

Bo Liu

☎ (412) 525-7592 | ✉ benjaminliu.eecs@gmail.com | 🏠 [Homepage](#) | 🔍 [Google Scholar](#) | 🌐 [GitHub](#) | **in** [LinkedIn](#)

EDUCATION

- B.S. in Intelligence Science and Technology** Sep. 2016 – Jul. 2020
School of Electronics Engineering and Computer Science (EECS), Peking University Beijing, China
- B.Econ. in Economics** Sep. 2017 – Jul. 2020
National School of Development, Peking University Beijing, China

PUBLICATIONS

* indicates equal contribution

- A Theoretical Understanding of Gradient Bias in Meta-Reinforcement Learning**
Bo Liu*, Xidong Feng*, Jie Ren, Luo Mai, Rui Zhu, Haifeng Zhang, Jun Wang, Yaodong Yang
Conference on Neural Information Processing Systems (**NeurIPS**), 2022
- TorchOpt: An Efficient Library for Differentiable Optimization**
Jie Ren*, Xidong Feng*, **Bo Liu***, Xuehai Pan*, Yao Fu, Luo Mai, Yaodong Yang
Conference on Neural Information Processing Systems OPT Workshop (**NeurIPS Workshop OPT**), 2022
- EnvPool: A Highly Parallel Reinforcement Learning Environment Execution Engine**
Jiayi Weng, Min Lin, Shengyi Huang, **Bo Liu**, Denys Makoviichuk, Viktor Makoviychuk, Zichen Liu, Yufan Song, Ting Luo, Yukun Jiang, Zhongwen Xu, Shuicheng Yan
Conference on Neural Information Processing Systems Datasets and Benchmarks (**NeurIPS D&B**), 2022
- Neural Auto-Curricula in Two-Player Zero-Sum Games**
Xidong Feng, Oliver Slumbers, Ziyu Wan, **Bo Liu**, Stephen McAleer, Ying Wen, Jun Wang, Yaodong Yang
Conference on Neural Information Processing Systems (**NeurIPS**), 2021
- Learning Correlated Communication Topology in Multi-Agent Reinforcement Learning**
Yali Du, **Bo Liu**, Vincent Moens, Ziqi Liu, Zhicheng Ren, Jun Wang, Xu Chen, Haifeng Zhang
International Conference on Autonomous Agents and Multiagent Systems (**AAMAS**), 2021 **Oral**

RESEARCH PROJECTS

- **University College London** Jul. 2020 – Mar. 2023
Research Assistant, Advisor: Jun Wang Beijing, China
 - A Theoretical Understanding of Gradient Bias in Meta-Reinforcement Learning**
 - Proposed a unified framework for variations of gradient-based meta-reinforcement learning (GMRL) algorithms
 - Derived upper bound for two newly identified biases in GMRL: compositional bias and multi-step hessian bias
 - Conducted ablation studies qualitatively and quantitatively to verify how the bias terms affect estimation quality
 - Designed two plug-and-play methods for bias mitigation: off-policy learning correction and LVC Hessian estimator
 - Completed a co-first-authored paper published at **NeurIPS 2022**
 - Learning Correlated Communication Topology in Multi-Agent Reinforcement Learning**
 - Proposed FlowComm to evaluate the correlation between agent communication interactions and learned message-augmented decentralized policies & graph reasoning policies through multi-agent reinforcement learning
 - Generalized coupling flow to model the interaction graph of MARL conditioning on the global states of all agents
 - Demonstrated through visualizations on Particle World that FlowComm learned meaningful communications
 - Completed a second-authored paper published at **AAMAS 2021**
- **Carnegie Mellon University** Oct. 2019 – Jan. 2020
Undergraduate Research Assistant, Advisor: Changliu Liu Pittsburgh, PA

Multi-UAV Collaborative Transportation

- Proposed a hierarchical training approach to explore cooperative strategy for multi-UAV collaborative transportation
- Established RL environment that combined OpenAI gym with real world UAV models within ROS Gazebo simulator
- Developed a modified MADDPG algorithm based on Bidirectional LSTM as centralized value function and then combined it with low-level control to identify the optimal strategy for UAVs
- Designed a novel RL learning agent to assign appropriate tasks to a decision-making module; Achieved improved generalization of lower level learned model

• Microsoft Research Asia

Jul. 2019 – Oct. 2019

Research Intern, Advisor: Qiwei Ye

Beijing, China

Multi-Agent Reinforcement Learning Game AI for Google Research Football

- Developed a Multi-agent Game AI for Google Research Football based on QMIX with pretrained model
- Surveyed and implemented various baseline MARL algorithms such as Independent PPO and QMIX
- Proposed pre-training agents through imitation learning, which significantly outperformed baselines and led to an improved model that was superior to benchmark

• Peking University

Jan. 2019 – Apr. 2019

Undergraduate Research Assistant, Advisor: Zongqing Lu

Beijing, China

Teacher-Student Curriculum Learning for Visual Active Tracking

- Proposed a novel teacher-student curriculum learning approach to train a strong tracker UAV with learning target UAV generating various moving patterns
- Formulated Multi-UAV tracking as Meta-Learning problem, with the target serving as a meta-agent
- Designed a RL learning target with two sub policies to generate experienced and novel moving patterns, which resulted in improved robustness of the tracker agent

OPEN SOURCE PROJECTS

• Peking University

Apr. 2022 – Present

Research Assistant, Advisor: Yaodong Yang

Beijing, China

TorchOpt: An Efficient Library for Differentiable Optimization

- Developed a differentiable AdamW optimizer for gradient-based Meta-Learning research
- Implemented MAML examples to verify the effectiveness of differentiable AdamW and compatibility with functorch
- Designed matrix inversion linear solver with neumann series approximation and implemented implicit MAML omniglot examples with corresponding linear solver
- Completed a co-first-authored paper published at **NeurIPS 2022 Workshop OPT**

• Sea AI Lab

Apr. 2022 – Oct. 2022

Outside Collaborator, Advisor: Zhongwen Xu

Beijing, China

EnvPool: A Highly Parallel Reinforcement Learning Environment Execution Engine

- Achieved 3M FPS throughput with MuJoCo physics engine on a single DGX-A100 machine
- Demonstrated high-performance RL agent training with Atari games and MuJoCo using EnvPool and RL Games
- Conducted MuJoCo environment alignment test, contributed bug reports, and performed debugging
- Completed a co-authored paper published at **NeurIPS 2022 Datasets and Benchmarks**

HONORS & AWARDS

Principal Scholarship For Undergraduate Research

Jul. 2018

For outstanding undergraduate researchers in Peking University chosen through comprehensive evaluation

Award For Scientific Research

Oct. 2017

For students (8 out of 392 in EECS College, PKU) demonstrating exceptional academic performance

TECHNICAL SKILLS

Programming Languages: Python, C/C++, Bash, Go, Stata

Tools and Frameworks: TensorFlow, PyTorch, Git, LaTeX, Docker, TravisCI, Google Cloud Platform, VS Code, Visual Studio, PyCharm