



Dimitar I. Dimitrov, Ph.D. Candidate

✉ dimitar.iliev.dimitrov@insait.ai

✂ @dimitrov_dimy

🌐 <https://www.sri.inf.ethz.ch/people/dimitadi>

🔗 <http://scholar.google.com/citations?user=phgGRaYAAAAJ/>

🌐 <http://www.linkedin.com/in/dimitar-i-dimitrov-44803924b/>

☎ +41 76 232 35 47

Education

Nov 2020 – May 2025

📖 PhD Student, Computer Science, *ETH Zurich*, Switzerland

- Published 12 papers
- Teaching assistant (TA) of multiple courses, including Head TA of **Research Topics in Software Engineering**
- Supervised many Bachelor's and Master's students
- Part of the team that trained **BgGPT 1.0**, a SotA LLM for Bulgarian based on the Gemma 2 family of models on par with **GPT-4o** in Bulgarian [15]
- Winners of the **PET worldwide AI privacy competition** organized by the **UK-US Summit for Democracy**
- Part of the **BRAIN++ AI factory grant proposal** that was awarded €90 million by the EU for creating a GPU cluster in Bulgaria
- PhD Thesis:
 - **Title:** Testing Federated Learning Privacy Through Gradient Leakage Attacks
 - **Supervisor:** Prof. Dr. Martin Vechev

Sep 2018 – Sep 2020

📖 Master Student, Computer Science, *ETH Zurich*, Switzerland GPA: 5.38/6.00

- Master's Thesis:
 - **Title:** Provably robust adversarial examples
 - **Supervisor:** Prof. Dr. Martin Vechev
 - **Grade:** 5.50/6.00
 - **Published:** @ICLR 2022 [11]

Sep 2012 – Jun 2016

📖 Bachelors Student, Computer Science, *The University of Edinburgh*, UK GPA: 83.9/100.0 (*Top of the Class*)

- JP Morgan Prize for Best Performance in Computer Science
- Class Prize for Top Performance in Computer Science
- Exchange Student at the University of California, Irvine
- Bachelor's Thesis:
 - **Title:** Image Inpainting with Gaussian Processes [17]
 - **Supervisor:** Prof. Dr. Christopher K. I. Williams
 - **Grade:** 80.0/100.0

Education (continued)

Sep 2014 – Jun 2015

■ **Exchange Student, Computer Science**, *University of California, Irvine, USA*
GPA: 3.9/4.0

- Obtained A+ grade in 6 courses
- Took several graduate courses in Machine Learning and Computer Vision.
- Took part in the UAVForge project where we constructed autonomous quadcopter. As part of the project I researched and implemented AI algorithms for path finding and obstacle avoidance. As a result, I was selected to be the leader of the AI/Computer Vision project's team. My duties included monitoring the work of the other members of the team and organizing weekly meetings.

Sep 2004 – Jul 2012

■ **High School Student**, *Sofia High School of Mathematics, Bulgaria*
GPA: 5.89/6.00

- Participated in many national and several international math competitions
- Developed a series of projects with my fellow students, learning different programming languages and frameworks

Employment History

Sep 2016 – Jun 2018

■ **Full-time Software Developer**, *Arista Networks*, Republic of Ireland
Managed a small team of 5 people developing custom Linux drivers for DMA memory allocation and management using IOMMU. Worked on PCIe drivers for packet routing between custom routing chips and CPU. Worked on PCIe drivers synchronizing state between CPU complexes for stateful switchover support.

Jun 2015 – Sep 2015

■ **Summer Intern**, *Arista Networks*, California, USA
Worked on low-level optimizations for one of the company products. In particular, programmed routing ASIC's L3 tables to store more routing paths by optimizing how they are stored in the hardware.

Dec 2014 – Jun 2015

■ **Teaching Assistant and Web Developer**, *Avid Academy for Gifted Youth*, California, USA
As a teaching assistant for Avid Academy, I prepared high school students for AMC (American Mathematics Competitions). I also worked on the company's web-based course management system using Django.

Jun 2013 – Sep 2014

■ **Lead Android Developer**, *The City of Edinburgh Council*, UK
I, alongside several fellow students, collaborated with the Edinburgh City Council to develop a public platform for promoting the sports facilities in the city. The collaboration emerged from our hackathon project, which was awarded the best UI prize.

- Designed and built a complex Android application from scratch
- Designed complex UI and improved it based on user testing
- Made the application compatible with a large family of Android devices
- Developed extensive knowledge in Android SDK, SQLite, OpenGL, and Java multithreading libraries

Languages



Native

Bulgarian



Proficient

English



Beginner

German

Technical Skills



Expert

Python
Pytorch
Pandas
Numpy
sklearn
Javascript
Java
Latex



Advanced

JAX
TensorFlow
Android
C/C++
Matlab
C#
Linux Kernel
Django
SQL



Worked with

OpenGL
OpenCV
MIPS

Miscellaneous Experience

Awards and Achievements

- 2024 Outstanding Reviewer @ **ICML 2024**
- 2023 Winners of the **PET worldwide AI privacy competition** organized by the **UK-US Summit for Democracy**
- 2022 Outstanding Reviewer @ **ICML 2022**
- 2019 Winners of the BETH Hackton for Sustainable Blockchain Solutions @ **ETH Zurich**
- 2016 JP Morgan Prize for Best Performance in Computer Science @ **The University of Edinburgh**
 Class Prize for Top Performance in Computer Science @ **The University of Edinburgh**
- 2014 Prize for the best UI design at the Smart Data Hack organized by the Edinburgh Council @ **The University of Edinburgh**
- 2011 3rd in the Division A PUMaC Team Competition @ **Princeton**

Certifications

- 2017 GRE General Test — Verbal Reasoning **162 (Top 9%)**; Quantitative Reasoning **169 (Top 4%)**; Writing **4.0 (Top 40%)**
- 2011 IELTS Academic **7.0/9.0**

Other








- 2025 Part of the **BRAIN++ AI factory** grant proposal that was awarded €90 million by the EU for creating a GPU cluster in Bulgaria
<https://insait.ai/bulgaria-will-have-its-own-ai-factory-a-project-for-90m-eur/>
- 2024 Our Bulgarian LLM model, **BgGPT 1.0**, was featured in a Google Blogpost on the use of Google's Gemma for modeling low-resource languages
<https://developers.googleblog.com/en/building-more-inclusive-llms-using-gemma-open-models/>

Miscellaneous Experience (continued)

2023  Our work was featured in an **ACM blogpost** on Privacy of Federated Learning
<https://cacm.acm.org/blogcacm/federated-learning-how-private-is-it-really/>





Research Publications

Conference Accepted Papers

- 1 M. Drencheva, I. Petrov, M. Baader, D. I. Dimitrov, and M. Vechev, "GRAIN: Exact graph reconstruction from gradients," in *The Thirteenth International Conference on Learning Representations, ICLR*, 2025.
 URL: <https://openreview.net/forum?id=7bAjVh3CG3>.
- 2 D. I. Dimitrov, M. Baader, M. Müller, and M. Vechev, "Spear: Exact gradient inversion of batches in federated learning," in *Advances in Neural Information Processing Systems, NeurIPS*, vol. 37, 2024, pp. 106 768–106 799.  URL: https://proceedings.neurips.cc/paper_files/paper/2024/file/c13cd7feab4beb1a27981e19e2455916-Paper-Conference.pdf.
- 3 K. Garov, D. I. Dimitrov, N. Jovanović, and M. Vechev, "Hiding in plain sight: Disguising data stealing attacks in federated learning," in *The Twelfth International Conference on Learning Representations, ICLR*, 2024.  URL: <https://openreview.net/forum?id=krx5512A6G>.
- 4 I. Petrov, D. I. Dimitrov, M. Baader, M. Müller, and M. Vechev, "Dager: Exact gradient inversion for large language models," in *Advances in Neural Information Processing Systems, NeurIPS*, vol. 37, 2024, pp. 87 801–87 830.  URL: https://proceedings.neurips.cc/paper_files/paper/2024/file/9ff1577a1f8308df1ccea6b4f64a103f-Paper-Conference.pdf.
- 5 N. Jovanović, M. Balunović, D. I. Dimitrov, and M. Vechev, "Fare: Provably fair representation learning with practical certificates," in *International Conference on Machine Learning, ICML*, 2023, pp. 15 401–15 420.
- 6 J. Lokna, A. Paradis, D. I. Dimitrov, and M. Vechev, "Group and attack: Auditing differential privacy," in *Proceedings of the 2023 ACM SIGSAC Conference on Computer and Communications Security, CCS*, 2023, pp. 1905–1918.
- 7 M. Vero, M. Balunović, D. I. Dimitrov, and M. Vechev, "Tableak: Tabular data leakage in federated learning," in *International Conference on Machine Learning, ICML*, 2023, pp. 35 051–35 083.
- 8 M. Balunović, D. I. Dimitrov, N. Jovanović, and M. Vechev, "Lamp: Extracting text from gradients with language model priors," in *Advances in Neural Information Processing Systems, NeurIPS*, vol. 35, 2022, pp. 7641–7654.
- 9 M. Balunović, D. I. Dimitrov, R. Staab, and M. T. Vechev, "Bayesian framework for gradient leakage," in *The Tenth International Conference on Learning Representations, ICLR*, 2022.  URL: <https://openreview.net/forum?id=f2lrIbGx3x7>.
- 10 D. I. Dimitrov, M. Balunović, N. Konstantinov, and M. Vechev, "Data leakage in federated averaging," in *Transactions on Machine Learning Research, TMLR*, 2022.  URL: <https://openreview.net/forum?id=e7A0B99zJf>.
- 11 D. I. Dimitrov, G. Singh, T. Gehr, and M. T. Vechev, "Provably robust adversarial examples," in *The Tenth International Conference on Learning Representations, ICLR*, 2022.  URL: <https://openreview.net/forum?id=UMfhoMtIaP5>.
- 12 M. Fischer, C. Sprecher, D. I. Dimitrov, G. Singh, and M. Vechev, "Shared certificates for neural network verification," in *International Conference on Computer Aided Verification, CAV*, 2022, pp. 127–148.

- 13 G. Bonaert, D. I. Dimitrov, M. Baader, and M. Vechev, “Fast and precise certification of transformers,” in *Proceedings of the 42nd ACM SIGPLAN International Conference on Programming Language Design and Implementation, PLDI*, 2021, pp. 466–481.

Other Publications

- 14 C. Dékány, S. Balauca, D. I. Dimitrov, R. Staab, and M. Vechev, “Mixat: Combining continuous and discrete adversarial training for llms,” *In Submission*, 2025.  URL: <https://drive.google.com/file/d/1ZA0b5DB809QpCmD50EUy90XGF0BDL-6i/view?usp=sharing>.
- 15 A. Alexandrov, V. Raychev, D. I. Dimitrov, C. Zhang, M. Vechev, and K. Toutanova, “Bggpt 1.0: Extending english-centric llms to other languages,” *arXiv preprint arXiv:2412.10893*, 2024.  URL: <https://arxiv.org/abs/2412.10893>.
- 16 T. Fedoseev, D. I. Dimitrov, T. Gehr, and M. Vechev, “LLM training data synthesis for more effective problem solving using satisfiability modulo theories,” *The 4th Workshop on Mathematical Reasoning and AI at NeurIPS*, 2024.  URL: <https://openreview.net/forum?id=hR4Hskr4GX>.
- 17 D. I. Dimitrov, “Image inpainting with gaussian processes,” *Bachelor’s Thesis, The University of Edinburgh, UK*, 2016.  URL: https://drive.google.com/file/d/1wdZZd3kR30q_DitCXF3KZqE7DbAbexYQ/view?usp=sharing.