is a project, and its goal is to give you a deeper understanding of how to think about, and how to solve a real-life problem from a privacy-engineering perspective.

For the project, you will choose a topic related to any area of privacy and security (including those not directly covered in this course). You can work on the project either individually, or in groups of

up to three persons. When working in a group, your end result should reflect the fact that it is a multi-person effort.

Your work on the project will consists of several milestones:

- Project proposal,
- Progress report,
- Final report, and
- Project presentation.

Grading:

Your grade in this course will be based on readings and discussion, homework assignments, privacy reviews and project. The expected grade breakdown is:

- Readings and discussion – 20%
- Privacy review – 10%
- Classroom presentation – 15%
- Project – 55%

Course Material:

There is no required textbook for this course, but some recommended books that you might want to consider include:

- K. Du, *Computer Security, A Hands On Approach*, CreateSpace, 2017
- Woodraw Hartzog, Privacy Blueprint: The Battle To Control and Design New Technologies
- MITRE Privacy Engineering Framework, online: https://www.mitre.org/publications/systems-engineering-guide/enterprise-engineering/engineering-informationintensive-enterprises/privacy-systems-engineering
- NIST, Privacy Engineering Program, online: https://www.nist.gov/itl/applied-cybersecurity/privacy-engineering
- NIST, An Introduction to Privacy Engineering and Risk Management in Federal Systems, online: https://nvlpubs.nist.gov/nistpubs/ir/2017/NIST.IR.8062.pdf

Course Policies:

Collaboration: In this course, we want you to learn from each other. Therefore, you are allowed (and encouraged!) to talk to your classmates and other students about all course assignments. You may also consult outside reference materials, and the instructor. However, all material that you decide to turn in should reflect your own understanding of the subject matter at the time of writing. This means that you should write your own posts about the assigned papers, your own privacy review, and your own homework. If you work with someone else on any assignment, please include their names on the material that you turn in.

Assignment Turn-in: Posts about the assigned papers should be submitted using the course discussion board. All other material (privacy reviews, homework, and project-related material) should be submitted in a PDF format, using course website on Canvas. Please, ***do not use*** email for assignment submissions.

Late Assignment Turn-in: Discussion board posts are due **by 4pm on the day of the class**, and no late turn-ins will be accepted. All other assignments are due **by 11:59pm PST on the assigned date,** but we understand that you may have to sometimes turn them in late. The grading penalty is 10% of the grade that you would otherwise receive for each day, or part of the day, that you are late. No submissions will be accepted after a week.

Checking grades: Grades will be posted to the course gradebook.