

Valeriia Cherepanova

✉ vkcherepanovabox@gmail.com | [Google Scholar](#)

Interests

My research goal is to develop reliable, robust, and fair machine learning systems, which can be safely and effectively used for practical applications.

Education

University of Maryland, College Park

PHD IN APPLIED MATHEMATICS

- Advisor: Prof. Tom Goldstein
- Dean's Fellowship

College Park
Aug 2018 - Aug 2023

University College London

MSc IN MODELING BIOLOGICAL COMPLEXITY (COMPLEX)

- Graduated with distinction

London
Sept 2017 - Sept 2018

National Research University Higher School of Economics

BSc IN MATHEMATICS

Moscow
Sept 2013 - June 2017

Industry Experience

Amazon, AWS Responsible AI

POSTDOCTORAL SCIENTIST

- I conduct research in machine learning with a focus on Responsible AI. In particular, I work on developing AI systems which operate according to the standards for fairness, robustness, privacy, security, transparency, and explainability.

Seattle
September 2023-Present

Amazon, Alexa Entertainment

APPLIED SCIENTIST INTERN

- Developed ML solutions to classify different types of Alexa mistakes for improving Alexa Voice Search on FireTV.
- Built ML models for predicting popularity of FireTV Voice Searches from time-series data.

Seattle
Jun 2022 - Aug 2022

Amazon, Alexa Monitoring

APPLIED SCIENTIST INTERN

- Developed NLP solutions to improve transparency of 3P Alexa Skills through detecting non-compliant privacy policy documents.
- Deployed the model in production and built an interactive dashboard.

Bellevue
Jun 2021 - Aug 2021

Teradata

DATA SCIENTIST INTERN

- Designed a machine learning training course for engineers at the company.

Moscow
Jul 2016 - Oct 2016

Selected Publications

LowKey: Leveraging Adversarial Attacks to Protect Social Media Users from Facial Recognition

V. Cherepanova, M. Goldblum, H. Foley, S. Duan, J. P. Dickerson, G. Taylor, T. Goldstein
International Conference on Learning Representations (ICLR), 2021, [paper], [webtool]

Transfer Learning with Deep Tabular Models

R. Levin*, V. Cherepanova*, A. Schwarzschild, A. Bansal, C. B. Bruss, T. Goldstein, A. G. Wilson, M. Goldblum
International Conference on Learning Representations (ICLR), 2023, [paper], [GitHub]

Spotting LLMs With Binoculars: Zero-Shot Detection of Machine-Generated Text

A. Hans, A. Schwarzschild, V. Cherepanova, H. Kazemi, A. Saha, M. Goldblum, J. Geiping, T. Goldstein
arXiv preprint, [paper]

TuneTables: Context Optimization for Scalable Prior-Data Fitted Networks

B. Feuer, R. T. Schirrmeister, **V. Cherepanova**, C. Hegde, F. Hutter, M. Goldblum, N. Cohen, C. White
arXiv preprint, [paper]

A Performance-Driven Benchmark for Feature Selection in Tabular Deep Learning

V. Cherepanova, R. Levin, G. Somepalli, J. Geiping, C. B. Bruss, A. G. Wilson, T. Goldstein, M. Goldblum
Conference on Neural Information Processing Systems Datasets and Benchmarks Track (NeurIPS), 2023, [paper], [GitHub]

Strong Data Augmentation Sanitizes Poisoning and Backdoor Attacks Without an Accuracy Tradeoff

E. Borgnia*, **V. Cherepanova***, L. Fowl*, A. Ghiasi*, J. Geiping*, M. Goldblum*, T. Goldstein*, A. Gupta*
The International Conference on Acoustics, Speech, & Signal Processing (ICASSP), 2021, [paper]

A Deep Dive into Dataset Imbalance and Bias in Face Identification

V. Cherepanova*, S. Reich*, S. Dooley, H. Souri, M. Goldblum, T. Goldstein
AAAI/ACM Conference on AI, Ethics, and Society, 2023 [paper]

Unraveling Meta-Learning: Understanding Feature Representations for Few-Shot Tasks

M. Goldblum, S. Reich*, L. Fowl*, R. Ni*, **V. Cherepanova***, T. Goldstein
International Conference on Machine Learning (ICML), 2020, [paper]

* indicates equal contribution

Conferences and Talks

Panel Discussion at the NeurIPS 2023 Table Representation Learning Workshop

A Performance-Driven Benchmark for Feature Selection in Tabular Deep Learning

- NeurIPS 2023
- NeurIPS 2023 Table Representation Learning Workshop

Transfer Learning with Deep Tabular Models

- Oral Presentation at the NeurIPS 2022 Table Representation Learning Workshop
- Invited Talk at Arthur AI

A Deep Dive into Dataset Imbalance and Bias in Face Identification

- NeurIPS 2022 Workshop on Trustworthy and Socially Responsible Machine Learning
- NeurIPS 2022 Workshop on Algorithmic Fairness through the Lens of Causality and Privacy
- NeurIPS 2022 Workshop on Machine Learning Safety

Technical Challenges for Training Fair Neural Networks

- ICLR 2021 Workshop on Responsible AI

LowKey: Leveraging Adversarial Attacks to Protect Social Media Users from Facial Recognition

- ICLR 2021
- NeurIPS 2020 Resistance AI Workshop
- NeurIPS 2020 Workshop on Dataset Curation and Security

Reviewer Service

NeurIPS 2023, ICML 2023, NeurIPS 2022, ICLR 2022, NeurIPS 2021, NeurIPS 2022 TSRML Workshop, ICLR 2021 RAI Workshop, IEEE TPAMI