

# Matthieu Meeus

[matthieumeeus@gmail.com](mailto:matthieumeeus@gmail.com) | [LinkedIn](#) | [Google Scholar](#) | [Personal website](#)

## Education

---

### **Imperial College, London – PhD student in Privacy of AI**

*Oct '22-Sept '26 – London, UK*

- Part of the [AI Security and Privacy Lab](#) under supervision of Prof Yves-Alexandre de Montjoye.
- Since the start of my PhD, I contributed to 10+ papers published/under review, including at ICML, Neurips, Nature Communications and Usenix Security. Details see *Publications*.
- Thesis (expected graduation September '26): “Auditing and Understanding Memorization in Large Language Models”.
- Main side-project: fine-tuning LLaMA-2/3 to Dutch (my native language) supported by a grant from the [Flemish Super Computer](#) for 40k GPU hours (released [here](#)).

### **Harvard University – M.Sc. in Computational Science & Engineering**

*Class of '20 – MA, USA*

- Mathematical programming ranging from numerical methods to data science, optimization, computational systems development – GPA 3.85/4.0.
- Teaching Fellow for the core graduate-level course in Advanced Numerical Methods during Fall 2020.
- Fellowship of \$100,000 – Belgian American Educational Foundation – Awarded to only 7 students.

### **University of California, Berkeley – M. Eng. in Mechanical Engineering**

*Class of '19 – CA, USA*

- One-year master with focus on Energy Technology – GPA 3.98/4.
- Fung Excellence Scholarship – UC Berkeley Fung Institute for Engineering Leadership.

### **University of Leuven – B.Sc. in Engineering – major Mechanical Engineering**

*Class of '18 – BE*

- Best of my class after the first year out of more than 500 – Received Best Student Award.

## Work Experience

---

### **Research Scientist Intern at Meta**

*March-August '26 – NYC, USA*

- Part of the Central Applied Science team, which applies privacy/memorization expertise to Meta products and often publishes insights as open-source contributions.
- Working on training data extraction from models trained on code.

### **Intern at Microsoft Research**

*May-August '24 – Cambridge, UK*

- Part of the privacy-preserving machine learning research group, working together with Lukas Wutschitz, Robert Sim and Reza Shokri.
- Researching privacy auditing of synthetic text data generated by LLMs, [work published at ICML 2025](#).

### **Senior Data Scientist at McKinsey & Company**

*March '21-July '22 – NYC, USA*

- As part of the People Analytics and Measurement team, I built machine learning models to enhance talent management and organizational effectiveness at McKinsey.
- Leveraged NLP models to improve the matching between applicants and job openings: (i) extracted features from free text resumes to improve the existing candidate ranking model by 10%; (ii) built a ‘job-matching’ algorithm to detect similar jobs across McKinsey (globally deployed on their website); (iii) built a ‘resume-to-job’ algorithm, which, given free-text resumes, recommends the top jobs based on your qualifications.
- Built a search engine for transcripts of expert interviews, based on language embeddings and named entity recognition.
- Many ad-hoc projects including NLP for McKinsey projects, prediction of attrition, the composition of ideal teams.

### **Energy Optimization Intern at Tesla Inc.**

*May-August '20 – CA, USA*

- Software development for Tesla’s Energy Optimization Team, working on optimal battery discharge strategies in Python.
- Simulated the performance of a new solar energy forecasting method (now implemented).
- Built the first parallelized performance tracking algorithm for one of the team’s products using Kubernetes and AWS.

## ***Publications***

---

- [1] **Meeus, M.**, Jain, S., Rei, M., & de Montjoye, Y. A. Did the neurons read your book? Document-level membership inference for large language models (USENIX Security 2024).
- [2] **Meeus\*, M.**, Shilov\*, I., Faysse, M., & de Montjoye, Y. A. Copyright Traps for Large Language Models (ICML 2024).
- [3] **Meeus, M.**, Shilov, I., Jain, S., Faysse, M., Rei, M., & de Montjoye, Y. A. SoK: Membership Inference Attacks on LLMs are Rushing Nowhere (and How to Fix It) (SaTML 2025, Best Paper Award).
- [4] **Meeus, M.**, Wutschitz, L., Zanella-Beguelin, S., Tople, S. & Shokri R. The Canary's Echo: Auditing Privacy Risks of LLM-Generated Synthetic Text (ICML 2025).
- [5] Shilov\*, I., **Meeus\*, M.**, & de Montjoye, Y. A. The Mosaic Memory of Large Language Models (Nature Communications 2026).
- [6] Hayes, J., Shumailov, I., Choquette-Choo, C. A., Jagielski, M., Kaissis, G., Lee, K., ... & Cooper, A. F. Exploring the limits of strong membership inference attacks on large language models (Neurips 2025).
- [7] **Meeus, M.**, Shilov, I., Kaissis, G., & de Montjoye, Y. A. Counterfactual Influence as a Distributional Quantity. [ICML2025 workshop](#).
- [8] **Meeus\*, M.**, Rathé\*, A., Remy, F., Delobelle, P., Decorte, J. & Demeester, T. (2024). ChocoLlama: Lessons Learned From Teaching Llamas Dutch. [ArXiv preprint](#).
- [9] **Meeus\*, M.**, Guepin\*, F., Crețu, A. M., & de Montjoye, Y. A. (2023, September). Achilles' heels: vulnerable record identification in synthetic data publishing (ESORICS 2023).
- [10] Guépin\*, F., **Meeus\*, M.**, Crețu, A. M., & de Montjoye, Y. A. (2023, September). Synthetic is all you need: removing the auxiliary data assumption for membership inference attacks against synthetic data (DPM Workshop at ESORICS 2023).
- [11] **Meeus, M.**, Jain, S., & de Montjoye, Y. A. (2023). Concerns about using a digital mask to safeguard patient privacy (Nature Medicine 2023).
- [12] Guépin, F., Krčo, N., **Meeus, M.**, & de Montjoye, Y. A. (2024). Lost in the Averages: A New Specific Setup to Evaluate Membership Inference Attacks Against Machine Learning Models (DPM Workshop at ESORICS 2023).
- [13] Yang, X., Stevanoski, B., **Meeus, M.**, & de Montjoye, Y. A. (2025). Checkpoint-GCG: Auditing and Attacking Fine-Tuning-Based Prompt Injection Defenses. [ArXiv preprint](#).
- [14] Krčo\*, N., Yao\*, Z., **Meeus\*, M.**, & de Montjoye, Y. A. (2026). RAT-Bench: A Comprehensive Benchmark for Text Anonymization. [ArXiv preprint](#).