Dmitrii Usynin

Trustworthy machine learning in medicine at Imperial College London and Technical University of Munich

Technische Universität München
Einsteinstr. 25, DE-81675 Munich, Germany
du216@ic.ac.uk | www.dmitrii.usyn.in/ | dimasquest | dusynin

Current Positions

2020-pres. **PhD Student in Privacy-Preserving Machine Learning**, Department of Computing, Imperial College London

2020-pres. **PhD Student in Trustworthy AI for Medicine**, Department of Diagnostic and Interventional Radiology, Technical University of Munich

2020-pres. **Research Assistant**, Institute for Artificial Intelligence in Healthcare, Technical University of Munich

Education

Imperial College London *Upper Second Class Honors (2.1)*

MEng Computing *2016 - 2020*

- Distinguished MEng Project: "*Privacy-Preserving Machine Learning in a Medical Domain*"
- Computing Department Representative
- Departmental award for excellence in project work
- Award for an outstanding computing topics project (Quantum logic)
- Award for the best presentation (Quantum logic)

Professional Experience

Microsoft Research *HuggingFace, PyTorch, opacus, Azure, AML*

Machine Learning Researcher *June 2024 – Aug 2024*

- Worked on memorisation and factuality in LLMs for radiology under differentially private training
- Developed a framework for differentially private model training and clinical evaluation of instruction-tuned LLMs

Brave Research *FLSim, PyTorch, Flower, Lambda*

Machine Learning Researcher *May 2023 – Aug 2023*

- Worked on a framework for efficient non-IID Federated Learning for data- and resource-constrained devices
- Implemented mechanisms for principled data and client selection in heterogeneous learning settings

CreatorFund *Python, Trello, Airtable*

Investment Partner *Oct 2022 – Oct 2024*

- Responsible for investment decisions
- Early-stage company sourcing and due diligence
- Managing key European relationships (QuantumDiamonds)

OpenMined *PySyft, PyTorch, scikit-learn*

Privacy Researcher *Jul 2020 – Oct 2024*

- Assistant lead of OpenMined Reseach (2021-2022)
- Researcher in Federated Learning and Differential Privacy for Healthcare (PriMIA)

Dynamify *Angular(Typescript), Groovy, HQL*

Software Engineer *Jul 2020 – Sep 2020*

- Developed a point-of-sale (POS) application for remote orders and delivery
- Worked on Stripe integration for POS

HSBC UK *Java, SQLServer, AngularJS, Python*

Client Insight Business Analyst *Apr 2019 – Sep 2019*

- Developed the post-trade performance, email analytics and operational service model dashboards
- Developed tools for client communication analytics

Credit Suisse *Flask(Python), Java, React Native*

Technical Analyst *Jun 2018 – Aug 2018*

- Developed an internal OA portal
- Created a novel online applicant assessment portal (nominated for ”Innovation of the year”)

Publications

* indicates shared first authorship Google Scholar

Highlights

[1] Dmitrii Usynin, Alexander Ziller, Marcus Makowski, Rickmer Braren, Daniel Rueckert, Ben Glocker, Georgios Kaissis, and Jonathan Passerat-Palmbach. “Adversarial interference and its mitigations in privacy-preserving collaborative machine learning”. In: *Nature Machine Intelligence* 3.9 (2021), pp. 749–758.

[2] Dmitrii Usynin, Daniel Rueckert, and Georgios Kaissis. “Beyond Gradients: Exploiting Adversarial Priors in Model Inversion Attacks”. In: *ACM Trans. Priv. Secur.* 26.3 (June 2023). ISSN: 2471-2566. DOI: 10.1145/3592800. URL: https://doi.org/10.1145/3592800.

[3] Dmitrii Usynin, Daniel Rueckert, Jonathan Passerat-Palmbach, and Georgios Kaissis. “Zen and the art of model adaptation: Low-utility-cost attack mitigations in collaborative machine learning”. In: *Proceedings on Privacy Enhancing Technologies* 2022.1 (2022), pp. 274–290.

[4] Georgios Kaissis, Alexander Ziller, Jonathan Passerat-Palmbach, Théo Ryffel, Dmitrii Usynin, Andrew Trask, Ionésio Lima, Jason Mancuso, Friederike Jungmann, Marc-Matthias Steinborn, et al. “End-to-end privacy preserving deep learning on multi-institutional medical imaging”. In: *Nature Machine Intelligence* 3.6 (2021), pp. 473–484.

2024

[1] Dmitrii Usynin, Moritz Knolle, and Giorgios Kaissis. “Memorisation in machine learning: A survey of results”. In: *Transactions on Machine Learning Research* (2024).

[2] Dmitrii Usynin, Daniel Rueckert, and Giorgios Kaissis. “Incentivising the federation: gradient-based metrics for data selection and valuation in private decentralised training”. In: *European Interdisciplinary Cybersecurity Conference* (2024).

[3] Jack Fitzsimons, Agustín Freitas Pasqualini, Robert Pisarczyk, and Dmitrii Usynin. “Naturally Private Recommendations with Determinantal Point Processes”. In: *Theory and Practice of Differential Privacy Workshop* (2024).

[4] Tamara T Mueller, Dmitrii Usynin, Johannes C Paetzold, Rickmer Braren, Daniel Rueckert, and Georgios Kaissis. “Differentially Private Guarantees for Analytics and Machine Learning on Graphs: A Survey of Results”. In: *Journal of Privacy and Confidentiality* 14.1 (2024).

[5] Felix Hsieh, Huy H Nguyen, AprilPyone MaungMaung, Dmitrii Usynin, and Isao Echizen. “Mitigating Backdoor Attacks using Activation-Guided Model Editing”. In: *arXiv preprint arXiv:2407.07662* (2024).

2023

[1] Tomas Chobola, Dmitrii Usynin, and Georgios Kaissis. “Membership inference attacks against semantic segmentation models”. In: *Proceedings of the 16th ACM Workshop on Artificial Intelligence and Security*. 2023, pp. 43–53.

2022

[1] Dmitrii Usynin, Helena Klause, Johannes C Paetzold, Daniel Rueckert, and Georgios Kaissis. “Can collaborative learning be private, robust and scalable?” In: *International Workshop on Distributed, Collaborative, and Federated Learning, Workshop on Affordable Healthcare and AI for Resource Diverse Global Health*. Springer. 2022, pp. 37–46.

[2] Dmitrii Usynin, Alexander Ziller, Daniel Rueckert, Jonathan Passerat-Palmbach, and Georgios Kaissis. “Distributed Machine Learning and the Semblance of Trust”. In: *The Third AAAI Workshop on Privacy-Preserving Artificial Intelligence* (2022).

[3] Tamara T Mueller, Johannes C Paetzold, Chinmay Prabhakar, Dmitrii Usynin, Daniel Rueckert, and Georgios Kaissis. “Differentially Private Graph Neural Networks for Whole-Graph Classification”. In: *IEEE Transactions on Pattern Analysis and Machine Intelligence* (2022).

[4] Tamara T Mueller, Stefan Kolek, Friederike Jungmann, Alexander Ziller, Dmitrii Usynin, Moritz Knolle, Daniel Rueckert, and Georgios Kaissis. “How Do Input Attributes Impact the Privacy Loss in Differential Privacy?” In: *The Forth AAAI Workshop on Privacy-Preserving Artificial Intelligence* (2023).

2021

[1] Dmitrii Usynin, Alexander Ziller, Moritz Knolle, Daniel Rueckert, and Georgios Kaissis. “An automatic differentiation system for the age of differential privacy”. In: *NeurIPS PriML Workshop* (2021).

[2] Georgios Kaissis, Moritz Knolle, Friederike Jungmann, Alexander Ziller, Dmitrii Usynin, and Daniel Rueckert. “A unified interpretation of the Gaussian mechanism for differential privacy through the sensitivity index”. In: *Journal of Privacy and Confidentiality* (2021).

[3] Moritz Knolle, Alexander Ziller, Dmitrii Usynin, Rickmer Braren, Marcus R Makowski, Daniel Rueckert, and Georgios Kaissis. “Differentially private training of neural networks with Langevin dynamics forcalibrated predictive uncertainty”. In: *ICML TPDP Workshop* (2021).

[4] *Tamara T Mueller, *Alexander Ziller, *Dmitrii Usynin, Moritz Knolle, Friederike Jungmann, Daniel Rueckert, and Georgios Kaissis. “Partial sensitivity analysis in differential privacy”. In: *arXiv preprint arXiv:2109.10582* (2021).

[5] Alexander Ziller, Dmitrii Usynin, Moritz Knolle, Kerstin Hammernik, Daniel Rueckert, and Georgios Kaissis. “Complex-valued deep learning with differential privacy”. In: *arXiv preprint arXiv:2110.03478* (2021).

[6] *Alexander Ziller, *Dmitrii Usynin, *Nicolas Remerscheid, Moritz Knolle, Marcus Makowski, Rickmer Braren, Daniel Rueckert, and Georgios Kaissis. “Differentially private federated deep learning for multi-site medical image segmentation”. In: *arXiv preprint arXiv:2107.02586* (2021).

[7] *Alexander Ziller, *Dmitrii Usynin, Rickmer Braren, Marcus Makowski, Daniel Rueckert, and Georgios Kaissis. “Medical imaging deep learning with differential privacy”. In: *Scientific Reports* 11.1 (2021), pp. 1–8.

[8] Alexander Ziller, Dmitrii Usynin, Moritz Knolle, Kritika Prakash, Andrew Trask, Rickmer Braren, Marcus Makowski, Daniel Rueckert, and Georgios Kaissis. “Sensitivity analysis in differentially private machine learning using hybrid automatic differentiation”. In: *ICML TPDP Workshop* (2021).

2020

[1] Alexander Ziller, Jonathan Passerat-Palmbach, Théo Ryffel, Dmitrii Usynin, Andrew Trask, Ionésio Da Lima Costa Junior, Jason Mancuso, Marcus Makowski, Daniel Rueckert, Rickmer Braren, et al. “Privacy-preserving medical image analysis”. In: *MedNeurIPS Workshop* (2020).

Awards, Scholarships & Grants

2024	**Privacy Enhancing Technologies for Public Health Grant**, data.org & OpenDP	*Grant*
2024	**Differential privacy workshop @Oxford University, invited participant**, Oblivious	*Travel stipend*
2024-2025	**Foresight Institute Fellowship**, The Foresight Institute	*Fellowship*
2023	**Intelligent Cooperation Workshop Scholarship**, The Foresight Institute	*Travel stipend*
2023	**PPAI-2023 Scholarship**, The 4th AAAI Workshop on Privacy-Preserving AI	*Travel stipend*
2020-2024	**Joint Academy of Doctoral Studies Scholarship**, Technical University of Munich	*Scholarship*
2020	**Corporate Partnership Award for excellence in project work**, Imperial College London	*Award*
2019	**Full Colors Award by City and Guilds College Union**, Imperial College London	*Award*
2016	**Duke of Edinburgh Gold Award**, Cheltenham College	*Award*
2015-2016	**Academic Scholarship**, Cheltenham College	*Scholarship*

Supervision & Mentoring

2023-2024	**Felix Hsieh**, Institute for AI in Healthcare, Technische Universität München	*MSc (Distinction)*
2023-2024	**Olivia Ma**, Department of Computing, Imperial College London	*MSc*
2023-2024	**Benson Zhou**, Department of Computing, Imperial College London	*MEng*
2022-2023	**Borja Sanchez**, Institute for AI in Healthcare, Technische Universität München	*MSc (Distinction)*
2022-2023	**Tomas Chobola**, Institute for AI in Healthcare, Technische Universität München	*MSc (Distinction)*
2022-2023	**Xingying Chen**, Institute for AI in Healthcare, Technische Universität München	*MSc*
2021-2022	**Tudor Cebere**, École Normale Supérieure Lyon	*MSc*
2021-2022	**Adelina Ioanna Filip**, Department of Computing, Imperial College London	*MEng (Distinction)*
2021-2022	**Aahil Mehta**, Department of Computing, Imperial College London	*MEng*
2020-2021	**Lucas Eckes**, School of Life Sciences, Ecole Polytechnique Federale de Lausanne	*MSc (Distinction)*
2020-2021	**Edward Hayes**, Department of Computing, Imperial College London	*MSc (Distinction)*

Teaching

2023-2024	**Group Projects: Software Engineering**, Department of Computing, Imperial College London
2022-2024	**Privacy, Security and Robustness in Machine Learning**, Institute for AI in Healthcare, Technical University of Munich
2022-2024	**Interpretability, Explainability and Uncertainty in Machine Learning**, Institute for AI in Healthcare, Technical University of Munich
2022-2024	**Group Projects: Applied Deep Learning in Medicine**, Institute for AI in Healthcare, Technical University of Munich
2022-2023	**Fairness in Machine Learning**, Institute for AI in Healthcare, Technical University of Munich
2021-2022	**Group Projects: Applications of Artificial Intelligence**, Department of Computing, Imperial College London

Academic Service

Conferences: ICLR 2025, AAAI 2024 (PC), NeurIPS 2024, ICML 2024, ICLR 2024, PoPETS 2023-2024, NeurIPS 2023, ICML 2023, AAAI PPAI 2023 (PC), PoPETS 2022-2023, NeurIPS 2022 (Top Reviewer), ICLR 2022 (Top Reviewer)

Journals: TMLR 2024, Statistics and Computing 2024, TMI 2023-2024, Scientific Reports 2023, TMI 2022-2023

Organisation: NeurIPS PPML Tutorial (Chair) 2024, MICCAI PPML (Chair) 2021, ICHack (Sponsorship Lead) 2018-2020, GDG DevFest 2018

Technologies

Languages: Python, Java, bash, C, TypeScript, Rust, Haskell, LaTeX

WebDev: Angular(JS), React Native, Flask

Cloud computing: Azure, AML

Other: Git, PyTorch, HuggingFace, TF, opacus, scikit-learn, pandas, FLSim, Flower