

Gaurav Dixit

<https://gdixit.com> – (425)-647-3679 – dixitg@oregonstate.edu

EDUCATION AND DEGREES AWARDED

- 08.09.2023 **Doctor of Philosophy (Ph.D.) in Robotics**
Oregon State University, Corvallis, OR, USA
Thesis: *Learning to coordinate in sparse asymmetric multiagent systems*
Advisor: Prof. Kagan Tumer
- 03.09.2021 **Master of Science (M.S.) in Computer Science**
Oregon State University, Corvallis, OR, USA
Advisor: Prof. Kagan Tumer
- 01.05.2016 **Honors Bachelor of Engineering (B.Eng.), *magna cum laude***
University of Pune, Pune, India
Pune Institute of Computer Technology

OTHER EDUCATION AND EXPERTISE

Advanced graduate coursework in multiagent systems, reinforcement learning, robotics, and learning-based control at Oregon State University.

Self-directed expertise with multi-robot simulation frameworks (MuJoCo, Robosuite, pagmo, custom RL environments).

LANGUAGE SKILLS

Native: English; Hindi

B1: German; Finnish

CURRENT EMPLOYMENT

- 2023–present **Postdoctoral Scholar**
AI-CARING Institute & Collaborative Robotics and Intelligent Systems Institute (CoRIS), Corvallis, OR
- Develop multi-objective multiagent learning methods that enable autonomous systems to reason across evolving stakeholder priorities, sustain coordinated behavior under uncertainty, and support human-centered decision-making in domains such as eldercare and multi-robot teaming.
 - Mentor graduate students; co-design and co-teach Multiagent Systems and Learning-Based Control.
 - Contributor to multi-institution grants with NSF, AFOSR, and ONR.

RESEARCH PROFILE

AI and Robotics researcher specializing in multiagent and multi-objective decision-making, adaptive and continual learning, and human-centered AI, with applications in assistive robotics, healthcare, and complex multiagent environments.

TEACHING MERITS

Course design and graduate teaching:

Co-designed and co-taught (2023-2025):

- Multiagent Systems (coordination, decision-making, learning)
- Learning-Based Control (control, reinforcement learning, embodied AI)

Teaching assistantships:

Machine Learning and Data Mining (2015);
 Data Structures and Algorithms (2015);
 Programming Paradigms with C++ (2014);
 Operating System Administration (2013);
 Functional Programming with Haskell (2013).

RESEARCH OUTPUT

Research expertise:

Multi-objective multiagent learning; adaptive multi-robot coordination; human-aligned decision-making.

Peer-reviewed venues include: AAMAS, GECCO, ECAI, VIS, IJCCI.

Full list of publications: <https://scholar.google.com/citations?user=1sh3qCAAAAAJ>

Peer-reviewed Publications:

1. Thakar, R.*, **Dixit, G.**, Iyer, S.*, & Tumer, K. (2025). *Multiagent Credit Assignment for Multi-Objective Coordination*. Proceedings of GECCO 2025, 663–672. <https://dl.acm.org/doi/10.1145/3712256.3726445>
2. Gonzalez, E., **Dixit, G.**, & Tumer, K. (2025). *Dynamic Influence for Coevolutionary Agents*. Proceedings of GECCO 2025, 368–376. <https://dl.acm.org/doi/10.1145/3712256.3726455>
3. Iyer, S.*, Aydeniz, A. A., **Dixit, G.**, & Tumer, K. (2025). *Multiagent Quality-Diversity for Effective Adaptation*. In ECAI 2025 (pp. 3323–3330). IOS Press. <https://ebooks.iospress.nl/doi/10.3233/FAIA251201>
4. Festa, A.*, **Dixit, G.**, & Tumer, K. (2024). *Influence-Focused Asymmetric Island Model*. Proceedings of AAMAS 2024, 2261–2263. <https://www.ifaamas.org/Proceedings/aamas2024/pdfs/p2261.pdf>
5. **Dixit, G.** & Tumer, K. (2024). *Informed Diversity Search for Learning in Asymmetric Multiagent Systems*. Proceedings of GECCO 2024, 313–321. <https://dl.acm.org/doi/10.1145/3638529.3654206>
6. Festa, A.*, **Dixit, G.**, & Tumer, K. (2024). *Reinforcing Inter-Class Dependencies in the Asymmetric Island Model*. Proceedings of GECCO 2024, 740–748. <https://dl.acm.org/doi/10.1145/3638529.3654213>
7. **Dixit, G.** & Tumer, K. (2024). *Objective-Informed Diversity for Multi-Objective Multiagent Coordination*. In ECAI 2024 (pp. 3660–3667). IOS Press. <https://ebooks.iospress.nl/doi/10.3233/FAIA240923>
8. **Dixit, G.** & Tumer, K. (2023). *Learning inter-agent synergies in asymmetric multiagent systems*. Proceedings of AAMAS 2023, 1569–1577. https://gdixit.com/assets/publications/AAMAS23_AIM.pdf
9. **Dixit, G.** & Tumer, K. (2023). *Learning synergies for multi-objective optimization in asymmetric multiagent systems*. Proceedings of GECCO 2023, 447–455. <https://dl.acm.org/doi/10.1145/3583131.3590524>
10. Nickelson, A., Zerbel, N., **Dixit, G.**, & Tumer, K. (2023). *Shaping the Behavior Space with Counterfactual Agents in Multi-Objective Map-Elites*. IJCCI 2023, 41–52. <https://www.scitepress.org/Papers/2023/121648/121648.pdf>
11. **Dixit, G.** & Tumer, K. (2022). *Balancing teams with quality-diversity for heterogeneous multiagent coordination*. Proceedings of the Genetic and Evolutionary Computation Conference Companion, 236–239. <https://dl.acm.org/doi/10.1145/3520304.3529062>
12. **Dixit, G.** & Tumer, K. (2022). *Behavior Exploration and Team Balancing for Heterogeneous Multiagent Coordination*. Proceedings of AAMAS 2022. <https://par.nsf.gov/servlets/purl/10359619>
13. **Dixit, G.**, Gonzalez, E., & Tumer, K. (2022). *Diversifying behaviors for learning in asymmetric multiagent systems*. Proceedings of the Genetic and Evolutionary Computation Conference, 350–358. <https://dl.acm.org/doi/10.1145/3512290.3528860>

14. **Dixit, G.**, Koll, C.*, & Tumer, K. (2021). *Heterogeneous agent coordination via adaptive quality diversity and specialization*. Proceedings of the Genetic and Evolutionary Computation Conference Companion, 95–96. <https://dl.acm.org/doi/10.1145/3449726.3459564>
15. Olson, M. L., Nguyen, T.-V., **Dixit, G.**, Ratzlaff, N., Wong, W.-K., & Kahng, M. (2021). *Contrastive identification of covariate shift in image data*. 2021 IEEE Visualization Conference (VIS), 36–40. IEEE. <https://arxiv.org/pdf/2108.08000>
16. **Dixit, G.**, Zerbé, N., & Tumer, K. (2020). *Gaussian processes as multiagent reward models*. AAMAS Conference Proceedings. <https://www.ifaamas.org/Proceedings/aamas2020/pdfs/p330.pdf>
17. **Dixit, G.**, Zerbé, N., & Tumer, K. (2019). *Dirichlet-multinomial counterfactual rewards for heterogeneous multiagent systems*. International Symposium on Multi-Robot and Multi-Agent Systems (MRS), 209–215. IEEE. https://gdixit.com/assets/publications/MRS19_DMCS.pdf

* Denotes supervised undergraduate or graduate student.

PREVIOUS WORK EXPERIENCE

2018–2023 **Graduate Research Assistant**

Oregon State University, Corvallis, OR

- Developed multiagent learning methods that address reward sparsity, credit assignment, and the emergence of coordinated behavior in environments with partially aligned or dynamic objectives.
- Designed diversity-search methods for improving zero-shot generalization to changes in task dynamics, agent policies and team composition.

2022–2025 **Artist Researcher (with Antti Ilvessuo)**

C37 Collective, Helsinki, Finland

- Explore computational and evolutionary techniques for generating aleatoric temporal art—particularly music—studying how stochastic structure and variation processes shape creative outcomes.
- Develop multimodal visualization tools to interpret transformations induced by generative models.

2019–2019 **AI Research Engineer**

RedLynx Oy, Helsinki, Finland

- Investigated approaches combining tree-based planning with reinforcement learning for zero-sum decision-making.
- Advanced quality-diversity techniques for learning latent behavior representations to support automated policy exploration and diversity-driven search.

2016–2018 **AI / Physics Programmer**

Ubisoft Entertainment SA, Pune, India

- Implemented policy-gradient and evolutionary optimization methods to develop an end-to-end pipeline for automated game testing.
- Designed and developed a web-based rigid-body physics engine in TypeScript/Node.js to support online simulation and interactive tools.

2015–2016 **Software Development Research Intern**

BMC Software, Pune, India

- Developed a real-time event analysis tool for root-cause diagnosis and event-association mining using an optimized variant of the Rete algorithm for network automation.

2014–2016 **Undergraduate Research Assistant**

Pune Institute of Computer Technology

- Developed methods to improve named-entity recognition for Hindi and regional dialects through data collection, feature engineering, and statistical modeling.

RESEARCH FUNDING AND GRANTS

Contributor to multi-institution research proposals funded by the US National Science Foundation (NSF), Air Force Office of Scientific Research (AFOSR), and Office of Naval Research (ONR).

Actively developed and submitted white papers to AFOSR with Prof. Kagan Tumer.

AWARDS AND HONOURS

Magna cum laude, B.Eng. in Computer Science, University of Pune (2016).

ACADEMIC SERVICE

Program committees & tutorial organization:

ECAI Tutorial on Multi-Objective Multiagent Decision-Making (2024).

Program Committee: ECAI (2024–2025); ALA workshop (2023–2025).

Organizer: Multi-Objective Decision Making (MODeM) workshop at ECAI/IJCAI.

Reviewing:

Neural Computing and Applications; IEEE TEC; AAMAS; GECCO; ECAI; UAI; AAAI; IJCAI.

SCIENTIFIC AND SOCIETAL IMPACT

Open-source contributions in multiagent reinforcement learning, multi-objective optimization, and evolutionary algorithms (pagmo). Work on responsible AI and beneficent intelligence for eldercare with the AI-CARING Institute. Creative collaborations exploring the interplay between humans and computers in art (C37 Collective).

Software and tools:

Created testbed for multi-arm manipulation with MuJoCo for the Office of Naval Research (ONR). Developed multiagent learning frameworks for asymmetric multiagent systems, multi-objective RL, and QD search.

Artistic and creative work:

Aleatoric temporal art and generative visualization tools (C37 Collective).