Caleb Chuck

Research Scientist, Synthefy, Austin, TX, United States
calebchuck@utexas.edu — +01 (206) 218-3019 — https://www.linkedin.com/in/caleb-chuck-12486085/ — calcharles.github.io

RESEARCH INTERESTS

• Sequential Decisionmaking • Reinforcement Learning • Robot Learning • Controllability • Hierarchical Skills • Robotics • Causal Reasoning • Object-based Reasoning • Goal-based Reinforcement Learning • Model-Based Reinforcement Learning • Representation Learning • Unsupervised Reinforcement Learning • Learning from Demonstration/Observation • Time Series Foundation Models • Tabular Foundation Models

AFFILIATIONS

Princeton University, Princeton, NJ — 07/2026 — Present
Postdoctoral Researcher, Computer Science
Affliated with **Benjamin Eysenbach**
Directed Exploration for Reinforcement Learning

University of Texas at Austin, Austin, TX — 08/2017 — 12/2024
Doctor of Philosophy, Computer Science — Cumulative GPA: 3.80/4.00
Advised by **Scott Niekum, Amy Zhang, Yuke Zhu**
Thesis Title: Learning Hierarchical Factor Control through Causal Interactions

University of California, Berkeley, Berkeley, CA — 08/2013 — 05/2017
Bachelors of Science, Electrical Engineering and Computer Science — Cumulative GPA: 3.90/4.00
Honors EECS, Dean's Honors

PUBLICATIONS AND PROJECTS

Learning Object Manipulation without Rewards or Demos via Contrastive Interaction
In Submission CoRL 2026

- Tongle Shen, **Caleb Chuck**, Fan Feng, Biwei Huang

Unsupervised RL: A Survey and Unified Perspective
In Submission JMLR 2026

- **Caleb Chuck***, Siddhant Agarwal*, Harshit Sikchi*, Jiaheng Hu, Max Rudolph, Scott Niekum, Peter Stone, Amy Zhang

Behavior-Centric Representations for Reuse in Hierarchical RL
In Submission NeurIPS 2026

- Sarthak Dayal, Abhinav Peri, Carl Qi, **Caleb Chuck**, Amy Zhang

MuSED-FM: A Benchmark for Evaluating Multivariate Time Series Foundation Models
Submitted ICLR 2026

- **Caleb Chuck**, Sai Shankar Narasimhan, Shubhankar Agarwal, Aditya Narayanan, Avantika Gupta, Raghav Mallampalli, Fan Feng, Anthony Bao, Jeffrey Lai, William Gilpin, Sandeep Chinchali, Sujay Sanghavi

A Unified Framework for Unsupervised Reinforcement Learning Algorithms
Accepted Reinforcement Learning Conference 2025 Workshop

- **Caleb Chuck***, Siddhant Agarwal*, Harshit Sikchi*, Jiaheng Hu, Max Rudolph, Scott Niekum, Peter Stone, Amy Zhang

RL Zero: Direct Policy Inference from Language Without In-Domain Supervision
Accepted ICLR Workshop 2025, Best Poster; Accepted NeurIPS 2025

- **Caleb Chuck***, Harshit Sikchi*, Siddhant Agarwal*, Pranaya Jajoo*, Samyak Parajuli*, Max Rudolph*, Peter Stone, Amy Zhang, Scott Niekum

Null Counterfactual Factor Interactions for Goal-Conditioned Reinforcement Learning
Accepted ICLR 2025

- **Caleb Chuck***, Fan Feng*, Carl Qi*, Chang Shi*, Sara Maglicane, Amy Zhang, Scott Niekum

A Dual Approach to Imitation Learning from Observations with Offline Datasets
Accepted CoRL 2024

- **Caleb Chuck***, Harshit Sikchi*, Amy Zhang, Scott Niekum

SkiLD: Unsupervised Skill Discovery Guided by Local Dependencies
Accepted NeurIPS 2024

- **Caleb Chuck***, Zizhao Wang*, Jiacheng Hu*, Roberto Martin-Martin, Amy Zhang, Scott Niekum, Peter Stone

Task Agnostic Representation Learning Using Control-based Equivalence Metrics for Reinforcement Learning
Accepted RLC 2024

- **Caleb Chuck***, Max Rudolph*, Kevin Black*, Amy Zhang, Scott Niekum

Robot Air Hockey: A Manipulation Testbed for Robot Learning with Reinforcement Learning
Accepted ICRA Workshop 2024

- **Caleb Chuck***, Carl Qi*, Michael J. Munje*, Shuozhe Li*, Max Rudolph*, Chang Shi*, Siddhant Agarwal*, Harshit Sikchi*, Abhinav Peri, Sarthak Dayal, Evan Kuo, Kavan Mehta, Anthony Wang, Peter Stone, Amy Zhang, Scott Niekum

Automated Identification of Functional Actual Causes in Complex Environments
In Submission TMLR 2026

- **Caleb Chuck***, Sankaran Vaidyanathan*, Stephen Guigere, Amy Zhang, David Jenson, Scott Niekum

Granger Causal Interaction Skill Chains
Accepted TMLR 2024

- **Caleb Chuck**, Kevin Black, Aditya Arjun, Yuke Zhu, Scott Niekum

Gaze Supervision for Mitigating Causal Confusion in Driving Agents
Accepted AAMAS 2024, IV 2024

- Abhijat Biswas, **Caleb Chuck**, Jarrett Holtz, Scott Niekum, Alessandro Allievi

ScrewNet: Category-Independent Articulation Model Estimation From Depth Images Using Screw Theory
Accepted ICRA 2021

- Ajinkya Jain, Rudolf Lioutikov, **Caleb Chuck**, and Scott Niekum

Hypothesis-Driven Skill Discovery for Hierarchical Deep Reinforcement Learning
Accepted IROS 2020

- **Caleb Chuck**, Supawit Chockchowwat, Scott Niekum

Statistical Data Cleaning for Deep Learning of Automation Tasks from Demonstrations
Accepted CASE 2017

- **Caleb Chuck**, Michael Laskey, Sanjay Krishnan, Ruta Joshi, Roy Fox, Ken Goldberg

Comparing Human-Centric and Robot-Centric Sample Efficiency for Robot Deep Learning from Demonstrations
Accepted ICRA 2017

- Michael Laskey, **Caleb Chuck**, Jonathan Lee, Jeffrey Mahler, Sanjay Krishnan, Kevin Jamieson, Anca Dragan, Ken Goldberg

Robot Grasping in Clutter: Using a Hierarchy of Supervisors for Learning from Demonstrations
Accepted CASE 2016

- Michael Laskey, Jonathan Lee, **Caleb Chuck**, David Gealy, Wesley Hsieh, Florian T. Pokorny, Anca D. Dragan, and Ken Goldberg

Paper links found at `http://calcharles.github.io/`

RESEARCH EXPERIENCE

Synthefy Inc. Austin, Texas
Head of Research January 2025 — March 2026

- Research in Time Series Foundation Models

Safe, Confident, And Aligned Learning + Robotics Lab University of Massachusetts, Amherst
PhD Researcher September 2022 — December 2024

- Research in Causal Reinforcement learning and Actual Cause Inference
- Server and Compute Manager
- Advised by Prof. Scott Niekum

Machine Intelligence through Decision-making and Interaction University of Texas, Austin
PhD Researcher September 2022 — December 2024

- Research in Reinforcement Learning for Dynamic Systems. Compute resources manager and support.
- Advised by Prof. Amy Zhang

UT Austin Robot Perception and Learning Lab University of Texas, Austin
PhD Researcher September 2021 — December 2024

- Research in causal hierarchical reinforcement learning and scalable methods
- Advised by Prof. Yuke Zhu

Personal Autonomous Robotics Lab University of Texas, Austin
PhD Researcher September 2017 — September 2022

- Research in object-based Hierarchical Reinforcement Learning
- Continued as a researcher in SCALAR Lab working with Prof. Scott Niekum

Berkeley Automation Science Lab University of California, Berkeley
Undegraduate Research October 2015 — May 2017

- Research in learning from demonstration in robotics. Implemented and tuned deep learning architectures
- Advised by Prof. Ken Goldberg and Dr. Michael Laskey

Biochemical Lab University of Washington
Volunteer Assistant May 2013 — August 2013

- Perform E-coli-based Protein Synthesis and Titration, DNA synthesis. Maintained HEPA Cell Cultures
- Worked with Research Scientist Alan Tu

WORK EXPERIENCE

Synthefy Austin, Texas
Head of Research January 2025 — March 2026

- Multimodal, Multivariate Time series foundation model learning at scale; both data and neural network modeling.

Centro Inc San Francisco, California
Data Science Intern June 2015 — August 2016

- Worked with the design and implementation of a Bayesian click-through prediction model
- Converting prototype C++ code to production code on an Apache Spark.

Enlearn Inc Seattle, Washington
Software Engineer June 2014 — September 2014

- Developed and implemented progression tree algorithms for the generation and solving of math problems.
- Designed and produced UI's for Android tablet.

AWARDS

National Defence Science and Engineering Graduate Fellow Washington DC, United States 2017-2021
Dean's Academic Honors Berkeley, CA 2013-2017
EECS Honors Berkeley, CA 2016-2017
National Society of Collegiate Scholars Washington DC, United States 2013-2017

SERVICE, OUTREACH

Workshop Organizer RLC Reinforcement Learning Beyond Rewards 2: Ingredients for Developing Generalist Agents

Main organizer for all elements, speaker coordination, social organization, event coordination and paper review.

Workshop Organizer RLC Reinforcement Learning Beyond Rewards

Main organizer for all elements, speaker coordination and paper review.

Workshop Organizer IROS Causality for Robotics: Answering the Question of Why

Organized the application, invitation of speakers, paper review and program

SCALAR Lab Server Manager

Maintains the NFS systems, server maintenance and machine specification, admin operations for 4 GPU clusters, and setup and maintenance of Slurm-scheduled cluster.

MIDI Lab Server Manager

Admin operations for GPU clusters, ordering, and installation of NFS system.

Reviewer

Reviewed for ICLR 2023, 2024, 2025 NeurIPS 2022, 2023, 2024, 2025 CoRL 2022, 2023, 2024, 2025 ICML 2022, AAAI 2022, 2025 ICRA 2020, 2021, 2022, 2023, 2024, 2025, IROS 2022, 2023, RLC 2026

Teaching Assistant

University of California Berkeley CS16a Data Science
University of Texas at Austin CS 312 Introduction to Computer Programming
University of Texas at Austin CS 329E Elements of Data Analytics
University of Texas at Austin ECE 404J Data Science Laboratory

Discussion Leader

Directed Undergraduate Reading Group for Reinforcement Learning 2023
Directed Undergraduate Reading Group for Reinforcement Learning 2024

Undergraduate Mentor

Advising for Kevin Black (Berkeley PhD), Supawit Chockchowwat (UIUC PhD), Aditya Arjun (Industry), Misha Lvovsky (Northeastern PhD)