

Yang Zhang | CV

✉ zhang@cispa.de • 🏠 yangzhangalmo.github.io
last update: May 8, 2023

Employment

CISPA Helmholtz Center for Information Security <i>Tenured Faculty, equivalent to full professor</i>	Saarbrücken, Germany 5/2023 -
CISPA Helmholtz Center for Information Security <i>Tenure-Track Faculty</i>	Saarbrücken, Germany 2/2020 - 4/2023
CISPA Helmholtz Center for Information Security <i>Research Group Leader</i>	Saarbrücken, Germany 1/2019 - 1/2020
CISPA, Saarland University <i>Postdoctoral Researcher</i>	Saarbrücken, Germany 1/2017 - 12/2018

Education

University of Luxembourg <i>Ph.D. in Computer Science, highest honor</i>	Luxembourg, Luxembourg 12/2012 - 11/2016
Shandong University <i>Master in Computer Science</i>	Jinan, China 9/2009 - 6/2012
University of Luxembourg <i>Master in Informatics, exchange student</i>	Luxembourg, Luxembourg 9/2010 - 10/2011
Shandong University <i>Bachelor in Software Engineering</i>	Jinan, China 9/2005 - 6/2009

Research Projects

- PriSyn: Repräsentative, synthetische Gesundheitsdaten mit starken Privatsphäringarantien; Funding agency: BMBF; Role: Co-PI; Time: 12/2022 - 11/2025; Budget: 310,106€
- Understanding the individual host response against Hepatitis D Virus to develop a personalized approach for the management of hepatitis D; Funding agency: Horizon Europe; Role: Co-PI; Time: 10/2022 - 9/2026; Budget: 318,375€
- The Norton Labs Graduate Fellowship;¹ Funding agency: Norton Lab; Role: Lead PI; Time: 5/2022 - 4/2023; Budget 20,000\$
- Dual-level Specification Inference; Funding agency: NGI Assure; Role: Lead PI; Time: 10/2022 - 9/2023; Budget: 34,391€
- Developing AI/Cloud-based User Data Security Technologies; Funding agency: Ministry of Science and ICT, Korea; Role: Lead PI; Time: 9/2021 - 9/2022; Budget: 7,275€
- Trustworthy Federated Data Analytics; Funding agency: Helmholtz Association; Role: Co-PI; Time: 1/2020 - 1/2023; Budget: 110,925€

¹This fellowship is given to my Ph.D. student Xinlei He.

Research Interests

- Trustworthy Machine Learning (Safety, Privacy, and Security)
- Misinformation, Hate Speech, and Memes
- Social Network Analysis

Service

- PC Member
 - 2024: IEEE S&P
 - 2023: IEEE S&P, CCS, NDSS, ICML, NeurIPS, ICLR, KDD, WWW, SaTML
 - 2022: USENIX Security, CCS, NeurIPS, ICLR, KDD, WWW, AAI, PETS, ASIACCS
 - 2021: USENIX Security, CCS, WWW, AAI, Euro S&P, PETS, ASIACCS
 - 2020: CCS, WWW, ICWSM, RAID, PETS
 - 2019: CCS, ISMB/ECCB
- Editorial Board
 - ACM Transactions on Privacy and Security (TOPS)
- PC Chair
 - International Workshop on Ethics in Computer Security (EthiCS 2022-2023)
- Organizer
 - Privacy and Security in ML Seminars
- Ph.D. Thesis Committee
 - Sinem Sav, EPFL, 2023
 - Ahmed Salem, Saarland University, 2022
 - Inken Hagestedt, Saarland University, 2021
 - Benjamin Zhao, University of New South Wales, 2021

Awards

- Best paper award honorable mention at CCS 2022
- Busy Beaver teaching award nomination for seminar "Privacy of Machine Learning" at Saarland University (2022 Winter)
- Busy Beaver teaching award nomination for advanced lecture "Machine Learning Privacy" at Saarland University (2022 Summer)
- Busy Beaver teaching award for seminar "Privacy of Machine Learning" at Saarland University (2021 Winter)
- Distinguished reviewer award at TrustML Workshop 2020 (co-located with ICLR 2020)
- Distinguished paper award at NDSS 2019
- Best paper award at ARES 2014

Publication

I have published more than 60 peer-reviewed papers so far. My publication list can also be found at DBLP and Google Scholar; however, they may not be up to date. Note that in the domain of information security, the most prestigious conferences are IEEE S&P, CCS, USENIX Security, and NDSS.

Conference

- [1] Zeyang Sha and Zheng Li and Ning Yu and **Yang Zhang**. DE-FAKE: Detection and Attribution of Fake Images Generated by Text-to-Image Generation Models. In *ACM SIGSAC Conference on Computer and Communications Security (CCS)*. ACM, 2023.
- [2] Yiting Qu and Xinyue Shen and Xinlei He and Michael Backes and Savvas Zannettou and **Yang Zhang**. Unsafe Diffusion: On the Generation of Unsafe Images and Hateful Memes From Text-To-Image Models. In *ACM SIGSAC Conference on Computer and Communications Security (CCS)*. ACM, 2023.
- [3] Boyang Zhang and Xinlei He and Yun Shen and Tianhao Wang and **Yang Zhang**. A Plot is Worth a Thousand Words: Model Information Stealing Attacks via Scientific Plots. In *USENIX Security Symposium (USENIX Security)*. USENIX, 2023.
- [4] Wai Man Si and Michael Backes and **Yang Zhang** and Ahmed Salem. Two-in-One: A Model Hijacking Attack Against Text Generation Models. In *USENIX Security Symposium (USENIX Security)*. USENIX, 2023.
- [5] Zheng Li and Ning Yu and Ahmed Salem and Michael Backes and Mario Fritz and **Yang Zhang**. UnGANable: Defending Against GAN-based Face Manipulation. In *USENIX Security Symposium (USENIX Security)*. USENIX, 2023.
- [6] Min Chen and Zhikun Zhang and Michael Backes and Tianhao Wang and **Yang Zhang**. FACE-AUDITOR: Data Auditing in Facial Recognition Systems. In *USENIX Security Symposium (USENIX Security)*. USENIX, 2023.
- [7] Haiming Wang and Zhikun Zhang and Tianhao Wang and Shibo He and Michael Backes and Jiming Chen and **Yang Zhang**. PrivTrace: Differentially Private Trajectory Synthesis by Adaptive Markov Model. In *USENIX Security Symposium (USENIX Security)*. USENIX, 2023.
- [8] Yihan Ma and Zhikun Zhang and Ning Yu and Xinlei He and Michael Backes and Yun Shen and **Yang Zhang**. Generated Graph Detection. In *International Conference on Machine Learning (ICML)*. PMLR, 2023.
- [9] Ziqing Yang and Xinlei He and Zheng Li and Michael Backes and Mathias Humbert and Pascal Berrang and **Yang Zhang**. Data Poisoning Attacks Against Multimodal Encoders. In *International Conference on Machine Learning (ICML)*. PMLR, 2023.
- [10] Kai Mei and Zheng Li and Zhenting Wang and **Yang Zhang** and Shiqing Ma. NOTABLE: Transferable Backdoor Attacks Against Prompt-based NLP Models. In *Annual Meeting of the Association for Computational Linguistics (ACL)*. ACL, 2023.
- [11] Zeyang Sha and Xinlei He and Ning Yu and Michael Backes and **Yang Zhang**. Can't Steal? Cont-Steal! Contrastive Stealing Attacks Against Image Encoders. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*. IEEE, 2023.
- [12] Yiting Qu and Xinlei He and Shannon Pierson and Michael Backes and **Yang Zhang** and Savvas Zannettou. On the Evolution of (Hateful) Memes by Means of Multimodal Contrastive Learning. In *IEEE Symposium on Security and Privacy (S&P)*. IEEE, 2023.
- [13] Rui Wen and Zhengyu Zhao and Zhuoran Liu and Michael Backes and Tianhao Wang and **Yang Zhang**. Is Adversarial Training Really a Silver Bullet for Mitigating Data Poisoning? In *International Conference on Learning Representations (ICLR)*, 2023.
- [14] Yugeng Liu and Zheng Li and Michael Backes and Yun Shen and **Yang Zhang**. Backdoor Attacks Against Dataset Distillation. In *Network and Distributed System Security Symposium (NDSS)*. Internet Society, 2023.
- [15] Xiaojian Yuan and Kejiang Chen and Jie Zhang and Weiming Zhang and Nenghai Yu and **Yang Zhang**. Pseudo Label-Guided Model Inversion Attack via Conditional Generative Adversarial Network. In *AAAI Conference on Artificial Intelligence (AAAI)*. AAAI, 2023.

- [16] Yufei Chen and Chao Shen and Yun Shen and Cong Wang and **Yang Zhang**. Amplifying Membership Exposure via Data Poisoning. In *Annual Conference on Neural Information Processing Systems (NeurIPS)*. NeurIPS, 2022.
- [17] Wai Man Si and Michael Backes and Jeremy Blackburn and Emiliano De Cristofaro and Gianluca Stringhini and Savvas Zannettou and **Yang Zhang**. Why So Toxic? Measuring and Triggering Toxic Behavior in Open-Domain Chatbots. In *ACM SIGSAC Conference on Computer and Communications Security (CCS)*, pages 2659–2673. ACM, 2022.
- [18] Hai Huang and Zhikun Zhang and Yun Shen and Michael Backes and Qi Li and **Yang Zhang**. On the Privacy Risks of Cell-Based NAS Architectures. In *ACM SIGSAC Conference on Computer and Communications Security (CCS)*, pages 1427–1441. ACM, 2022.
- [19] Yiyong Liu and Zhengyu Zhao and Michael Backes and **Yang Zhang**. Membership Inference Attacks by Exploiting Loss Trajectory. In *ACM SIGSAC Conference on Computer and Communications Security (CCS)*, pages 2085–2098. ACM, 2022.
- [20] Zheng Li and Yiyong Liu and Xinlei He and Ning Yu and Michael Backes and **Yang Zhang**. Auditing Membership Leakages of Multi-Exit Networks. In *ACM SIGSAC Conference on Computer and Communications Security (CCS)*, pages 1917–1931. ACM, 2022.
- [21] Min Chen and Zhikun Zhang and Tianhao Wang and Michael Backes and Mathias Humbert and **Yang Zhang**. Graph Unlearning. In *ACM SIGSAC Conference on Computer and Communications Security (CCS)*, pages 499–513. ACM, 2022.
- [22] Tianshuo Cong and Xinlei He and **Yang Zhang**. SSLGuard: A Watermarking Scheme for Self-supervised Learning Pre-trained Encoders. In *ACM SIGSAC Conference on Computer and Communications Security (CCS)*, pages 579–593. ACM, 2022.
- [23] Yun Shen and Yufei Han and Zhikun Zhang and Min Chen and Ting Yu and Michael Backes and **Yang Zhang** and Gianluca Stringhini. Finding MNEMON: Reviving Memories of Node Embeddings. In *ACM SIGSAC Conference on Computer and Communications Security (CCS)*, pages 2643–2657. ACM, 2022.
- [24] Xinlei He and Hongbin Liu and Neil Zhenqiang Gong and **Yang Zhang**. Semi-Leak: Membership Inference Attacks Against Semi-supervised Learning. In *European Conference on Computer Vision (ECCV)*, pages 365–381. Springer, 2022.
- [25] Yugeng Liu and Rui Wen and Xinlei He and Ahmed Salem and Zhikun Zhang and Michael Backes and Emiliano De Cristofaro and Mario Fritz and **Yang Zhang**. ML-Doctor: Holistic Risk Assessment of Inference Attacks Against Machine Learning Models. In *USENIX Security Symposium (USENIX Security)*, pages 4525–4542. USENIX, 2022.
- [26] Yufei Chen and Chao Shen and Cong Wang and **Yang Zhang**. Teacher Model Fingerprinting Attacks Against Transfer Learning. In *USENIX Security Symposium (USENIX Security)*, pages 3593–3610. USENIX, 2022.
- [27] Zhikun Zhang and Min Chen and Michael Backes and Yun Shen and **Yang Zhang**. Inference Attacks Against Graph Neural Networks. In *USENIX Security Symposium (USENIX Security)*, pages 4543–4560. USENIX, 2022.
- [28] Xinyue Shen and Xinlei He and Michael Backes and Jeremy Blackburn and Savvas Zannettou and **Yang Zhang**. On Xing Tian and the Perseverance of Anti-China Sentiment Online. In *International Conference on Weblogs and Social Media (ICWSM)*, pages 944–955. AAAI, 2022.
- [29] Yun Shen and Xinlei He and Yufei Han and **Yang Zhang**. Model Stealing Attacks Against Inductive Graph Neural Networks. In *IEEE Symposium on Security and Privacy (S&P)*, pages 1175–1192. IEEE, 2022.
- [30] Ahmed Salem and Rui Wen and Michael Backes and Shiqing Ma and **Yang Zhang**. Dynamic Backdoor Attacks Against Machine Learning Models. In *IEEE European Symposium on Security and Privacy (Euro S&P)*, pages 703–718. IEEE, 2022.

- [31] Ahmed Salem and Michael Backes and **Yang Zhang**. Get a Model! Model Hijacking Attack Against Machine Learning Models. In *Network and Distributed System Security Symposium (NDSS)*. Internet Society, 2022.
- [32] Junhao Zhou and Yufei Chen and Chao Shen and **Yang Zhang**. Property Inference Attacks Against GANs. In *Network and Distributed System Security Symposium (NDSS)*. Internet Society, 2022.
- [33] Xinlei He and **Yang Zhang**. Quantifying and Mitigating Privacy Risks of Contrastive Learning. In *ACM SIGSAC Conference on Computer and Communications Security (CCS)*, pages 845–863. ACM, 2021.
- [34] Min Chen and Zhikun Zhang and Tianhao Wang and Michael Backes and Mathias Humbert and **Yang Zhang**. When Machine Unlearning Jeopardizes Privacy. In *ACM SIGSAC Conference on Computer and Communications Security (CCS)*, pages 896–911. ACM, 2021.
- [35] Minxing Zhang and Zhaochun Ren and Zihan Wang and Pengjie Ren and Zhumin Chen and Pengfei Hu and **Yang Zhang**. Membership Inference Attacks Against Recommender Systems. In *ACM SIGSAC Conference on Computer and Communications Security (CCS)*, pages 864–879. ACM, 2021.
- [36] Zheng Li and **Yang Zhang**. Membership Leakage in Label-Only Exposures. In *ACM SIGSAC Conference on Computer and Communications Security (CCS)*, pages 880–895. ACM, 2021.
- [37] Xiaoyi Chen and Ahmed Salem and Michael Backes and Shiqing Ma and Qingni Shen and Zhonghai Wu and **Yang Zhang**. BadNL: Backdoor Attacks Against NLP Models with Semantic-preserving Improvements. In *Annual Computer Security Applications Conference (ACSAC)*, pages 554–569. ACSAC, 2021.
- [38] Xinlei He and Jinyuan Jia and Michael Backes and Neil Zhenqiang Gong and **Yang Zhang**. Stealing Links from Graph Neural Networks. In *USENIX Security Symposium (USENIX Security)*, pages 2669–2686. USENIX, 2021.
- [39] Zhikun Zhang and Tianhao Wang and Jean Honorio and Ninghui Li and Michael Backes and Shibo He and Jiming Chen and **Yang Zhang**. PrivSyn: Differentially Private Data Synthesis. In *USENIX Security Symposium (USENIX Security)*, pages 929–946. USENIX, 2021.
- [40] Fatemeh Tahmasbi and Leonard Schild and Chen Ling and Jeremy Blackburn and Gianluca Stringhini and **Yang Zhang** and Savvas Zannettou. “Go eat a bat, Chang!”: On the Emergence of Sinophobic Behavior on Web Communities in the Face of COVID-19. In *The Web Conference (WWW)*. ACM, 2021.
- [41] Rui Wen and Yu Yu and Xiang Xie and **Yang Zhang**. LEAF: A Faster Secure Search Algorithm via Localization, Extraction, and Reconstruction. In *ACM SIGSAC Conference on Computer and Communications Security (CCS)*, pages 1219–1232. ACM, 2020.
- [42] Dingfan Chen and Ning Yu and **Yang Zhang** and Mario Fritz. GAN-Leaks: A Taxonomy of Membership Inference Attacks against Generative Models. In *ACM SIGSAC Conference on Computer and Communications Security (CCS)*, pages 343–362. ACM, 2020.
- [43] Ahmed Salem and Apratim Bhattacharya and Michael Backes and Mario Fritz and **Yang Zhang**. Updates-Leak: Data Set Inference and Reconstruction Attacks in Online Learning. In *USENIX Security Symposium (USENIX Security)*, pages 1291–1308. USENIX, 2020.
- [44] Inken Hagestedt and Mathias Humbert and Pascal Berrang and Irina Lehmann and Roland Eils and Michael Backes and **Yang Zhang**. Membership Inference Against DNA Methylation Databases. In *IEEE European Symposium on Security and Privacy (Euro S&P)*, pages 509–520. IEEE, 2020.
- [45] **Yang Zhang** and Mathias Humbert and Bartłomiej Surma and Praveen Manoharan and Jilles Vreeken and Michael Backes. Towards Plausible Graph Anonymization. In *Network and Distributed System Security Symposium (NDSS)*. Internet Society, 2020.
- [46] Jinyuan Jia and Ahmed Salem and Michael Backes and **Yang Zhang** and Neil Zhenqiang Gong. MemGuard: Defending against Black-Box Membership Inference Attacks via Adversarial Examples. In *ACM SIGSAC Conference on Computer and Communications Security (CCS)*, pages 259–274. ACM, 2019.

- [47] Zheng Li and Chengyu Hu and **Yang Zhang** and Shanqing Guo. How to Prove Your Model Belongs to You: A Blind-Watermark based Framework to Protect Intellectual Property of DNN. In *Annual Computer Security Applications Conference (ACSAC)*, pages 126–137. ACSAC, 2019.
- [48] Zhiqiang Zhong and **Yang Zhang** and Jun Pang. A Graph-Based Approach to Explore Relationship Between Hashtags and Images. In *International Conference Web Information Systems Engineering (WISE)*, pages 473–488. Springer, 2019.
- [49] Tahleen Rahman and Bartłomiej Surma and Michael Backes and **Yang Zhang**. Fairwalk: Towards Fair Graph Embedding. In *International Joint Conferences on Artificial Intelligence (IJCAI)*, pages 3289–3295. IJCAI, 2019.
- [50] **Yang Zhang**. Language in Our Time: An Empirical Analysis of Hashtags. In *The Web Conference (WWW)*, pages 2378–2389. ACM, 2019.
- [51] Ahmed Salem and **Yang Zhang** and Mathias Humbert and Pascal Berrang and Mario Fritz and Michael Backes. ML-Leaks: Model and Data Independent Membership Inference Attacks and Defenses on Machine Learning Models. In *Network and Distributed System Security Symposium (NDSS)*. Internet Society, 2019.
- [52] Inken Hagestedt and **Yang Zhang** and Mathias Humbert and Pascal Berrang and Haixu Tang and Xiaofeng Wang and Michael Backes. MBeacon: Privacy-Preserving Beacons for DNA Methylation Data. In *Network and Distributed System Security Symposium (NDSS)*. Internet Society, 2019.
- [53] Fanghua Zhao and Linan Gao and **Yang Zhang** and Zeyu Wang and Bo Wang and Shanqing Guo. You Are Where You App: An Assessment on Location Privacy of Social Applications. In *International Symposium on Software Reliability Engineering (ISSRE)*, pages 236–247. IEEE, 2018.
- [54] **Yang Zhang** and Mathias Humbert and Tahleen Rahman and Cheng-Te Li and Jun Pang and Michael Backes. Tagvisor: A Privacy Advisor for Sharing Hashtags. In *The Web Conference (WWW)*, pages 287–296. ACM, 2018.
- [55] Pascal Berrang and Mathias Humbert and **Yang Zhang** and Irina Lehmann and Roland Eils and Michael Backes. Dissecting Privacy Risks in Biomedical Data. In *IEEE European Symposium on Security and Privacy (Euro S&P)*, pages 62–76. IEEE, 2018.
- [56] Michael Backes and Mathias Humbert and Jun Pang and **Yang Zhang**. walk2friends: Inferring Social Links from Mobility Profiles. In *ACM SIGSAC Conference on Computer and Communications Security (CCS)*, pages 1943–1957. ACM, 2017.
- [57] Jun Pang and **Yang Zhang**. Quantifying Location Sociality. In *ACM Conference on Hypertext and Social Media (HT)*, pages 145–154. ACM, 2017.
- [58] Jun Pang and **Yang Zhang**. DeepCity: A Feature Learning Framework for Mining Location Check-Ins. In *International Conference on Weblogs and Social Media (ICWSM)*, pages 652–655. AAAI, 2017.
- [59] Yan Wang and Zongxu Qin and Jun Pang and **Yang Zhang** and Xin Jin. Semantic Annotation for Places in LBSN Using Graph Embedding. In *ACM International Conference on Information and Knowledge Management (CIKM)*, page 2343–2346. ACM, 2017.
- [60] **Yang Zhang** and Minyue Ni and Weili Han and Jun Pang. Does #like4like Indeed Provoke More Likes? In *International Conference on Web Intelligence (WI)*, pages 179–186. ACM, 2017.
- [61] Minyue Ni and **Yang Zhang** and Weili Han and Jun Pang. An Empirical Study on User Access Control in Online Social Networks. In *ACM Symposium on Access Control Models and Technologies (SACMAT)*, pages 12–23. ACM, 2016.
- [62] Jun Pang and Polina Zablotskaia and **Yang Zhang**. On Impact of Weather on Human Mobility in Cities. In *International Conference Web Information Systems Engineering (WISE)*, pages 247–256. Springer, 2016.
- [63] Jun Pang and **Yang Zhang**. Location Prediction: Communities Speak Louder than Friends. In *ACM Conference on Online Social Networks (COSN)*, pages 161–171. ACM, 2015.

- [64] **Yang Zhang** and Jun Pang. Distance and Friendship: A Distance-based Model for Link Prediction in Social Networks. In *Asia-Pacific Web Conference (APWeb)*, pages 55–66. Springer, 2015.
- [65] Jun Pang and **Yang Zhang**. Event Prediction with Community Leaders. In *Conference on Availability, Reliability and Security (ARES)*, pages 238–243. IEEE, 2015.
- [66] Marcos Cramer and Jun Pang and **Yang Zhang**. A Logical Approach to Restricting Access in Online Social Networks. In *ACM Symposium on Access Control Models and Technologies (SACMAT)*, pages 75–86. ACM, 2015.
- [67] Jun Pang and **Yang Zhang**. Cryptographic Protocols for Enforcing Relationship-based Access Control Policies. In *Annual IEEE Computers, Software and Applications Conference (COMPSAC)*, pages 484–493. IEEE, 2015.
- [68] Jun Pang and **Yang Zhang**. Exploring Communities for Effective Location Prediction. In *International Conference on World Wide Web (WWW)*, pages 87–88. ACM, 2015.
- [69] **Yang Zhang** and Jun Pang. Community-driven Social Influence Analysis and Applications. In *International Conference on Web Engineering (ICWE)*. Springer, 2015.
- [70] Jun Pang and **Yang Zhang**. A New Access Control Scheme for Facebook-style Social Networks. In *Conference on Availability, Reliability and Security (ARES)*, pages 1–10. IEEE, 2014.

Journal

- [71] Xiaokuan Zhang and **Yang Zhang** and Yinqian Zhang. VERITRAIN: Validating MLaaS Training Efforts via Anomaly Detection. *IEEE Transactions on Dependable and Secure Computing*, 2023.
- [72] Cheng-Te Li and Cheng Hsu and **Yang Zhang**. FairSR: Fairness-aware Sequential Recommendation through Multi-Task Learning with Preference Graph Embeddings. *ACM Transactions on Intelligent Systems and Technology*, 2022.
- [73] Xinlei He and Qingyuan Gong and Yang Chen and **Yang Zhang** and Xin Wang and Xiaoming Fu. DatingSec: Detecting Malicious Accounts in Dating Apps Using a Content-Based Attention Network. *IEEE Transactions on Dependable and Secure Computing*, 2021.
- [74] Bo-Heng Chen and Cheng-Te Li and Kun-Ta Chuang and Jun Pang and **Yang Zhang**. An Active Learning-based Approach for Location-aware Acquaintance Inference. *Knowledge and Information Systems*, 2018.
- [75] Jun Pang and **Yang Zhang**. A New Access Control Scheme for Facebook-style Social Networks. *Computers & Security*, 2015.

Teaching

- Advanced Lecture: Attacks Against Machine Learning Models (2023 Summer)
- Seminar: Data-driven Understanding of the Disinformation Epidemic (2023 Summer)
- Seminar: Privacy of Machine Learning (2022 Winter)
- Advanced Lecture: Machine Learning Privacy (2022 Summer)
- Seminar: Data-driven Understanding of the Disinformation Epidemic (2022 Summer)
- Seminar: Privacy of Machine Learning (2021 Winter)
- Advanced Lecture: Privacy Enhancing Technologies (2021 Summer)
- Seminar: Data-driven Understanding of the Disinformation Epidemic (2021 Summer)
- Seminar: Data Privacy (2020 Winter)
- Advanced Lecture: Privacy Enhancing Technologies (2020 Summer)
- Seminar: Data-driven Approaches on Understanding Disinformation (2020 Summer)
- Seminar: Data Privacy (2019 Winter)

- Advanced Lecture: Privacy Enhancing Technologies (2019 Summer)
- Seminar: Biomedical Privacy (2019 Summer)
- Seminar: Data Privacy (2018 Winter)
- Advanced Lecture: Privacy Enhancing Technologies (2018 Summer)
- Seminar: Adversarial Machine Learning (2018 Summer)

Students

Ph.D. Students.....

Zeyang Sha	3/2023 -
Junjie Chu <i>(co-supervised with Michael Backes)</i>	11/2022 -
Yixin Wu <i>(co-supervised with Michael Backes)</i>	11/2022 -
Ziqing Yang <i>(co-supervised with Michael Backes)</i>	11/2022 -
Xinyue Shen <i>(co-supervised with Michael Backes)</i>	10/2022 -
Yugeng Liu <i>(co-supervised with Michael Backes)</i>	1/2022 -
Boyang Zhang	12/2021 -
Hai Huang <i>(co-supervised with Michael Backes)</i>	11/2021 -
Yiting Qu <i>(co-supervised with Michael Backes)</i>	11/2021 -
Wai Man Si <i>(co-supervised with Michael Backes)</i>	11/2021 -
Rui Wen <i>(co-supervised with Michael Backes)</i>	10/2021 -
Yihan Ma <i>(co-supervised with Michael Backes)</i>	7/2021 -
Zheng Li	2/2021 -
Xinlei He	2/2020 -

Ph.D. Preparatory Phase.....

Yukun Jiang	11/2022 -
--------------------	------------------

Alumni

Tianshuo Cong <i>visiting Ph.D. student, Tsinghua University</i>	08/2021 - 12/2022
Jingcan Chen <i>intern, Fudan University</i>	7/2022 - 10/2022
Runjie Jin <i>intern, Fudan University</i>	7/2022 - 10/2022
Ziyan Qu <i>intern, KTH</i>	7/2022 - 10/2022
Shannon Pierson <i>visiting researcher, Saarland University, now at LSE</i>	3/2022 - 6/2022
Ahmed Salem <i>co-supervised Ph.D. student, CISPA, now at MSR</i>	2/2017 - 1/2022
Fan Zhang <i>intern, Tsinghua University, now at Penn State</i>	7/2021 - 10/2021
Bartlomiej Surma <i>co-supervised Ph.D. student, CISPA, now at Google</i>	6/2016 - 9/2021
Joann Chen <i>intern, UC Irvine</i>	7/2021 - 9/2021
Ge Han <i>visiting Ph.D. student, Shandong University</i>	10/2019 - 8/2021
Suliya <i>visiting Ph.D. student, Chinese Academy of Science, now at JD</i>	1/2020 - 1/2021
Xiaoyi Chen <i>visiting Ph.D. student, Peking University</i>	10/2019 - 10/2020
Yuhao Mao <i>intern, Zhejiang University, now at ETH</i>	6/2020 - 9/2020
Leonard Schild <i>research assistant, Saarland University, now at KU Leuven</i>	3/2017 - 7/2020
Zeyu Yang <i>visiting Ph.D. student, Zhejiang University</i>	10/2019 - 3/2020