

Alexander Bukharin

✉ alexbukharin18@gmail.com ☎ 404-200-6361 🌐 abukharin3 🌐 <https://abukharin3.github.io/>

Education

Aug 2021 – Present	Georgia Institute of Technology <i>Ph.D. in Machine Learning</i>
Aug 2017 – May 2021	Georgia Institute of Technology <i>B.S. in Industrial and Systems Engineering, Concentration in Advanced Studies in Operations Research and Statistics</i>

Work Experience

May 2024 - Present	Research Intern, NVIDIA Alignment Team: Work on (1) implementing scalable RL algorithms for language model training in NeMo-Aligner (2) training language models with reinforcement learning to improve instruction following, code generation, and mathematical reasoning abilities (3) training reward models with different types of feedback, focusing on ranking and regression data.
Sep 2023 - May 2024	Applied Scientist Intern, Amazon Rufus Team: Work on (1) implementing scalable online RL algorithms in NeMo-Aligner (2) using RL to improve LLM capabilities in math, coding, and instruction following.
May 2023 - Aug 2023	Applied Research Intern, LinkedIn Foundation Models Team: Work on continual pre-training of embedding models.
Jan 2023 - May 2023	Research Intern (Part Time), Microsoft: Work on training large language models with reinforcement learning for code generation.
May 2022 - Aug 2022	Applied Machine Learning Research Intern, Bytedance: Work on improving the performance of large graph neural networks for molecular dynamics modeling.
Jun 2021 - Aug 2021	Research Intern, NEC Labs America: Developed novel methods for distribution-based multiple instance learning with applications to pothole detection and fiber optic cable mapping.

Publications and Preprints

- Robust Reinforcement Learning From Corrupted Human Feedback
Alexander Bukharin, Ilgee Hong, Haoming Jiang, Zichong Li, Qingru Zhang, Zixuan Zhang, and Tuo Zhao
In submission, 2024
- Adaptive Preference Scaling for Reinforcement Learning with Human Feedback
Ilgee Hong, Zichong Li, **Alexander Bukharin**, Yixiao Li, Haoming Jiang, Tianbao Yang and Tuo Zhao
In submission, 2024
- Data Diversity Matters for Robust Instruction Tuning
Alexander Bukharin, Shiyang Li, Zhengyang Wang, Jingfeng Yang, Bing Yin, Xian Li, Chao Zhang, Tuo Zhao, Haoming Jiang
In submission, 2024
- Deep Reinforcement Learning from Hierarchical Reward Modeling
Alexander Bukharin, Yixiao Li, Pengcheng He, Tuo Zhao
In submission, 2024
- RNR: Teaching Large Language Models to Follow Roles and Rules

Alexander Bukharin*, Kuan Wang*, Haoming Jiang, Qingyu Yin, Zhengyang Wang, Tuo Zhao, Jingbo Shang, Chao Zhang, Bing Yin, Xian Li, Jianshu Chen and Shiyang Li
ICML 2024 Workshop on Foundation Models in the Wild

- Robust Multi-Agent Reinforcement Learning via Adversarial Regularization: Theoretical Foundation and Stable Algorithms
Alexander Bukharin, Yue Yu, Qingru Zhang, Zhehui Chen, Simiao Zuo, Yan Li, Chao Zhang, Tuo Zhao
Conference on Neural Information Processing Systems, 2023
- Machine Learning Force Fields with Data Cost Aware Training
Alexander Bukharin, Tianyi Liu, Shengjie Wang, Simiao Zuo, Weihao Gao, Wen Yan, Tuo Zhao
International Conference of Machine Learning (ICML) 2023
**A short version was presented at the NeurIPS AI4Science Workshop 2023*
- Ambient Noise based Weakly Supervised Manhole Localization Methods over Deployed Fiber Networks
Alexander Bukharin, Shaobo Han, Yuheng Chen, Ming-Fang Huang, Yue-Kai Huang, Yao Xie, Ting Wang
Optics Express, March 2023
- Adaptive Budget Allocation for Parameter-Efficient Fine-Tuning
Qingru Zhang, Minshuo Chen, **Alexander Bukharin**, Pengcheng He, Yu Cheng, Weizhu Chen, Tuo Zhao
International Conference on Learning Representations (ICLR), 2023
- PLATON: Pruning Large Transformer Models with Upper Confidence Bound of Weight Importance
Qingru Zhang, Simiao Zuo, Chen Liang, **Alexander Bukharin**, Pengcheng He, Weizhu Chen, Tuo Zhao
International Conference of Machine Learning (ICML), 2022
- Early Detection of COVID-19 Hotspots Using Spatio-Temporal Data.
Shixiang Zhu, **Alexander Bukharin**, Liyan Xie, Shihao Yang, Pinar Keskinocak, Yao Xie
IEEE Journal of Selected Topics in Signal Processing, February 2022
**Best Paper Award (Honorable Mention) at ICML Time Series Workshop 2021*
**A short version is accepted for oral presentation and highlighted as a contributed talk by ICML Time Series Workshop 2021*
**Finalist of Best Applied Paper Competition at 2021 INFORMS Workshop on Data Mining and Decision Analytics*
**Excellent Poster Award at Georgia Statistics Day 2021*
- High-resolution Spatio-temporal Model for County-level COVID-19 Activity in the US.
Shixiang Zhu, **Alexander Bukharin**, Liyan Xie, Mauricio Santillana, Shihao Yang, Yao Xie
ACM Transactions on Management Information Systems, July 2021
- Five-Year Project-Level Statewide Pavement Performance Forecasting Using a Two-Stage Machine Learning Approach Based on Long Short-Term Memory.
Alexander Bukharin, Zhongyu Yang, and Yichang Tsai.
Transportation Research Record, May 2021
- Data-Driven Optimization for Police Beat Design in South Fulton, Georgia
Shixiang Zhu, **Alexander W Bukharin**, Le Lu, He Wang, Yao Xie
KDD Workshop on Data Science for Social Good 2021

Research Experience

Jan 23 - Present	Training Language Models with Reinforcement Learning: Work on training large language models via reinforcement learning, with a focus on improving code generation, mathematical reasoning, instruction following, and chat capabilities.
Aug 23 - Present	Instruction Tuning Language Models: Work on automatic dataset selection for instruction tuning of large language models. Also work on improving rule-following ability of language models.
Nov 2021 - May 2023	Efficient Transformer Training: Work on compressing large language models with pruning algorithms and low-rank adaptation. Conduct experiments on pre-trained Vision Transformer models and large language models (i.e. BLOOM). Also working on training large language models with reinforcement learning.
May 2022 - Jan 2023	Pre-Training Graph Neural Networks: Work on improving the pre-training of large graph neural networks for molecular science applications.
Jan 2021 - May 2022	Robust Multi-Agent Reinforcement Learning: Study the effect of adversarial training on the robustness of MARL algorithms to observation noise and different environment changes with applications to traffic light control, autonomous driving, and robotics.
Aug 2019 - May 2021	Spatio-Temporal Data Mining: Develop methods for spatio-temporal modelling and decision making by combining machine learning, statistics, and operations research. This work was motivated by high-impact problems from police operations and epidemiology.

Teaching Experience

Teaching Assistant:

- Stochastic Manufacturing and Service Systems (ISyE 3232), Fall 2021, Spring 2022
- Engineering Economy (ISyE 3025), Spring 2020
- Probability with Applications (ISyE 2027), Fall 2019

Honors

- Presidents Fellowship at Georgia Tech, 2021
- College of Engineering Undergraduate Research Award, 2020
- Presidents Undergraduate Research Award at Georgia Tech, 2018, 2020

Skills

Programming: Python, PyTorch, R (intermediate), Java (intermediate), C (intermediate), SQL (intermediate)

Patents

- Weakly-supervised Learning for Manhole Localization Based on Ambient Noise
Shaobo Han, Yuheng Chen, Ming-Fang Huang, Ting Wang, **Alexander Bukharin**
US Patent No. 20240102833A1