

# Cevahir Koprulu

cevahir.koprulu@utexas.edu · 1-512-902-3307 · [Website](#) · [Github](#) · [LinkedIn](#)

## RESEARCH INTERESTS

**Keywords:** Reinforcement learning (RL), Generalization, and Curriculum Learning

My research focuses on RL: which tasks an agent trains on and what signal it learns from. My work proposes automated curriculum generation methods for heavy-tailed task distributions, temporal specifications, and cost constraints, and for the scale of batched simulators. In addition, I developed methods for reward synthesis from task-agnostic experience with a few demonstrations, and from physics-informed dynamics models with distance-aware uncertainty estimates. This work spans offline RL, learning from demonstrations, unsupervised environment design, and neuro-symbolic RL, with applications in autonomous driving. Ongoing work extends curriculum generation to composing traffic behaviors and to unlearning in large language models.

## EDUCATION

**University of Texas at Austin**, Austin, TX, USA  
Electrical and Computer Engineering

M.S. 2023, PhD 2027  
GPA: 3.93/4

**Bilkent University**, Ankara, Turkey  
Electrical and Electronics Engineering

B.Sc. June 2021  
GPA: 3.73/4

**Honours/Awards:** (i) Bilkent University Comprehensive Scholarship (ii) Scholarship of Turkish Prime Ministry (iii) Bilkent University EEE Department High Honours (iv) Ranked 24<sup>th</sup> among 2 million students in university entrance exam (v) FIRST Robotics Competition 2015: Rookie All-Star

## WORK EXPERIENCE

### Waymo

*Software Engineering Intern* (all work is under NDA) || Python, SQL, JAX

Mountain View, CA  
June 2026 - August 2026

- Researched **learning from unsupervised driving data within Waymo's data flywheel** for onboard planner models, with a focus on undesirable driving behaviors. Surveyed state-of-the-art methods, implemented and compared a selection within the training infrastructure, and proposed an algorithm and data roadmap for future deployment.

### Bosch Center for Artificial Intelligence

*Research Intern* (all work is under NDA) || Python, PyTorch

Sunnyvale, CA  
May 2025 - August 2025

- Developed **CL4AD**, the *first* integration of curriculum learning into GPUdive, a batched autonomous driving simulator, reducing wall clock time to achieve 99% success rate by 77%, more than a billion steps earlier than domain randomization, i.e., the default uniform scenario sampling method.

### Honda Research Institute - US

*Research Intern* (all work is under NDA) || Python, PyTorch, RAY

Ann Arbor, MI  
May 2024 - August 2024

- Developed an **action advising framework**, GEN2SPEC, that distills knowledge of generalist agents to train specialist agents in a continual learning setting. This framework accelerates the training of an RL agent via action advice from a transformer-based meta-RL agent that can adapt to unseen tasks.

### Eatron Technologies

*Engineering Intern* (all work is under NDA) || Python, C++, PyTorch, ROS

Istanbul, Turkey  
June 2020 - June 2021

- Designed **graph convolutional** and **self-attention**-based neural network architectures to extract spatio-temporal features of a traffic scene for trajectory prediction in a Level-2+ ADAS powered vehicle.

### ROKETSAN (in collaboration with Bilkent University)

*Industrial Design Project* (all work is under NDA) || Python, C++, PyTorch, ROS

Ankara, Turkey  
Sept 2019 - June 2020

- Developed a mobile robot that follows a human leader, combining **YOLOv3 for object detection** and **artificial potential field method for path planning** in a mapped area with unknown dynamic obstacles.

## RESEARCH WORK

[Scaling Curriculum Learning For Autonomous Driving](#)

*ArXiv, 2026*

*Cevahir Koprulu, David Paz, Feng Tao, Yuliang Guo, Xinyu Huang, Ufuk Topcu, Liu Ren*

[Neural Stochastic Differential Equations for Uncertainty-Aware, Offline Reinforcement Learning](#)

*ICLR, 2025*

*Cevahir Koprulu, Franck Djeumou, Ufuk Topcu*

[Safety Prioritizing Curricula for Constrained Reinforcement Learning](#)

*ICLR, 2025*

*Cevahir Koprulu, Thiago D. Simão, Nils Jansen, Ufuk Topcu*

|                                                                                                                                                                                                                                   |                               |
|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-------------------------------|
| Dense Dynamics-Aware Reward Synthesis: Integrating Prior Experience with Demonstrations<br><b>Cevahir Koprulu</b> , Po-han Li, Tianyu Qiu, Ruihan Zhao, Tyler Westenbroek, David Fridovich-Keil, Sandeep P. Chinchali, Ufuk Topcu | L4DC, 2025                    |
| Joint Learning of Reward Machines and Policies in Environments with Partially Known Semantics<br>Christos Verginis, <b>Cevahir Koprulu</b> , Sandeep Chinchali, Ufuk Topcu                                                        | Artificial Intelligence, 2024 |
| Risk-Aware Curriculum Generation for Heavy-tailed Task Distributions<br><b>Cevahir Koprulu</b> , Thiago D. Simão, Nils Jansen, Ufuk Topcu                                                                                         | UAI, 2023                     |
| Reward-Machine-Guided, Self-Paced Reinforcement Learning<br><b>Cevahir Koprulu</b> , Ufuk Topcu                                                                                                                                   | UAI, 2023                     |
| Act to Reason: A Dynamic Game Theoretical Driving Model for Highway Merging Applications<br><b>Cevahir Koprulu</b> , Yildiray Yildiz                                                                                              | CCTA, 2021                    |

TECHNICAL SKILLS

- Programming languages: (Expert) Python, (Proficient) C++, Julia
- Frameworks: PyTorch, JAX, TensorFlow, ROS, RAY, verl, vLLM, Hugging Face Transformers

RELEVANT COURSEWORK

- **UT Austin:** Online Learning, Causality and Reinforcement Learning, Statistical Machine Learning, Learning-based Optimal Control, Game-Theoretic Modeling of Multi-Agent Systems, Program Synthesis, Cyber-Physical Systems, Reinforcement Learning, Convex Optimization, Probability and Statistics, and Complex Networks in the Real World.
- **École polytechnique fédérale de Lausanne** (Exchange - Spring 2019): Image Analysis and Pattern Recognition, Convex Optimization, Deep Learning, and Biological Modelling of Neural Networks.
- **Bilkent University:** Statistical Learning and Data Analytics, Robust Feedback Theory, Introduction to Financial Mathematics, Neural Networks, and Game Theory.

RESEARCH AND LEADERSHIP EXPERIENCES

|                                                                                                                                                                                                                                                                                                                                                              |                                         |
|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------------------------------------|
| <b>Reinforcement Learning Reading Group</b><br><i>Organizer</i>                                                                                                                                                                                                                                                                                              | Austin, TX<br>Aug 2025 - Ongoing        |
| <ul style="list-style-type: none"> <li>• RLRG at the CS Department of UT Austin is a student-run group that discusses research papers on RL.</li> </ul>                                                                                                                                                                                                      |                                         |
| <b>Center for Autonomy</b><br><i>Ph.D. Student</i>                                                                                                                                                                                                                                                                                                           | Austin, TX<br>Aug 2021 - Ongoing        |
| <b>Systems Laboratory</b><br><i>Undergraduate Researcher</i>                                                                                                                                                                                                                                                                                                 | Ankara, Turkey<br>Sept 2019 - July 2021 |
| <ul style="list-style-type: none"> <li>• Developed <b>human driver models</b> from real-traffic data that can change its reasoning level dynamically by combining <b>level-k game theory</b> and <b>reinforcement learning</b>.</li> </ul>                                                                                                                   |                                         |
| <b>IEEE Robotics and Automation Society at Bilkent University</b><br><i>Chairman</i>                                                                                                                                                                                                                                                                         | Ankara, Turkey<br>May 2017 - June 2018  |
| <ul style="list-style-type: none"> <li>• Organized “Mühendis Kafası” in cooperation with Technology Development Foundation of Turkey: Series of sessions on Computer Vision and Deep Learning.</li> <li>• Gave lectures on robotics, control techniques, and related micro-controller programming: EE-101: Introduction to Robotics with Arduino.</li> </ul> |                                         |

LANGUAGE SKILLS

- Turkish: Native proficiency
- English: TOEFL 110/120 (Fall 2020)
- French: DELF B1 (Spring 2015)

RECREATIONAL INTERESTS

I enjoy climbing, cycling, watching/playing soccer (football :)), and learning about history and psychology.