

Private Machine Learning via Multi-Party Computation

Thomas Schneider

schneider@encrypto.cs.tu-darmstadt.de

Technical University of Darmstadt / ENCRYPTO Group

Darmstadt, Germany

Abstract

Private Machine Learning (PPML) allows to privately evaluate or even train machine learning models on sensitive data while simultaneously protecting the data and the model. In this keynote, I will start with some historic works (which are often ignored) in today's rapidly growing field of PPML up to very recent results. Among the several available techniques enabling PPML such as Homomorphic Encryption (HE), Trusted Execution Environments (TEEs), Differential Privacy (DP), and Multi-Party Computation (MPC), I will put most emphasis on MPC. The main message and running theme of this keynote will be that PPML is much more than just evaluating or training neural networks under encryption. I will give an overview of a large variety of research works with special focus on those by the ENCRYPTO Group at TU Darmstadt, and will conclude with an outlook on future directions in the area of PPML.

In more detail, this keynote surveys our research towards engineering practical PPML protocols that protect data and models. First of all, there is no need to design PPML protocols for too simple models such as Support Vector Machines (SVMs) or Support Vector Regression Machines (SVRs), because they can be stolen easily [26]. Complex models can be protected and evaluated using Trusted Execution Environments (TEEs), which we demonstrated for speech recognition using Intel SGX [8] and for keyword recognition using ARM TrustZone [4]. Our goal is to build tools for non-experts in cryptography to automatically generate highly optimized mixed PPML protocols from a high-level specification in a ML framework like TensorFlow. Towards this, we have built tools to automatically generate optimized mixed protocols that combine HE and different MPC protocols [9–13, 18, 25]. There is a large body of works on private inference of neural networks using MPC, starting with the now 20 years old work by Barni, Orlandi and Piva [3] and several other earlier works that date back to long before the current hype on PPML started [2, 6, 27, 28]. But there are many more classifiers that can be evaluated securely, including decision trees [1, 2, 20, 21, 29], sum product networks [31], and even transformer models [16, 24]. For private training, there is a large body of works on private clustering [7, 17, 19], and private federated learning [5, 15, 22, 23, 30]. Finally, we have shown how to privately attest properties of ML training data to ensure fairness and non-discrimination [14]. The keynote concludes with open challenges in PPML: scalability, performance, energy efficiency, and usability.

CCS Concepts

• **Security and privacy** → **Privacy-preserving protocols**; • **Computing methodologies** → *Classification and regression trees; Neural networks.*

Keywords

privacy-preserving machine learning; tools for secure computation; optimization

ACM Reference Format:

Thomas Schneider. 2026. Private Machine Learning via Multi-Party Computation. In *Proceedings of the 12th ACM International Workshop on Security and Privacy Analytics (IWSPA '26)*, June 23–25, 2026, Frankfurt am Main, Germany. ACM, New York, NY, USA, 2 pages. <https://doi.org/10.1145/3806007.3810964>

Biography

Thomas Schneider is full professor for Cryptography and Privacy Engineering in the Department of Computer Science at the Technical University of Darmstadt. Before, he was independent research group leader at TU Darmstadt (2012-2018), did a PhD in IT Security at Ruhr-University Bochum (2008-2011), and wrote his Master thesis during a research internship at Alcatel-Lucent Bell Labs, NJ, USA (2007).



His research interests include privacy-enhancing technologies, cryptographic protocols, applied cryptography, and computer security. He heads the Cryptography and Privacy Engineering Group (ENCRYPTO), whose mission is to demonstrate that privacy can be efficiently protected in real-world applications. For this, his group combines applied cryptography and algorithm engineering to build protocols and tools for protecting sensitive data and algorithms.

For his research, especially in the areas of multi-party computation, private set intersection, and private function evaluation, he was awarded an ERC Starting Grant 2019 and an ERC Consolidator Grant 2023, the highest-profile research funding in Europe. This keynote summarizes some of the research works that are most relevant to privacy-preserving machine learning (see <https://encrypto.de/topics/PPML>), for which the ENCRYPTO group received a research award by Intel Corporation in 2019, and project funding via the Intel/Avast/VMware Private AI Collaborative Research Institute for Project “Engineering Private AI Systems (EPAI)” in 2021-2023.



This work is licensed under a Creative Commons Attribution-NonCommercial-NoDerivatives 4.0 International License.

IWSPA '26, Frankfurt am Main, Germany

© 2026 Copyright held by the owner/author(s).

ACM ISBN 979-8-4007-2609-5/2026/06

<https://doi.org/10.1145/3806007.3810964>

Acknowledgments

The research of the ENCRYPTO group underlying this keynote was generously funded by the European Research Council (ERC) under the European Union's research and innovation programs Horizon Europe (PRIVTOOLS/101124778) and Horizon 2020 (PSOTI/850990). It was co-funded by the Deutsche Forschungsgemeinschaft (DFG) within SFB 1119 CROSSING/236615297 and GRK 2050 Privacy & Trust/251805230, and by the German Federal Ministry of Research, Technology and Space (BMFT) and the Hessian Ministry of Science and Research, Arts and Culture (HMWK) within the National Research Center for Applied Cybersecurity ATHENE.

References

- Mauro Barni, Pierluigi Failla, Vladimir Kolesnikov, Riccardo Lazzeretti, Ahmad-Reza Sadeghi, and Thomas Schneider. 2009. Secure Evaluation of Private Linear Branching Programs with Medical Applications. In *14. European Symposium on Research in Computer Security (ESORICS'09) (LNCS, Vol. 5789)*. Springer, 424–439. doi:10.1007/978-3-642-04444-1_26 Full version: <https://ia.cr/2009/195>.
- Mauro Barni, Pierluigi Failla, Riccardo Lazzeretti, Ahmad-Reza Sadeghi, and Thomas Schneider. 2011. Privacy-Preserving ECG Classification with Branching Programs and Neural Networks. *IEEE Transactions on Information Forensics and Security (TIFS)* 6, 2 (June 2011), 452–468. doi:10.1109/TIFS.2011.2108650
- Mauro Barni, Claudio Orlandi, and Alessandro Piva. 2006. A Privacy-Preserving Protocol for Neural-Network-based Computation. In *8. Workshop on Multimedia and Security (MMSEC'06)*. ACM, 146–151.
- Sebastian P. Bayerl, Tommaso Frassetto, Patrick Jauernig, Korbinian Riedhammer, Ahmad-Reza Sadeghi, Thomas Schneider, Emmanuel Stempf, and Christian Weinert. 2020. Offline Model Guard: Secure and Private ML on Mobile Devices. In *23. Design, Automation & Test in Europe Conference & Exhibition (DATE'20)*. IEEE. doi:10.23919/DATE48585.2020.9116560 Online: <https://arxiv.org/abs/2007.02351>.
- Yaniv Ben-Itzhak, Helen Möllering, Benny Pinkas, Thomas Schneider, Ajith Suresh, Oleksandr Tkachenko, Shay Vargaftik, Christian Weinert, Hossein Yalame, and Avishay Yanai. 2024. ScionFL: Secure Quantized Aggregation for Federated Learning. In *2. IEEE Conference on Secure and Trustworthy Machine Learning (SaTML'24)*. IEEE, Toronto, Canada, 490–511. doi:10.1109/SaTML59370.2024.00031
- Runner-up Distinguished Paper Award.** Online: <https://ia.cr/2023/652>.
- Fabian Boemer, Rosario Cammarota, Daniel Demmler, Thomas Schneider, and Hossein Yalame. 2020. MP2ML: A Mixed-Protocol Machine Learning Framework for Private Inference. In *15. International Conference on Availability, Reliability and Security (ARES'20)*. ACM, 14:1–14:10. doi:10.1145/3407023.3407045 Full version: <https://ia.cr/2020/721>. Code: <https://github.com/IntelAI/he-transformer>.
- Beza Bozdemir, Sébastien Canard, Orhan Ermiş, Helen Möllering, Melek Önen, and Thomas Schneider. 2021. Privacy-preserving Density-Based Clustering. In *16. ACM ASIA Conference on Computer and Communications Security (ASIACCS'21)*. ACM, Virtual Event, 658–671. doi:10.1145/3433210.3453104 Online: <https://ia.cr/2021/612>. Code: <https://crypto.de/code/ppDBSCAN>.
- Ferdinand Brasser, Tommaso Frassetto, Korbinian Riedhammer, Ahmad-Reza Sadeghi, Thomas Schneider, and Christian Weinert. 2018. VoiceGuard: Secure and Private Speech Processing. In *19. Conference of the International Speech Communication Association (INTERSPEECH'18)*. International Speech Communication Association (ISCA), Hyderabad, India, 1303–1307. doi:10.21437/Interspeech.2018-2032
- Lennart Braun, Daniel Demmler, Thomas Schneider, and Oleksandr Tkachenko. 2022. MOTION – A Framework for Mixed-Protocol Multi-Party Computation. *ACM Transactions on Privacy and Security (TOPS)* 25, 2 (2022). doi:10.1145/3490390 Online: <https://ia.cr/2020/1137>. Code: <https://crypto.de/code/MOTION>.
- Lennart Braun, Moritz Huppert, Nora Khayata, Thomas Schneider, and Oleksandr Tkachenko. 2023. FUSE – Flexible File Format and Intermediate Representation for Secure Multi-Party Computation. In *18. ACM ASIA Conference on Computer and Communications Security (ASIACCS'23)*. ACM. doi:10.1145/3579856.3590340 Full version: <https://ia.cr/2023/563>. Code: <https://crypto.de/code/FUSE>.
- Niklas Büscher, Daniel Demmler, Stefan Katzenbeisser, David Kretzmer, and Thomas Schneider. 2018. HyCC: Compilation of Hybrid Protocols for Practical Secure Computation. In *25. ACM Conference on Computer and Communications Security (CCS'18)*. ACM, Toronto, Canada, 847–861. doi:10.1145/3243734.3243786 Code: <https://gitlab.com/securityengineering/HyCC>.
- Daniel Demmler, Thomas Schneider, and Michael Zohner. 2015. ABY – A Framework for Efficient Mixed-Protocol Secure Two-Party Computation. In *22. Network and Distributed System Security Symposium (NDSS'15)*. Internet Society. Code: <https://crypto.de/code/ABY>.
- Henri Dohmen, Robin William Hundt, Nora Khayata, and Thomas Schneider. 2025. SEEC: Memory Safety Meets Efficiency in Secure Two-Party Computation. In *20. ACM ASIA Conference on Computer and Communications Security (ASIACCS'25)*. ACM, Ha Noi, Vietnam, 118–135. doi:10.1145/3708821.3736224 Full version: <https://ia.cr/2025/930>. Code: <https://crypto.de/code/SEEC>.
- Vasishth Duddu, Anudeep Das, Nora Khayata, Hossein Yalame, Thomas Schneider, and N. Asokan. 2024. Attesting Distributional Properties of Training Data for Machine Learning. In *29. European Symposium on Research in Computer Security (ESORICS'24) (LNCS, Vol. 14982)*. Springer, Bydgoszcz, Poland, 3–23. doi:10.1007/978-3-031-70879-4_1 Full version: <https://arxiv.org/abs/2308.09552>.
- Till Gehlhar, Felix Marx, Thomas Schneider, Ajith Suresh, Tobias Wehrle, and Hossein Yalame. 2023. SafeFL: MPC-Friendly Framework for Private and Robust Federated Learning. In *6. Deep Learning Security and Privacy Workshop (DLSP'23)*. IEEE. doi:10.1109/SPW59333.2023.00012 Full version: <https://ia.cr/2023/555>.
- Meng Hao, Hongwei Li, Hanxiao Chen, Pengzhi Xing, Guowen Xu, and Tianwei Zhang. 2022. Iron: Private Inference on Transformers. In *36. International Conference on Neural Information Processing Systems (NeurIPS'22)*.
- Aditya Hegde, Helen Möllering, Thomas Schneider, and Hossein Yalame. 2021. SoK: Efficient Privacy-Preserving Clustering. *Proceedings on Privacy Enhancing Technologies (PoPETs)* 2021, 4 (July 2021), 225–248. doi:10.2478/popets-2021-0068 Online: <https://ia.cr/2021/809>. Code: https://crypto.de/code/SoK_ppClustering.
- Wilko Henecka, Stefan Kögl, Ahmad-Reza Sadeghi, Thomas Schneider, and Immo Wehrenberg. 2010. TASTY: Tool for Automating Secure Two-party computations. In *17. ACM Conference on Computer and Communications Security (CCS'10)*. ACM, Chicago, IL, USA, 451–462. doi:10.1145/1866307.1866358 Full version: <https://ia.cr/2010/365>. Code: <https://crypto.de/code/TASTY>.
- Hannah Keller, Helen Möllering, Thomas Schneider, and Hossein Yalame. 2021. Balancing Quality and Efficiency in Private Clustering with Affinity Propagation. In *18. International Conference on Security and Cryptography (SECRYPT'21)*. SciTePress, Virtual Event, 173–184. doi:10.5220/0010547801730184 Full version: <https://ia.cr/2021/825>. Code: <https://crypto.de/code/ppAffinityPropagation>.
- Ágnes Kiss, Masoud Naderpour, Jian Liu, N. Asokan, and Thomas Schneider. 2019. SoK: Modular and Efficient Private Decision Tree Evaluation. *Proceedings on Privacy Enhancing Technologies (PoPETs)* 2019, 2 (April 2019), 187–208. doi:10.2478/popets-2019-0026 Full version: <https://ia.cr/2018/1099>. Code: <https://crypto.de/code/PDTE>.
- Louis Kruger, Somesh Jha, Eu-Jin Goh, and Dan Boneh. 2006. Secure Function Evaluation with Ordered Binary Decision Diagrams. In *CCS*. ACM, 410–420.
- Felix Marx, Thomas Schneider, Ajith Suresh, Tobias Wehrle, Christian Weinert, and Hossein Yalame. 2026. WW-FL: Secure and Private Large-Scale Federated Learning. *IACR Transactions on Cryptographic Hardware and Embedded Systems (CHES)* 2026, 1 (January 16, 2026), 185–224. doi:10.46586/tches.v2026.i1.185-224 Full version: <https://arxiv.org/abs/2302.09904>.
- Thien Duc Nguyen, Phillip Rieger, Huili Chen, Hossein Yalame, Helen Möllering, Hossein Fereidooni, Samuel Marchal, Markus Miettinen, Azalia Mirhoseini, Shaza Zeitouni, Farinaz Koushanfar, Ahmad-Reza Sadeghi, and Thomas Schneider. 2022. FLAME: Taming Backdoors in Federated Learning. In *31. USENIX Security Symposium (USENIX Security'22)*. USENIX. Online: <https://ia.cr/2021/025>.
- Qi Pang, Jinhao Zhu, Helen Möllering, Wenting Zheng, and Thomas Schneider. 2024. BOLT: Privacy-Preserving, Accurate and Efficient Inference for Transformers. In *45. IEEE Symposium on Security and Privacy (IEEE S&P'24)*. IEEE, 4753–4771. doi:10.1109/SP54263.2024.00130 Online: <https://ia.cr/2023/1893>.
- Arpita Patra, Joachim Schmidt, Thomas Schneider, Ajith Suresh, and Hossein Yalame. 2026. SynCirc: Efficient Synthesis of Depth-Optimized Circuits from High-Level Languages. *IEEE Transactions on Computers (TC)* (March 16, 2026), 1–12. doi:10.1109/TC.2026.3673169 Online: <https://ia.cr/2026/561>. Code: <https://crypto.de/code/SynCirc>.
- Robert Nikolai Reith, Thomas Schneider, and Oleksandr Tkachenko. 2019. Efficiently Stealing your Machine Learning Models. In *18. Workshop on Privacy in the Electronic Society (WPES'19)*. ACM, 198–210. doi:10.1145/3338498.3358646
- M. Sadegh Riazi, Christian Weinert, Oleksandr Tkachenko, Ebrahim M. Songhori, Thomas Schneider, and Farinaz Koushanfar. 2018. Chameleon: A Hybrid Secure Computation Framework for Machine Learning Applications. In *13. ACM ASIA Conference on Computer and Communications Security (ASIACCS'18)*. ACM, 707–721. doi:10.1145/3196494.3196522 Preliminary version: <https://ia.cr/2017/1164>.
- Ahmad-Reza Sadeghi and Thomas Schneider. 2008. Generalized Universal Circuits for Secure Evaluation of Private Functions with Application to Data Classification. In *11. International Conference on Information Security and Cryptology (ICISC'08) (LNCS, Vol. 5461)*. Springer. doi:10.1007/978-3-642-00730-9_21
- Thomas Schneider. 2008. *Practical Secure Function Evaluation*. Master's thesis. Friedrich-Alexander University Erlangen-Nürnberg, Germany. <https://crypto.de/papers/S08Thesis.pdf>
- Thomas Schneider, Ajith Suresh, and Hossein Yalame. 2023. Comments on "Privacy-Enhanced Federated Learning Against Poisoning Adversaries". *IEEE Transactions on Information Forensics and Security (TIFS)* 18 (January 20, 2023), 1407–1409. doi:10.1109/TIFS.2023.3238544
- Amos Treiber, Alejandro Molina, Christian Weinert, Thomas Schneider, and Kristian Kersting. 2020. CryptoSPN: Privacy-Preserving Sum-Product Network Inference. In *24. European Conference on Artificial Intelligence (ECAI'20)*. doi:10.3233/FALA200313 Code: <https://crypto.de/code/CryptoSPN>.