Ehsan Futuhi

futuhiehsan@gmail.com
Google Scholar
ehsan-futuhi
ehsan-fmb

Education

University of Alberta — Edmonton, Canada
Master of Science, Computing Science — 2021–2024
Research Area: Reinforcement Learning and Search
Supervisor: Martin Mueller
GPA: 3.9/4

Amirkabir University of Technology - Tehran Polytechnic — Tehran, Iran
Bachelor of Science, Computer Engineering — 2016–2021
GPA: 17.92/20

Experience

Moving AI Lab — Edmonton, Canada
Research Assistant (Part-time) — 02/2026 – present
- Topic: Evaluating reasoning capabilities of LLMs in planning.

MedWatch Technologies — Florida, USA
Data Scientist — 01/2025 - present
- Topic: Non-invasive blood glucose monitoring using AI

Moving AI Lab — Edmonton, Canada
Research Assistant — 04/2024 - 01/2025
- Topic: Designing Admissible Heuristics for Single-Agent Search Algorithms

University of Alberta — Edmonton, Canada
Graduate Research Fellow — 09/2021 - 04/2024
- Topic: Improving sample efficiency of RL algorithms by Search and definition of Subgoals

University of Alberta — Edmonton, Canada
Graduate Teaching Assistant — 09/2021 - 05/2022
- Courses: Intelligent Systems, Foundations of Computation I

Amirkabir University of Technology - Tehran Polytechnic — Tehran, Iran
Teaching Assistant — 08/2020 - 04/2021
- Courses: Algorithm Design, Operating Systems

Institute for Research in Fundamental Sciences (IPM) — Tehran, Iran
Research Internship — 07/2019 - 06/2020
- Research area: solving games by AI

Publications

1. **Ehsan Futuhi**, Nathan R. Sturtevant. "Planning with Transformers: Chain of Computation and Structured Context Windows.", *preprint version available on arXiv; presented as an invited talk by Nathan R. Sturtevant at the LM4Plan Workshop at ICML 2026.*

2. **Ehsan Futuhi**, Nathan R. Sturtevant. "Learning Admissible Heuristics for A*: Theory and Practice", *International Conference on Learning Representations (ICLR)*, 2026 (**28%** acceptance rate).

3. **Ehsan Futuhi**, Nathan R. Sturtevant. "A Parallel CPU-GPU Framework for Batching Heuristic Operations in Depth-First Heuristic Search", *AAAI Conference on Artificial Intelligence (AAAI 2026)*, Accepted as Oral Presentation (Acceptance Rate: **17.6%**).

4. **Ehsan Futuhi**, Shayan Karimi, Chao Gao, Martin Müller. "ETGL-DDPG: A Deep Deterministic Policy Gradient Algorithm for Sparse Reward Continuous Control", *Transactions on Machine Learning Research (TMLR)*, 2025.

5. **Ehsan Futuhi**. "Improving Deep Deterministic Policy Gradient for Sparse Reward and Goal-Conditioned Continuous Control", *Master's Thesis, University of Alberta, 2024.*

6. Kiarostami, Mohammad Sina, Mohammadreza Daneshvaramoli, Saleh Khalaj Monfared, Aku Visuri, Helia Karisani, Simo Hosio, Hamed Khashehchi, **Ehsan Futuhi**, Dara Rahmati, and Saeid Gorgin. "On using Monte- Carlo tree search to solve puzzles.", *In 2021 7th International Conference on Computer Technology Applications, pp. 18-26. 2021.*

7. Sharareh Alipour, **Ehsan Futuhi**, Shayan Karimi. "On Distributed Algorithms For Minimum Dominating Set Problem, from theory to application.", *preprint version available on arXiv.*

8. **Ehsan Futuhi**, Shayan Karimi. "Solving Nurikabe with Monte-Carlo Tree Search.", *in 2020 5th International Conference on Computational Intelligence and Applications (ICCIA), IEEE, 2020, pp. 33–38.*

Selected Projects

See full list of projects on my github

- Reasoning capabilities of LLMs in planning tasks (2026)
  - Developed novel environments for benchmarking LLM planning.
  - Introduced a specialized prompting strategy achieving 100% plan validity.
  - Fine-tuned state-of-the-art LLMs to establish robust performance baselines.
  - Conducted extensive experiments across diverse baseline models.
- Learning Admissible Heuristics (2025)
  - Developed a novel loss function for optimizing admissible heuristics in search.
  - Conducted sample complexity analysis for the learnability of admissible models.
  - Executed extensive experiments achieving admissible models for the Rubik's Cube.
- ETGL-DDPG: Enhanced DDPG for sparse reward tasks (2025)
  - Design a new exploration algorithm.
  - Design a new replay buffer mechanism.
  - Extensive experiments on robotics benchmarks: Mujoco, and Meta-World.
- Batch IDA* (2024)
  - Implemented CPU and GPU parallelism to accelerate large-scale heuristic search.
  - Introduced a novel Batch IDA* variant capable of integrating ML-driven heuristics.
  - Leveraged generative models, including LLMs, to synthesize high-quality solution paths.

- Safe-MCTS (2023)
  - A MCTS approach for stochastic constrained MDPs.
  - A safety critic and policy prior attained by actor-critic algorithm.
- Policy Gradient (PG) Algorithms (2022)
  - Implement the state-of-the-art PG algorithms.
  - Evaluate exploration strategies in RL by their performance combined with PG algorithms.
- Analysing the game of Connect 4 (2021)
  - Solve the game by proof-number search.
  - Design MCTS and AlphaZero players.

Scholarships and Awards

- B.Sc. Full scholarship for being in the top 700 of the National Public Universities Exam 2016

Areas of Specialization

- **Reinforcement Learning & Robotics:** Deep RL, Policy Gradients, and sample efficiency with sparse-reward.
- **AI Planning & Search:** Heuristic Search and Monte Carlo Tree Search (MCTS).
- **Foundation Models:** LLM/VLM reasoning for planning, and RL fine-tuning.
- **Time Series Forecasting:** Blood glucose level prediction from non-invasive biosensor signals.

Technical Skills

- **Coding:** Python, Jax, Java, C/C++.
- **Deep Learning:** Pytorch/Libtorch, Tensorflow, PDDL.
- **Cloud & Infrastructure:** GCP, AWS, Azure , Docker, Kubernetes.

Conference Attendance

- International Conference on Learning Representations (ICLR) 2026
- Association for the Advancement of Artificial Intelligence (AAAI) 2026
- International Symposium on Combinatorial Search (SoCS) 2024
- Upper Bound AI Conference 2024
- Upper Bound AI Conference 2023

Peer Reviewing Experience

- International Conference on Learning Representations (ICLR) 2025-2026
- Association for the Advancement of Artificial Intelligence (AAAI) 2026
- International Conference on Automated Planning and Scheduling (ICAPS) 2025
- Uncertainty in Artificial Intelligence (UAI) 2024

Languages

- **English:** Fluent
- **Persian:** Native
- **Turkish:** Native