Beyazit Yalcinkaya

Ph.D. Candidate, UC Berkeley | Reinforcement Learning | Multi-Agent Systems | Safe Autonomy

beyazit@berkeley.edu | beyazit-y.github.io | LinkedIn | GitHub | Google Scholar

Summary

Ph.D. candidate at **UC Berkeley**, building reinforcement learning (RL) agents that generalize across long-horizon tasks, and verifying the safety of AI-based autonomous systems. Industry experience at **Zoox** (self-play RL for simulation, **one first-author publication and two patent applications**) and **Amazon Web Services** (formal tools for distributed systems). **14 publications** (11 as first or co-first author), including **ICML**, **NeurIPS**, **IEEE TCAD**, and **FMCAD**; creator of open-source JAX libraries for RL.

Education

University of California, Berkeley — Berkeley, CA
Ph.D. in Computer Science | GPA: 3.925/4.0 | Advisor: Sanjit A. Seshia — *2021 – May 2027 (expected)*
Focus: RL, multi-agent systems, representation learning, robotics, simulation, statistical verification, and runtime monitoring

Middle East Technical University (METU) — Ankara, Türkiye
B.S. in Computer Engineering | GPA: 3.93/4.0 — *2015 – 2020*
Valedictorian: ranked 1st of 249 in the department | TÜBİTAK Research Fellowship

Industry Experience

Zoox (an Amazon company) — Foster City, CA
Machine Learning Engineer Intern | Multi-Agent Simulation | Mentor: Marcell Vazquez-Chanlatte — *Summer 2026*

- Designed new goal representations for goal-conditioned RL policies trained through **large-scale self-play**, expanding the behavioral capabilities of simulation agents and enabling new test scenarios for validating autonomous-driving software.
- Contributed to a **counterfactual scenario generation** pipeline combining self-play policies with a **multimodal world model**.
- Implemented in JAX and Ray; contributions resulted in **one first-author publication and two patent applications**.

Amazon Web Services (AWS) — East Palo Alto & Cupertino, CA
Applied Scientist Intern (returning) | Availability of Distributed Systems | Mentor: Ankush Desai — *Summer 2023*

- Developed **PAvail**, an internal tool for analyzing the availability of distributed systems, built on **P**, a state-machine-based language for formally modeling and specifying distributed systems.
- Applied PAvail to AWS distributed protocols to analyze their availability under varying network traffic conditions; contributed improvements to P's core implementation in C#.

Applied Scientist Intern | Observability of Distributed Systems | Mentor: Ankush Desai — *Summer 2022*

- Built **PObserve**, an internal tool that checks distributed systems against their formal P specifications using real system logs; applied it to distributed protocols used across AWS.
- Implemented a streaming pipeline on **Apache Flink** to process large-scale log data from Amazon S3.

Research Experience

University of California, Berkeley – Learn & Verify Group — Berkeley, CA
Graduate Student Researcher | Advisor: Sanjit A. Seshia — *2021 – Present*

- Lead research on **goal-conditioned and multi-agent reinforcement learning** for tasks given as formal specifications, pairing theoretical results with practical implementations; first author at **ICML 2026**, **NeurIPS 2024**, and NeuS 2025/2026.
- Created **RAD Embeddings**, pretrained graph neural network encoders of formal task specifications with provable correctness guarantees, enabling **zero-shot generalization** to unseen tasks and faster policy learning (NeurIPS 2024, NeuS 2025).
- Developed a **cooperative multi-agent RL** method that learns decentralized team policies for long-horizon tasks (ICML 2026).
- Translated formal task specifications into natural-language instructions via counterexample-guided synthesis with an **LLM** as the oracle, then benchmarked language embeddings of these instructions against RAD Embeddings (NeuS 2026).
- Built and maintain open-source, pip-installable **JAX** libraries that serve as the modular foundation for all RL research code: dfax (JAX-native automata), dfa-gym (single- and multi-agent RL environments with automata objectives), and rad-embeddings (pretrained encoders with CUDA support); code for every RL publication, including end-to-end JIT-compiled single- and multi-agent PPO training, is available at github.com/rad-dfa.
- Developed compositional, simulation-based **statistical verification** methods for AI-based autonomous systems (RV 2023, RV 2026) and **ULGEN**, a runtime monitoring framework for programming safe autonomous systems (IEEE TCAD 2023).

École Polytechnique Fédérale de Lausanne (EPFL) – Dependable Systems Lab — Lausanne, Switzerland
Research Intern | Summer@EPFL Fellow | Advisor: George Candea — *Summer 2019*

- Built a tool that computes performance contracts for software network functions across their entire software stack, along with a Python library supporting a range of network functions.

Max Planck Institute for Software Systems (MPI-SWS) – Real-Time Systems Group — Kaiserslautern, Germany
Research Intern | Advisor: Björn B. Brandenburg — *Summer 2018*

- Developed the **first exact schedulability test** for non-preemptive self-suspending real-time tasks by modeling them as timed automata; the test has become a major benchmark in this literature (first author, **DATE 2019**).

Middle East Technical University (METU) – Cyber-Physical Systems Research Group — Ankara, Türkiye
Undergraduate Research Assistant | Advisor: Ebru Aydin Gol — *2017 – 2020*

- Designed a clock-reduction method for timed automata (first author, FormaliSE 2019; 1st place in METU's CENG 488 Guided Research); co-developed an automated system-repair framework using Signal Temporal Logic (Acta Informatica 2022).

Publications

*Equal contribution. Full list on Google Scholar.

IEEE IV 2027 **Learning to Interact in Multi-Agent Self-Play for Controllable Autonomous Driving Scenario Generation** – **In review**
B. Yalcinkaya, E. Pronovost, M. Naumann, M. Vazquez-Chanlatte

ICML 2026 **Automata-Conditioned Cooperative Multi-Agent Reinforcement Learning**
B. Yalcinkaya, M. Vazquez-Chanlatte, A. Shah, H. Krasowski, S. A. Seshia

NeuS 2026 **Learning Task-Conditioned Policies with Natural Language and Graph Embeddings of Formal Specifications**
B. Yalcinkaya, M. Vazquez-Chanlatte, S. A. Seshia

NeuS 2026 **Controllability in Preference-Conditioned Multi-objective Reinforcement Learning**
P. de las Heras Molins, **B. Yalcinkaya**, L. Peters, D. Fridovich-Keil, G. Bakirtzis

RV 2026 **Scenario-Based Compositional Statistical Model Checking for Safety Specifications**
A. Pomalapally, A. Raeesi, K. K.-C. Chang, **B. Yalcinkaya**, S. A. Seshia

NeuS 2025 **Provably Correct Automata Embeddings for Optimal Automata-Conditioned Reinforcement Learning**
B. Yalcinkaya, N. Lauffer, M. Vazquez-Chanlatte, S. A. Seshia

NeurIPS 2024 **Compositional Automata Embeddings for Goal-Conditioned Reinforcement Learning**
B. Yalcinkaya*, N. Lauffer*, M. Vazquez-Chanlatte*, S. A. Seshia

NeurIPS-W 2023 **Automata-Conditioned Reinforcement Learning with Experience Replay** – **Spotlight**
B. Yalcinkaya*, N. Lauffer*, M. Vazquez-Chanlatte, S. A. Seshia | NeurIPS Workshop on Goal-Conditioned RL

RV 2023 **Compositional Simulation-Based Analysis of AI-Based Autonomous Systems for Markovian Specifications**
B. Yalcinkaya, H. Torfah, D. J. Fremont, S. A. Seshia

IEEE TCAD 2023 **ULGEN: A Runtime Assurance Framework for Programming Safe Cyber-Physical Systems**
B. Yalcinkaya, H. Torfah, A. Desai, S. A. Seshia

FMCAD 2022 **Learning Deterministic Finite Automata Decompositions from Examples and Demonstrations**
B. Yalcinkaya*, N. Lauffer*, M. Vazquez-Chanlatte, A. Shah, S. A. Seshia

Acta Inf. 2022 **An Automated System Repair Framework with Signal Temporal Logic**
M. Ergurtuna, **B. Yalcinkaya**, E. Aydin Gol

DATE 2019 **An Exact Schedulability Test for Non-Preemptive Self-Suspending Real-Time Tasks**
B. Yalcinkaya, M. Nasri, B. B. Brandenburg

FormaliSE 2019 **Clock Reduction in Timed Automata while Preserving Design Parameters**
B. Yalcinkaya, E. Aydin Gol

Honors & Awards

EECS Department Fellowship, University of California, Berkeley — 2021
Valedictorian, METU Computer Engineering (ranked 1st of 249) — 2020
5th Place, METU Computer Engineering Demo Day — 2020
Summer@EPFL Fellowship, École Polytechnique Fédérale de Lausanne — 2019
1st Place, Guided Research course (CENG 488), METU Computer Engineering — 2019
TÜBİTAK Research Fellowship, Scientific and Technological Research Council of Türkiye — 2018 – 2020

Teaching & Service

Graduate Student Instructor, CS 189/289A: Introduction to Machine Learning, UC Berkeley — Spring 2026
Graduate Student Instructor, EECS 149/249A: Introduction to Embedded Systems, UC Berkeley — Fall 2022
Undergraduate Teaching Assistant, CENG 230: Introduction to C Programming, METU — Fall 2017

Reviewer: ICML 2026, NeurIPS 2025, VerifAI Workshop @ ICLR 2025 & 2026, NeuS 2025, TACAS 2025, RV 2025, TIME 2023
Mentoring: Abhinav Pomalapally, Arya Raeesi, Pau de las Heras Molins, Rudy Colato, Emir Elzein, Marcos Gomez Villafane, Narelle Ho Kai Ting, Henry (Juncheng) He, Hanjoo (Andy) Kim, Isaac Andres Castillo

Technical Skills

Languages: Python, C/C++, C#, Haskell, Prolog, Bash, Assembly
ML & RL: JAX (Flax, Optax, Chex, Distrax, Jraph), PyTorch, DGL, Gym, scikit-learn, ONNX Runtime, Ollama
Data & Visualization: NumPy, SciPy, pandas, NetworkX, Matplotlib, Seaborn
Tools: Ray, Weights & Biases, ROS, Apache Flink, AWS (S3), Git, uv, LaTeX