Kick-off: Data Privacy Technologies

Chair for IT Security / I20
Prof. Dr. Claudia Eckert
Technical University of Munich

Georg Bramm
georg.bramm@aisec.fraunhofer.de

Martin Schanzenbach
martin.schanzenbach@aisec.fraunhofer.de

February 2, 2021

Outline

Organization

The seminar will be organized as a scientific conference:

1. Familiarization phase (2 Weeks)
2. Writing phase (12 Weeks)
3. Review phase (2 Weeks)
4. Improvement phase (1 Week)
5. Talk preparation (min 1 Week)
6. Talk and Discussion

Requirements

- Report
  - Written report in the form of a scientific paper
  - Mandatory length of 6 pages (references don't count)
  - Usage of LaTeX is mandatory
  - Formatting with the provided LaTeX-Style (IEEE 2-column)
- Review
  - Every Student creates two anonymous reviews
  - Review template will be provided
  - Approximately 1/2 page
  - Every Student writes a rebuttal
- Presentation
  - Presentation with slides
  - 30 minutes presentation
  - 15 minutes discussion

Grading

Grading considers all contributions to this seminar:

1. Report (50%)
   - Contents, Accuracy, Style, Effort, Grasp
2. Presentation (30%)
   - Slides, Execution, Contents, Understandability
3. Reviews (15%)
   - Written Reviews and Rebuttal
4. Participation and discussion (5%)

Time Table (tentative)

02.02.21	Kick-off meeting (today)
15.03.21	Topic Assignment
02.04.21	Introduction to scientific writing (recommended)
18.06.21	Deadline for report (pre-final) submission
21.06.21	Review Assignments
02.07.21	Deadline for review submission
09.07.21	Deadline for rebuttal submission
09.07.21	Deadline for final report submission
09.07.21	Deadline for presentation submission
until 17.07.21	Presentations and discussion

Before we go on....

... any questions so far?

Topics

- Building Blocks of Privacy-enhancing technologies
  1. A comparison between identity standards (OpenID, DID).
  2. Verifiable Random Functions and their Applications.
  3. Distributed/decentralized Private Information Retrieval.
  4. A Survey on Hierarchical Deterministic Key Derivation.
  5. Privacy at the Transport Layer.
- Privacy in Machine Learning
  1. Training data privacy
  2. Input/Output privacy
  3. Model privacy
  4. Pivacy-preserving Federated Learning
  5. Privacy attacks on Machine Learning Models

Topics

- Building Blocks of Privacy-enhancing technologies
  1. A comparison between identity standards (OpenID, DID).
  2. Verifiable Random Functions and their Applications.
  3. Distributed/decentralized Private Information Retrieval.
  4. A Survey on Hierarchical Deterministic Key Derivation.
  5. Privacy at the Transport Layer.
- Privacy in Machine Learning
  1. Training data privacy
  2. Input/Output privacy
  3. Model privacy
  4. Pivacy-preserving Federated Learning
  5. Privacy attacks on Machine Learning Models
- Bring your own interesting topic.

Building blocks of privacy-enhancing technologies

A comparison between identity standards (OpenID, DID)

There are currently two standards for identity and attribute representation: JSON-Web-Tokens (OpenID Connect) and Decentralized Identifiers (DIDs).
Goals:

- Understand both standards/protocols and their differences.
- Show what privacy considerations exists in both concepts.
- Compare both approaches.

Figure: High-level overview of VRF.

http://cryptowiki.net/index.php?title=Verifiable_Random_Functions.

A VRF is a cryptographic concept that can be used to create publicly verifiable proofs or commitments on data in a privacy-preserving fashion. It allows a prover to calculate a function $y = f(x)$ and provide a proof π. Any verifier may use π that the y is actually the result of $f(x)$ without being able to calculate it.
Goals:

- Understand and present generalized concepts of VRF.
- Survey applications and uses of VRFs in PETs.

A Survey on Hierarchical Deterministic Key Derivation.

Hierarchical Deterministic Key Derivation (HDKD) are cryptographic key derivation schemes which are used for key blinding as well as derivation of crypto wallets.
Goals:

- Understand and present generalized concepts of HDKD.
- Present and compare existing HDKDs.
- Present use cases of HDKD.

Distributed Private Information Retrieval Schemes

PIR is used to protect user privacy when working with outsourced data.
It allows users to retrieve data from a remote store without revealing to third parties which item was retrieved.
Goals:

- Understand and present generalized concepts of PIR.
- Survey the state of the art in decentralized/distributed PIR schemes.
- Research and discuss current applications of the above.

Privacy at the Transport Layer

Privacy at the transport layer is a difficult endeavour. Especially protection of metadata is crucial but difficult. Recent proposals from leveaging TCP Fast Open and TLS 1.3. try to tackle the issue.
Goals:

- Research existing transport layer privacy approaches.
- Systematically present and compare the approaches.

Fraunhofer AISEC

Privacy in Machine Learning

Pivacy-preserving Machine Learning (1-3)

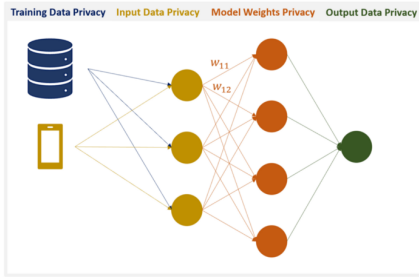

Figure: https://towardsdatascience.com/perfectly-privacy-preserving-ai-c14698f322f5.

Research and discuss latest improvements and techniques in perfectly-privacy preserving ML:

1. Training Data Privacy. (protect data creator.)
2. Input/Ouput Privacy. (protect user input/output.)
3. Model Privacy. (protect model creator.)

Pivacy-preserving Federated Learning (4)

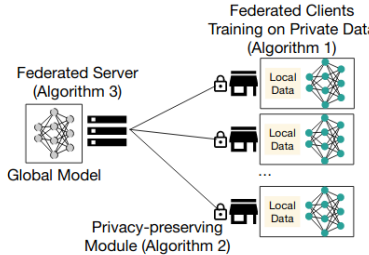

Figure: Federated learning system https://arxiv.org/pdf/1910.00962.pdf.

Research, compare and discuss existing pivacy-preserving federated learning architectures, like NVIDIAs approach[1] or the sherpa.ai approach[2] or others of your choice.

[1] https://arxiv.org/abs/1910.00962
[2] https://arxiv.org/abs/2007.00914

Privacy attacks on Machine Learning Models (5)

Fraunhofer AISEC

Research, compare and discuss privacy attacks against machine learning systems:

- Membership inference attack.[3] [4]
- Model inversion attack.[5]
- Measuring unintended neural network extraction attack.[6]
- Others of your choice

[3] https://arxiv.org/abs/1610.05820
[4] https://www.comp.nus.edu.sg/ reza/files/Shokri-SP2019.pdf
[5] https://www.cs.cmu.edu/ mfredrik/papers/fjr2015ccs.pdf
[6] https://arxiv.org/abs/1802.08232

Procedure

1. Matching and Topic assignment
   - After the matching concludes, we'll get in touch with the participants.
   - If you want to deregister
     - do so timely to avoid penalty or brace yourself for a 5.0.
   - Participants send top 3 topics via email, we'll assign the topics.
2. Familiarization phase:
   - Literature research.
   - Get an overview of your topic.
   - Create report structure.
3. Introduction to scientific writing possibly provided by chair.
4. Writing phase.
   - The first version for review must be acceptable!
   - No submission $\Rightarrow$ 5.0.
   - Violation of page limit $\Rightarrow$ 5.0.
   - No "buffering" of pages using images with little informational value or oversize.
5. Review phase.
6. Presentation.

Q&A

See first slide for contact emails.