曹露凡 (LuFan Cao)

(+86)19180980126
caolf23@mails.tsinghua.edu.cn

EDUCATION

B.S. in Mathematics and Applied Mathematics **2023 – Present**
Tsinghua University, Qiuzhen College Expected Graduation: June 2027

Related Coursework (All Qiuzhen College Courses unless specified):

CS: Machine Learning Theory, Deep Learning, Python for Machine Learning, Object-Oriented Programming in C++, Programming Fundamentals (CS Dept)

STATS: Probability Theory I&II (PhD), Statistical Inference (PhD)

MATH: Mathematical Analysis 1&2, Real & Complex Analysis, Functional Analysis, Topology, Linear Algebra, Abstract Algebra, Differential Geometry, Numerical Analysis 1&2, Number Theory (Math Dept), Discrete Math (Math Dept)

PHYSICS: Quantum Mechanics, Electrodynamics, Theoretical Physics, Physics 1&2

AWARDS: 2022 Silver Medal in the Chinese Mathematical Olympiad (CMO). First Prize in the National High School Mathematics League.

RESEARCH PROJECTS

L1-Coverage Exploration for Reinforcement Learning Repository

- Reproduced the CODEX.W and PSDP (Policy Search by Dynamic Programming) algorithms from *Scalable Online Exploration via Coverability*, implementing hierarchical policy-cover exploration in Tightrope Grid-World environments.
- Proposed a trajectory-wide intrinsic reward mechanism to address the limitation of optimizing only terminal-layer coverage, significantly improving overall state-space coverage efficiency under limited policy budgets.
- Replaced the original transition-based intrinsic reward with a state-based reward formulation, reducing reward-estimation complexity while improving sample efficiency, training stability, and policy convergence speed.
- Designed neural-network-based Q-function approximation using Sinusoidal Positional Encoding with MLP architectures, extending tabular exploration methods to larger state spaces and demonstrating stronger generalization in 5×100 environments.
- Conducted systematic experiments comparing Tabular, Neural, and Uniform Policy baselines, together with ablation studies on reward design and exploration behavior.

Deep Reinforcement Learning for Hex Grid Game Repository

- Game AI environment: Built Gymnasium-compatible Hex grid game with game-theoretic properties and configurable reward structures for two agent self-play training.
- RL agents: Trained PPO (w/ PyTorch DDP for multi-GPU: 4 A800's) and Actor-Critic MCTS pipelines; ResNet policy network with residual connections captures spatial board structure; GAE advantage estimation with action masking.

Mathematical Contest in Modeling (includes 23 Page Report)

- Developed a data-driven wear model by extending the Archard equation and applying Simulated Annealing to predict historical foot traffic and tourist flow patterns.
- Implemented Unsupervised Learning techniques, specifically X-Means clustering and Z-score analysis, to detect

anomalies and classify staircase maintenance states.

- Proved that optimal approximations for tread edge and center depressions equal the number of independent users, while simultaneous two-person use converges to zero.

Video Generation Model

- Architecture Analysis: Reverse-engineered the Sora framework, dissecting the shift from U-Nets to Diffusion Transformers and "Patch n' Pack" strategies for variable resolution inputs.
- Empirical Evaluation: Benchmarked Open-Sora on RTX 4090 hardware, quantifying model sensitivity to prompt complexity and identifying bottlenecks in physical realism.
- Research Proposal: Proposed novel architectural improvements integrating differentiable physics engines and LLM-based temporal planning to address "physics hallucination" and coherence issues in generative video.