CMPUT 626 - A2
Machine Learning and Practical Privacy

Fall 2023, Tuesday/Thursday 3:30-4:50pm

Today

- Some logistics
- What is this course
- A bit on privacy and machine learning

Instructors

Bailey Kacsmar

- kacsmar@alberta.ca
  - https://bkacsmar.github.io//MLandPracticalPrivacy/
- Instructor office hours:
  - Thursdays 5pm in ATH 3-17, or by appointment

Instructors

Bailey Kacsmar

- kacsmar@alberta.ca
  - https://bkacsmar.github.io//MLandPracticalPrivacy/
- Instructor office hours:
  - Thursdays 5pm in ATH 3-17, or by appointment

Course Website

- The course website is at:
  - https://bkacsmar.github.io//MLandPracticalPrivacy/
  - eclass: https://eclass.srv.ualberta.ca/enrol/instances.php?id=90128
- We will use eclass
  - will be updated regularly, syllabus, calendar, lecture notes, additional materials, assignments/graded work, and policies.
  - It is your responsibility to keep up with the information on both eclass and the course site.
- Questions can be posted to eclass discussion forum.
- Paper reviews (more later) will be posted to eclass
- Some communication may be sent to your ualberta email

Course Syllabus

- Be familiar with the content in the course syllabus
- It is available on the instructors website and e-class
- It contains a selection of papers related to this course. We will cover a **subset** of them this term.

IF you haven't reviewed the syllabus, do so after this lecture.

Plagiarism and Academic Offenses

We take academic offenses very seriously

- Plagiarism applies to both text and code
- You are free (and encouraged) to exchange ideas, but no sharing code or text
- See syllabus for more information and advice

Plagiarism Con’t

- Common mistakes
  - Excess collaboration with other students
  - Using solutions from other sources (like for previous offerings of this course, maybe written by yourself)
  - Asking public questions containing (partial) solutions online
  - Posting (partial) solutions to public websites (e.g.,github)
- Possible penalties
  - First offense (for assignments; exams are harsher), 0% for that assignment, -5% on final grade
  - Second offense, more severe penalties
  - More information on course syllabus

A note on security…

- In this course, you will be exposed to information about security problems and vulnerabilities with computing systems and networks
- You are not to use this or any other similar information to test the security of, break into, compromise, or otherwise attack, any system or network without the express consent of the owner
- You will comply with all applicable laws and uAlberta policies

What is this course? Learning Outcomes

By the end of this course students should be able to:

- Explain the notion of privacy within the machine learning space.
- Analyze security and privacy of machine learning protocols.
- Evaluate **research on privacy of machine learning** and articulate advantages and limitations.

Grading Scheme

- 20% Seminar-style presentations as discussion lead
  - (2-3 throughout the term)
- 5% Quality of feedback on peers
- 20% Paper reviews (2 per week)
- 10% Participation in paper discussions
- 35% Research project
  - 10% Project proposal, (Due October 13, 2023 at 4pm MT)
  - 24% Final project report (Research paper style, Due December 8, 2023 at 4pm MT)
- 10% Final project presentation
  - (To be scheduled. Nov 28, 30, Dec 5, 7)

Late Policy Graded Material - Excluding Reviews

- Due 4pm on the day of the deadline
- Late submissions will be accepted up to 72 hours after the deadline (without penalty) and no documentation needed
- Note:
  - No assistance is available after the deadline
  - No submissions after the 72 hour window
- **Important**: This policy does not apply to paper reviews or peer feedback. Reviews are due 30 minutes before class

Paper Reviews

- **One** due each day before class that has a presentation
  - So, two per week
- Review consists of
  - Write a short paper summary (3-6 sentences)
  - Identify the contributions of the paper
  - Identify the research question(s) of the paper
  - Identify the strong components of the paper (e.g., reasons likely associated with its acceptance, strong executions, etc). Identify 2-4 such attributes
  - Identify weak components of the paper (e.g., revisions that would improve validity, aspects that could improve breadth/impact, etc). Identify 2-4 such attributes.
  - A 1-2 sentence statement as to why you think this paper was included in the course/it's relationship to the content thus far.
  - Identify one possible research question that could be follow up for this paper.

Paper Presentations

- Each student will do **2-3** paper presentations this term
- Each presenter will create slides for a 20 minute presentation
- Lead a 15-25 minute discussion section.
- The presentation should highlight the research questions, methodology, results, and take-aways/impacts of the work.

Peer Review

- Each student will receive feedback on their seminar presentations from each course participant.
- The presenting students' grade is not affected by these evaluations.
- Quality of feedback is graded
  - “your presentation was really bad", versus
  - "Slide n: the amount of content made it difficult to know what to focus on. Consider splitting into two, using boxes for emphasis, or removing any non-critical content"
- This feedback will be submitted via e-class to each presenter before the next class.

Project: Topic Approval

- Focus on academic venues
  - e.g., USENIX Security Symposium, ACM CCS, IEEE Symposium on Security and Privacy, Privacy Enhancing Technologies Symposium (PETS), or the NDSS Symposium.
- Your topic must be approved in advance by the instructor before you submit your proposal.
- Email the instructor at least one week before the proposal deadline (October 5, 2023) with a brief (1-3 sentence) description of your intended topic.

Project: Proposal

- One page in length
- Include at least 10 references,
    - preferably including (but not limited to) papers from the aforementioned venues.
- It is recommended but not required that you discuss the proposal with the instructor first.
- Email your proposal to the instructors by Friday October 13, 2023 by 4pm MT.
- **Presentation**, October 17th, 2023.

Final Project

- Projects will be done in teams of two.
    - Exceptions (e.g., groups of one or three) are permissible only if they have acquired prior approval from the instructor before the proposal deadline. Note that a group of three would be expected to accomplish ``more" than a group of two proportionately.
- Your paper should:
    - Include related work (a summary of past and current work on your topic
    - Provide a concise summary of work, emphasizing major accomplishments, rather than a detailed accounting of individual pieces of research activity.
    - Be like a small research paper (or a workshop style submission) with preliminary results, but perhaps not large datasets/large iterations of experiments for example.
- Final presentation will be a conference style presentation of your work

Paper Selections - For Presentations

Send me a list of:

- Five papers you want to present (indicate your favorite)
- Five papers you would be okay presenting

Send by September 12th, 2023

Let me know any scheduling conflicts

Privacy and Machine Learning?

A Technical Privacy Family: Something to Learn

Do other cats get fed more than me?

Wait a minute…my humans are the ones that feed me…who's the subject?

Can Bob be trusted with my food habits…

I don't want Bob to know I have a cooler food dish…

Data and Abstraction

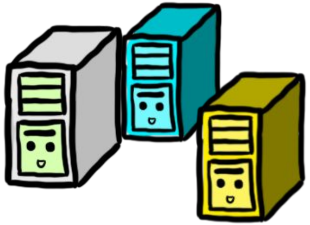

A company wants to analyze data

But the data has privacy implications for the data subjects

Researchers develop technical solutions

Data, Beyond the Abstraction

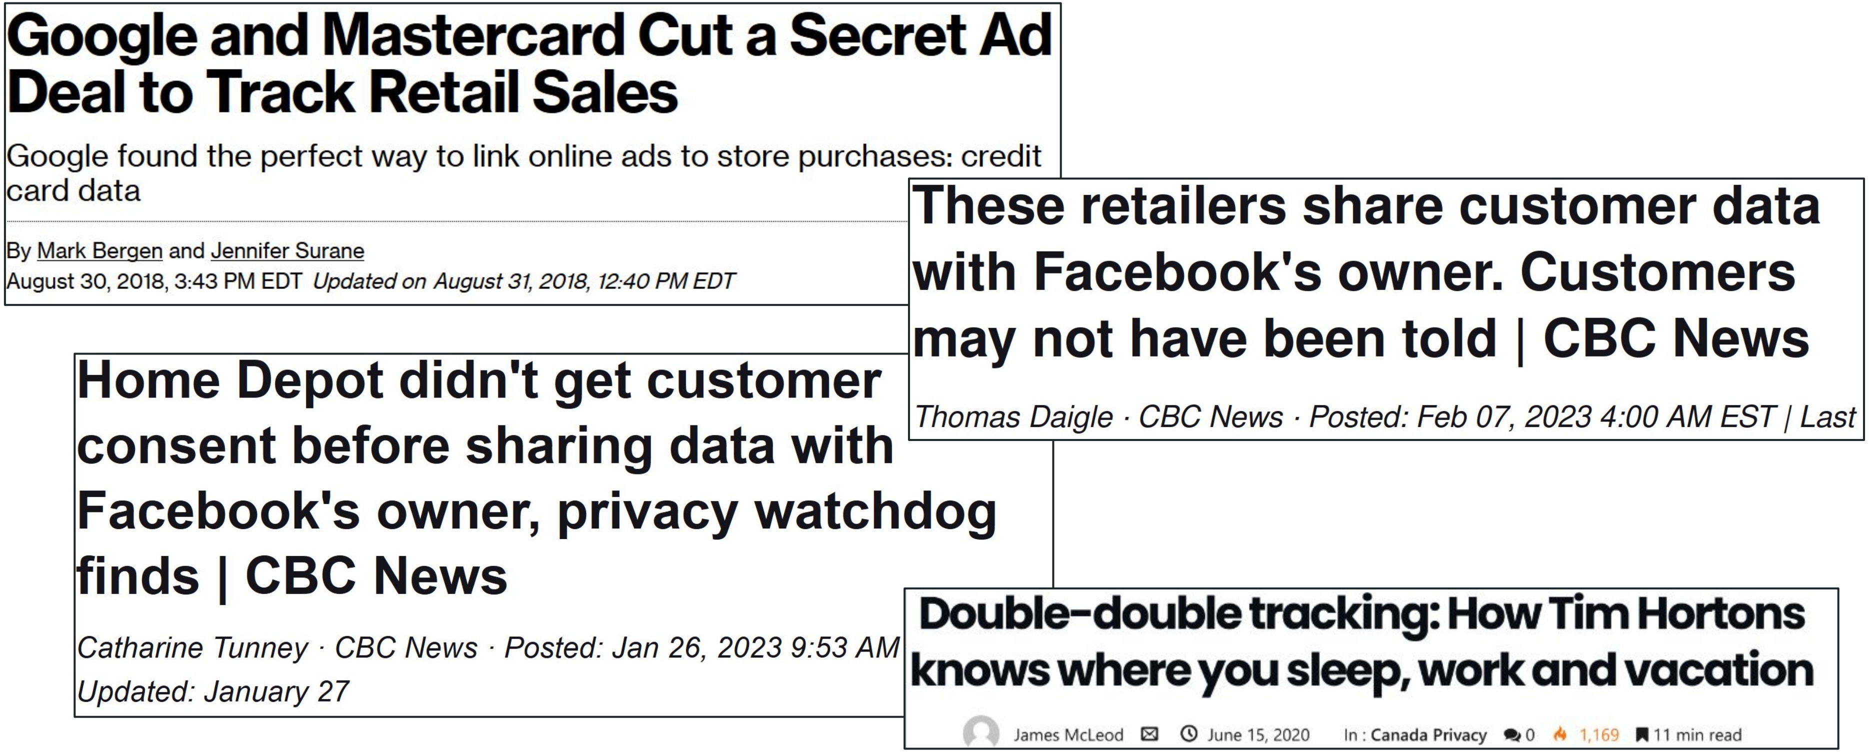

Google and Mastercard Cut a Secret Ad Deal to Track Retail Sales

Google found the perfect way to link online ads to store purchases: credit card data

By Mark Bergen and Jennifer Surane
August 30, 2018, 3:43 PM EDT *Updated on August 31, 2018, 12:40 PM EDT*

These retailers share customer data with Facebook's owner. Customers may not have been told | CBC News

Thomas Daigle · CBC News · Posted: Feb 07, 2023 4:00 AM EST | Last

Home Depot didn't get customer consent before sharing data with Facebook's owner, privacy watchdog finds | CBC News

Catharine Tunney · CBC News · Posted: Jan 26, 2023 9:53 AM
Updated: January 27

Double-double tracking: How Tim Hortons knows where you sleep, work and vacation

James McLeod ✉ June 15, 2020 In : Canada Privacy 0 1,169 11 min read

How Data is Used Continues to Evolve

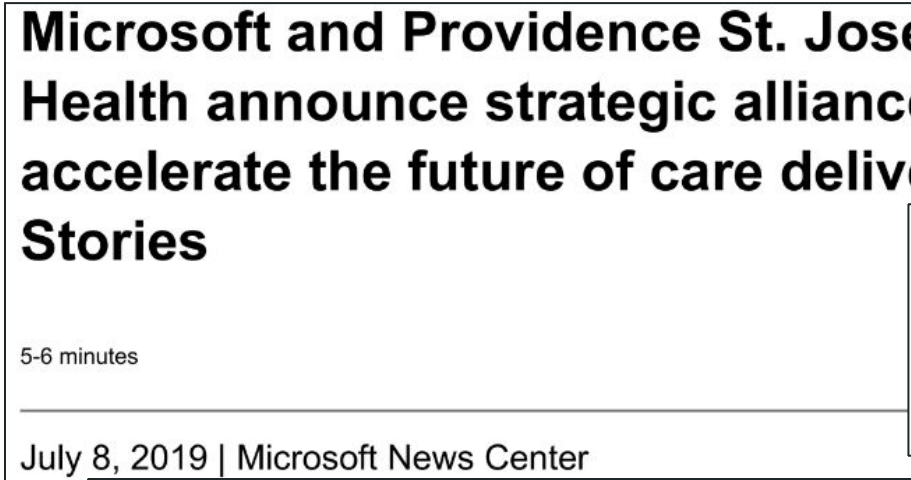

Microsoft and Providence St. Joseph Health announce strategic alliance to accelerate the future of care delivery - Stories

5-6 minutes

July 8, 2019 | Microsoft News Center

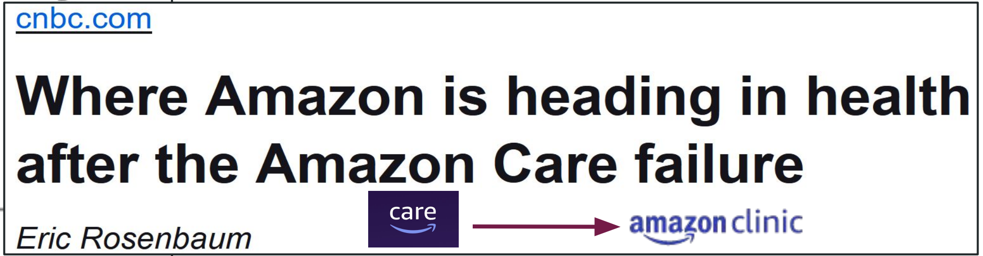

cnbc.com

Where Amazon is heading in health after the Amazon Care failure

Eric Rosenbaum

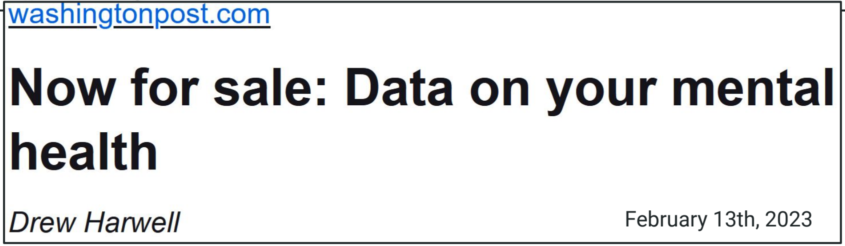

washingtonpost.com

Now for sale: Data on your mental health

Drew Harwell February 13th, 2023

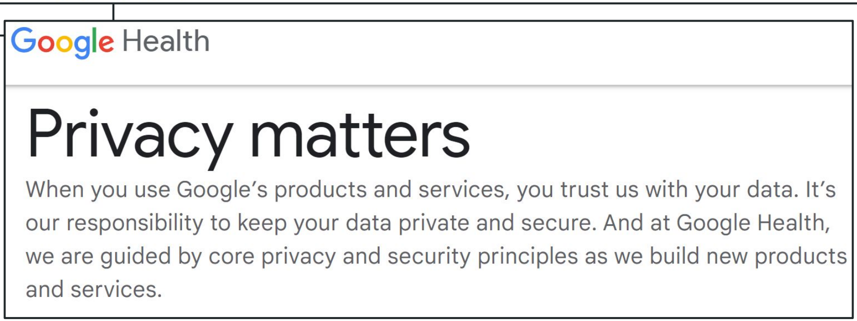

Google Health

Privacy matters

When you use Google’s products and services, you trust us with your data. It’s our responsibility to keep your data private and secure. And at Google Health, we are guided by core privacy and security principles as we build new products and services.

The Use of Health Data

P20986: “It depends. I think it can be beneficial **under certain circumstances**, but I would be hesitant having any healthcare data shared outside my practitioners. However, I recognize how it can improve goods/services, but there **has to be a lot of protection** in place **anytime data is shared**"

P94865: “**Repugnant**, especially in light of for profit health systems attempting to maximize profitability from patient interactions"

B. Kacsmar, K. Tilbury, M. Mazmudar, and F. Kerschbaum. "Caring about Sharing: User Perceptions of Multiparty Data Sharing." In 31st USENIX Security. 2022.

A Wider View of Technical Privacy

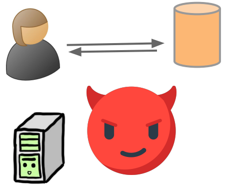

Technical Privacy

Conceptual Privacy

Legal Privacy

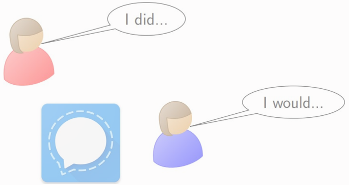

Usable Privacy

Define, **what** is being protected, from **who**, and under what **conditions** this protection will hold.

Technical Solutions for Privacy Problems

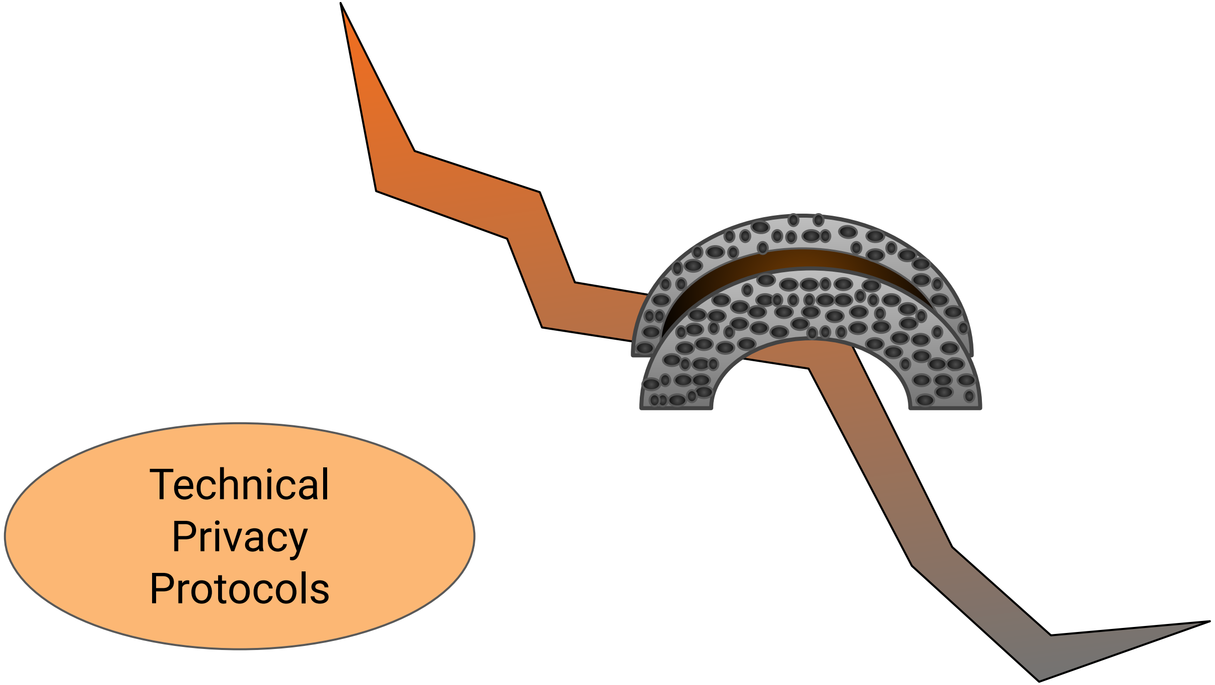

Technical Privacy for Machine Learning?

Training Data

Models

Inferences/Outputs

Define, **what** is being protected, from **who**, and under what **conditions** this protection will hold.

Privacy for Machine Learning

Training Data

Models

Inferences/Outputs

Unintentional Leakage

Intentional Leakage

Define, **what** is being protected, from **who**, and under what **conditions** this protection will hold.

Privacy for Machine Learning

Training Data	**Models**	**Inferences/Outputs**
Unintentional Leakage		**Intentional Leakage**
Data Subject	**Data Owner**	**Access Control**

Define, **what** is being protected, from **who**, and under what **conditions** this protection will hold.

Is this enough?

Define, **what** is being protected, from **who**, and under what **conditions** this protection will hold.

Data Security and Privacy: Vulnerabilities

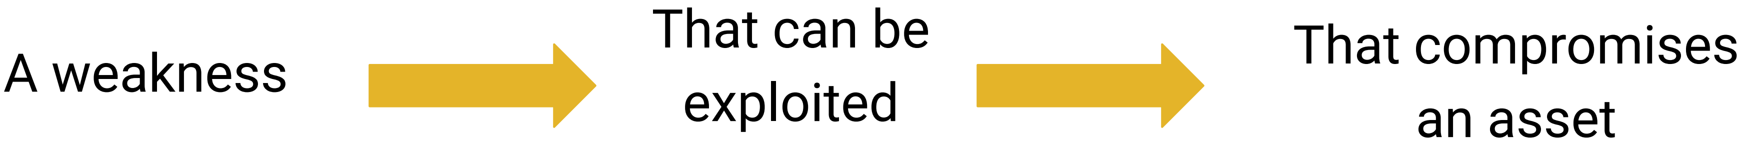

A weakness

That can be exploited

That compromises an asset

Data Security and Privacy: Threats

- Loss or harm
- Interception
- Interruption
- Modification
- Fabrication

These **threats** are part of a **threat model**. Recall the **what** is being protected, from **who**, and under what **conditions**

Data Security and Privacy: Attack

Exploit a vulnerability

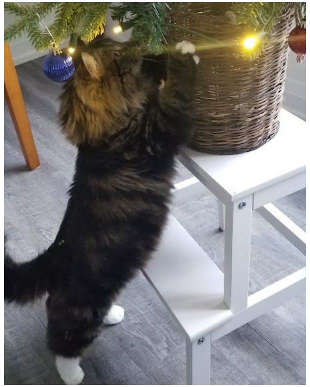

Execute a threat

Data Security and Privacy: Control and Defense

“Security” Tape

Remove or reduce a vulnerability

Control to prevent attacks and defend against threats

Dealing with Attacks

- Prevent it
- Deter it
- Deflect it
- Detect it
- Recover from it

Is this enough?

Define, **what** is being protected, from **who**, and under what **conditions** this protection will hold.

Is this enough?

Who decides?

Define, **what** is being protected, from **who**, and under what **conditions** this protection will hold.

Is this enough?

Who decides?

What conditions?

Define, **what** is being protected, from **who**, and under what **conditions** this protection will hold.

Technical Solutions for Privacy Problems

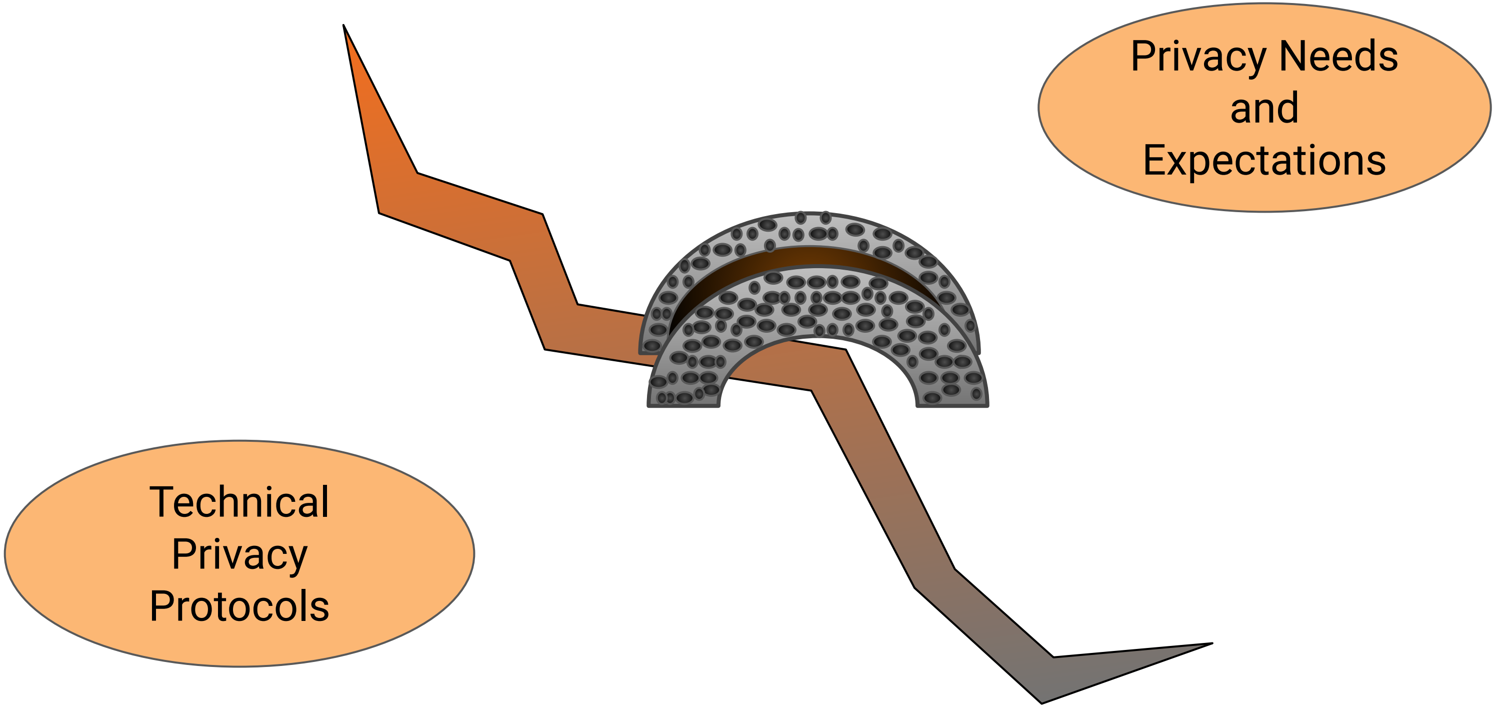

Privacy Needs and Expectations

Technical Privacy Protocols

A Wider View of Technical Privacy

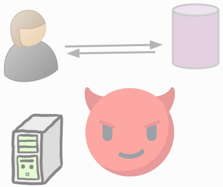

Technical Privacy

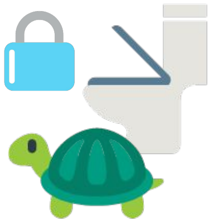

Conceptual Privacy

Legal Privacy

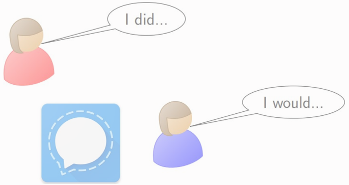

Usable Privacy

Understanding privacy notions and behaviours, **right to privacy**, and privacy expectations

M. Oates, et al. Turtles, locks, and bathrooms: Understanding mental models of privacy through illustration." Proceedings on Privacy Enhancing Technologies 2018.

A Wider View of Technical Privacy

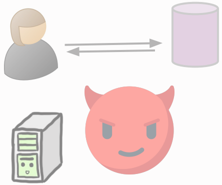

Technical Privacy

Conceptual Privacy

California Consumer Privacy Act
PIPEDA
Bill C-27
GDPR.EU

Legal Privacy

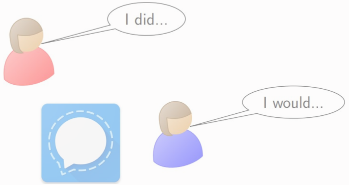

Usable Privacy

“Trusted-third parties”, “Partners”,

A Wider View of Technical Privacy

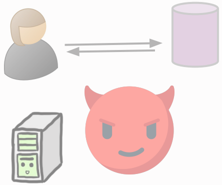

Technical Privacy

Conceptual Privacy

Legal Privacy

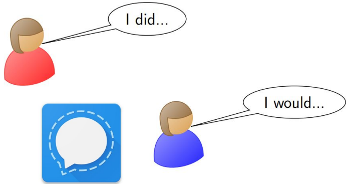

Usable Privacy

What do users actually do? What do they want to do?

A Wider View of Technical Privacy

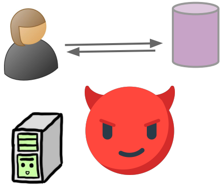

Technical Privacy

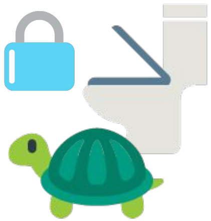

Conceptual Privacy

Legal Privacy

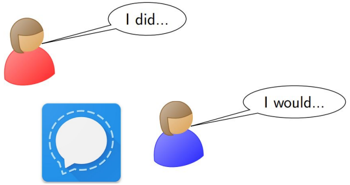

Usable Privacy

Develop Technical Privacy Solutions Informed by the Breadth of Privacy Notions

Questions? Day one mini office hours