Casey Meehan

Email : cmeehan@eng.ucsd.edu

Publications

- **Reconstruction attacks on SSL models** — Preprint
  Generative, diffusion-based data extraction attack on self-supervised vision models — *post*
- **Sentence-level Privacy for Document Embeddings** — ACL 2022
  Novel mechanism offering pure local DP at the sentence level for LLM document embeddings
- **Privacy of Generalized Shuffling** — ICLR 2022
  Formalizing a non-DP privacy notion offered by general shuffling distributions
- **Location Trace Privacy Under Conditional Priors** — AISTATS, 2021
  How to sanitize a sequence of highly correlated locations from a single user — *blog post*
- **A Non-Parametric Test to Detect Data-Copying in Generative Models** — AISTATS, 2020
  Exploring what constitutes 'overfitting' in generative models and how to detect it — *blog post*
- **Privacy Amplification by Subsampling in the Time Domain** — AISTATS, 2022
  Time-domain subsampling benefits the privacy/utility tradeoff for temporal aggregate data

Education

- **University of California, San Diego** — La Jolla, CA
  PHD candidate ***currently defending*** *studying machine learning privacy & methods* — *Fifth Year*
  - **Reconstruction attacks on SSL models** Demonstrated that self-supervised models memorize their training images by implementing a generative diffusion-based data extraction attack that leverages the SSL model to reconstruct select training samples.
  - **Sentence-level local privacy** Proposed the new, strong privacy definition of Sentence DP. Developed Tukey median based mechanism for generating sentence-private embeddings of documents.
  - **Non-Uniform Shuffling for Local Privacy:** Formalized how shuffling of private data prevents inferential threats e.g. correlation attacks. Proposed novel non-uniform shuffling mechanism that blocks such attacks while enabling trend-learning not available to uniform shuffling.
  - **Local Privacy for Location Traces:** Local privacy framework for sequences of highly dependent data, accentuating the balance between utility and realistic dependence. Developed SDP for optimizing covariance of added noise to thwart inference of any Gaussain process adversary.
  - **Nonparametric Hypothesis Test for Evaluating Generative Models:** Developed novel hypothesis testing framework for evaluating the generalization of generative models along with an efficient test statistic. Results are promising for KDEs, VAEs, and GANs.
  - **Organizer for NeurIPS privacy workshops 2019/20/21** Helped coordinate and AC multiple of NeurIPS' privacy workshops, which has been a fantastic opportunity to connect and engage with the ML privacy community on a personal level.
- **Harvard University** — Cambridge, MA
  M.S. Computational Science & Engineering (Applied Math & CS) — *Aug 2017 – May 2018*
- **Brown University** — Providence, RI
  Bachelor of Science in Electrical Eng. & Signal Processing — *Aug. 2011 – May 2015*
  - **Brown Space Engineering** lead a group of undergraduate engineers in designing/launching Brown's first satellite

Experience

- **Facebook AI Research** — San Francisco, CA
  Research intern advised by Chuan Guo — *Summer 2022*
  - **ML Privacy Risks** Investigated data reconstruction attacks on large ML models
- **Tumult Labs**
  Research intern advised by Ashwin Machanavajjhala — *Spring 2022*
  - **Privacy for large query workloads** Developed novel adaptive privacy mechanism for large scale application
- **Other things:** surfing, cooking, short fiction