Naicheng He(Arnie)

Undergraduate Student at Brown University

naicheng_he@brown.edu https://arnie-he.github.io/Arnie-He/

EDUCATION

09/2022 – Present **Sc.B. in Applied Math-CS & A.B. in Mathematics**
Brown University

RESEARCH

2025 **Spectral Collapse Drives Loss of Plasticity in Deep Continual Learning**
Neurips ARLET workshop 2025, Now Under Review
He, N., *Guo, K., *Prakash, A., Tiwari, S., Sapio-Kirk, T., Tao, R.Y., Greenwald, A., & Konidaris, G.

We investigate why deep neural networks suffer from loss of plasticity in deep continual learning, failing to learn new tasks without reinitializing parameters. We show that this failure is preceded by Hessian spectral collapse at new-task initialization, where meaningful curvature directions vanish and gradient descent becomes ineffective. To characterize the necessary condition for successful training, we introduce the notion of τ-trainability and show that current plasticity preserving algorithms can be unified under this framework. Targeting spectral collapse directly, we then discuss the Kronecker factored approximation of the Hessian, which motivates two regularization enhancements: maintaining high effective feature rank and applying L2 penalties. Experiments on continual supervised and reinforcement learning tasks confirm that combining these two regularizers effectively preserves plasticity.

2025 **Inverse Reinforcement Learning on GPUDrive**
NYRL workshop 2025, In Progress
He, N., Prakash, A., Swamy, G., Greenwald A., & Vinitsky E.
Investigating imitation for autonomous driving in GPUDrive. Implemented GAIL-based methods to learn from both human and agent demonstrations across diverse driving scenarios. Ongoing work developing game-theoretic approaches to enable robust multi-agent imitation learning and reward generalization.

2024 – 2025 **Bi-Level Policy Optimization with Nyström Hypergradients**
Under Review
Prakash, A.*, **He, N.***, Goktas, D, Greenwald, A.,

The dependency of the actor on the critic in actor-critic (AC) reinforcement learning means that AC can be characterized as a bilevel optimization (BLO) problem, also called a Stackelberg game. This characterization motivates two modifications to vanilla AC algorithms. First, the critic's update should be nested to learn a best response to the actor's policy. Second, the actor should update according to a hypergradient that takes changes in the critic's behavior into account. Computing this hypergradient involves finding an inverse Hessian vector product, a process that can be numerically unstable. We thus propose a new algorithm, Bilevel Policy Optimization with Nyström Hypergradients (BLPO), which uses nesting to account for the nested structure of BLO, and leverages the Nyström method to compute the hypergradient. Theoretically, we prove BLPO converges to (a point that satisfies the necessary conditions for) a local strong Stackelberg equilibrium in polynomial time with high probability, assuming a linear parametrization of the critic's objective. Empirically, we demonstrate that BLPO performs on par with or better than PPO on a variety of discrete and continuous control tasks.

PROFESSIONAL EXPERIENCE

02/2025 – 05/2025 **Undergraduate Teaching Assistant for CSCI 1440/2440: Algorithmic Game Theory**
Brown University

02/2024 – 05/2024 **Undergraduate Teaching Assistant for CSCI 1470/2470: Deep Learning**
Brown University

09/2023 – 01/2024 **Part-time Full-Stack Engineer**
Revvity

- **Developed a mobile app** with React Native frontend, Ruby services, and AWS backend to answer Revvity product customer queries.
- **Fine-tuned LLaMA-2** on internal handouts and historical support emails to provide accurate, automated responses.
- **Reduced technical support emails by 50%** post-deployment, significantly lowering customer support workload.

AWARDS

05/2024 **2nd place in one-shot automated negotiation at the Supply Chain Management League**
AAMAS 2024

09/2023 **Undergraduate Teaching and Research Awards**
Funded for drone downwash research at ACT lab,