

ELITA LOBO

WEBPAGE: <https://elitalobo.github.io>

elobo@umass.edu

[LinkedIn](#) ◇ [Github](#) ◇ [Google Scholar](#)

RESEARCH INTERESTS

I design algorithms to address robustness and fairness challenges in reinforcement learning, resource allocation, and large language models.

Machine Learning, Reinforcement Learning, Convex Optimization, Natural Language Processing

EDUCATION

University of Massachusetts Amherst

Sept '22 - Current

PhD Computer Science

(advised by [Yair Zick](#))

University of Massachusetts Amherst

Sept '18 - Dec '20

Masters in Computer Science

Advisors: Dr. Marek Petrik , Dr. Shlomo Zilberstein

National Institute of Technology Durgapur

July '12 - Jun '16

Bachelor of Technology (BTech) in Electronics and Communication Engineering, National Institute of Technology Durgapur

(Advisors: Dr. Rajib Kar)

INDUSTRY EXPERIENCE

Amazon, Central ML Team

Jun – Aug '25

Research Intern

Seattle, WA

- Developed a hierarchical planning framework inspired by And-Or* search for long-horizon web-browsing LLM agents, enabling structured decision-making over complex, multi-step tasks. Fine-tuned a sub-planner reward model to effectively distinguish between good and bad subplans, improving intermediate decision quality and overall agent performance.

Harvard Business School

Jan – May '24

Research Intern

Cambridge, MA

- Investigated effects of finetuning on reasoning abilities of large language models (LLMs).

Microsoft Research India

Jun – Aug '24

Research Intern under Gaurav Sinha, Nagarajan Natarajan

Bangalore, India

- Investigated robustness of alignment algorithms for LLMs to noise in AI Feedback.

IBM Research

May – Aug '23

Research Intern

Yorktown Heights, NY

- Designed LLM-based methods for accurate metadata to business glossary matching.
- Proposed framework for finetuning LLMs with RLHF and contrastive loss to enhance metadata mapping accuracy.

IBM Research*Research Intern*

- Designed algorithms for efficient hyperparameter tuning in reinforcement learning.

May – Aug '22

Yorktown Heights, NY

IBM Research*Research Intern*

- Integrated Off-Policy Policy Evaluation in Automated Dynamic Optimization Framework.
- Developed a method for minimizing variance of risk estimators in RL via robust statistics (influence function).

May – Aug '21

Yorktown Heights, NY

Harvard Business School*Research Intern*

- Developed a data-poisoning attack framework to analyze sensitivity of off-policy evaluation methods.

Nov '20 – Feb '21

Cambridge, MA

Flipkart*Software Engineer*

- Built DL model to detect anomalous payouts in accounting systems.
- Developed stock ledger generator API and invoice register API.
- Provided on-call support for inventory valuation and warehouse transfer systems.

Aug '17 – Jul '18

Bangalore, India

Endurance International Group*Software Engineer*

- Created APIs for web orchestration, smart search, and session management.
- Developed ML system to detect parked domains.
- Built Imagio: a fast keyword- and color-filtered image search tool.

Jul '16 – Aug '17

Bangalore, India

Indian Institute of Science*Research Trainee*

- Integrated MSR codes into Ceph.
Link: shorturl.at/sFRW9

Sept '16 – Mar '17

Bangalore, India

HackerRank*Part-time Problem Setter*

- Authored problems for competitive programming contests.

Sept '15 – Jan '16

GoIbibo*Software Engineer Intern*

- Built ML framework to predict TTL of flight search results for caching optimization.
- Created distributed in-memory cache with bus protocol communication.

May – Aug '15

Bangalore, India

Google Summer of Code*Software Engineer Intern*

- Redesigned UI of Gnome-Calculator; implemented Keyboard Mode and History View.

May – Aug '14

RESEARCH PAPERS IN PROGRESS

Counterfactual LLM Verifiers for Math and Logic Reasoning Tasks

Elita Lobo, Shiv Shankar, Chirag Agarwal, Yair Zick

(In Progress)

PUBLICATIONS

Peer-Reviewed Publications

1. [On the Impact of Fine-Tuning on Chain-of-Thought Reasoning in LLMs](#)
Elita Lobo^{*1}, Chirag Agarwal, Hima Lakkaraju
NAACL 2025
2. [Alternate Preference Optimization for Unlearning Factual Knowledge in Large Language Models](#)
Anmol Mekala, Vineeth Dorna, Shreya Dubey, Abhishek Lalwani, David Koleczek, Mukund Rungta, Sadid Hasan, **Elita Lobo**
COLING 2024
3. [Fair and Welfare-Efficient Resource Allocation under Uncertainty](#)
Elita Lobo*, **Justin Payan*** Cyrus Cousins, Yair Zick
NeurIPS 2024
4. [On Welfare-Centric Fair Reinforcement Learning](#)
Cyrus Cousins, **Elita Lobo**, Kavosh Asadi, Michael L. Littman
Reinforcement Learning Conference 2024 (Outstanding Paper Award)
5. [Axiomatic Aggregations of Abductive Explanations](#)
Vignesh Viswanathan*, **Elita Lobo***, Yacine Izza, Gagan Biradar, Yair Zick
AAAI 2024
6. [Percentile Criterion Optimization in Offline Reinforcement Learning](#)
Elita Lobo*, Cyrus Cousins, Marek Petrik, Yair Zick
NeurIPS 2023
7. [Data Poisoning Attacks on Off-Policy Policy Evaluation Methods](#)
Elita Lobo*, Harvineet Singh, Cynthia Rudin, Himabindu Lakkaraju
UAI 2022 (Oral Presentation, Top 5%)
8. [A Metahyperparameter Tuning Framework for Reinforcement Learning](#)
Elita Lobo*, Nhan Pham, Dharmashankar Subramanian, Tejaswini Pedapati
Patent, 2023 [Link]
9. [A Novel System for Metadata to Glossary Matching in Data Lakes Using Human Feedback and Generative Models](#)
Elita Lobo*, Nhan Pham, Oktie Hassanzadeh, Dharmashankar Subramanian, Nandana Sampath Mihindukulasooriya, Long Vu
Patent (Under Review), 2024

Workshops

10. [A Hierarchical Planning Framework for LLM-based Web Agents](#)
Elita Lobo*, Frank Chen, Jingjing Meng, Yang Jiao, Nan Xi, Yan Gao
Efficient Reasoning Workshop, Neurips 2025

¹* indicates papers where I was the lead author

11. [Matching Table Metadata with Business Glossaries Using Large Language Models](#)
Elita Lobo*, Oktie Hassanzadeh, Nhan Pham, Nandana Mihindukulasooriya, Dharmashankar Subramanian, Horst Samulowitz
International Workshop on Ontology Matching, 2023
12. [Soft-Robust Algorithms for Batch Reinforcement Learning](#)
Elita Lobo*, Mohammad Ghavamzadeh, Marek Petrik
R2AW Workshop, IJCAI 2021
13. [Behavior Policy Search for Risk Estimators in RL](#)
Elita Lobo*, Yash Chandak, Dharmashankar Subramanian, Josiah Hanna, Marek Petrik
NeurIPS Workshop on Safe and Robust Control, 2021

OTHER RESEARCH PROJECTS

- **Independent Study, Perceptual Robotics Lab, UMass Amherst** *Jan 2019 – May 2019*
Designed a Hierarchical Reinforcement Learning framework to learn diverse task-level skills using deep embedded encodings.
[Link](#)
Advisor: Dr. Roderic Grupen
- **Independent Study, Center for Smart and Connected Societies, UMass Amherst** *Sep 2018 – Dec 2018*
Implemented deep learning-based forecasting for peak electricity demand days to enable peak shaving in energy grids.
Advisor: Dr. Prashant Shenoy

HONORS AND AWARDS

- Recipient of Thesis Writing Fellowship in Computer Science *2024*
\$15,000
- Recipient of Anuradha and Hanuma Kodavalla Graduate Scholarship in Computer Science *2023*
\$10,000
- Recipient of UMass Robin Popplestone Fellowship in Robotics and Artificial Intelligence *2019*
\$5,000
- 1st Place in Hackday 10 (Marketplace Category), Flipkart *2018*
- 3rd Place in ML Challenge 3, Flipkart *2018*
- 2nd Place in ML Challenge 2, Flipkart *2017*
- 1st Place in Hackathon, Endurance International Group *2017*
- 98 rank in Google Code Jam to I/O for Women *2017*
- 1st Place in Trickology (Competitive Coding), Department of MCA, NIT Durgapur *2013*
- 99 percentile in All India Engineering Entrance Exam (State Rank 9) *2012*

- State Rank 11 (99.9 percentile) in Goa Engineering Entrance Exam

2012

TEACHING / MENTORING / OTHER SERVICE

- **Teaching Assistant**, University of Massachusetts Amherst
Assisted in instruction and grading for the following Computer Science courses:
Operating Systems (CS377, Fall 2018), Reasoning under Uncertainty (CS240, Spring 2019), Numerical Optimization (CS590OP, Fall 2019), Convex Optimization (CS690OP, Spring 2020)
- **Reviewer for Peer-Reviewed Venues**
UAI 2023, ICML 2023, ICML 2024, ICLR 2024, JAIR 2024
- **Mentoring and Outreach**
 - Mentor, PhD Applicant Support Program (2023, 2024)
Supported underrepresented students with graduate school applications
 - Mentor, CS 696DS Industry Mentorship Program, UMass Amherst (2024, 2025)
Guided masters students on NLP projects that lead to research publications.

RELEVANT GRADUATE COURSEWORK

- Reinforcement Learning, Convex Optimization, Probabilistic Graphical Models, Machine Learning, Artificial Intelligence, Advanced Algorithms, Information Assurance, Advanced Machine Learning Seminar, Statistical Methods for Research, Principles of Statistical Inference, Data Mining and Predictive Analytics

SKILLS

- **Programming Languages and Tools:**
C/C++, Python, Java, GoLang, Ruby, R, MySQL, Spring Boot, ElasticSearch
- **Frameworks and Libraries:**
PyTorch, Keras, Pandas
- **Expertise:**
Machine Learning, Deep Learning, Reinforcement Learning, Large Language Models (LLMs)

REFERENCES

- **PhD Advisor:** Dr. Yair Zick
Email: yzick@umass.edu
- **Master's Thesis Advisors:**
Dr. Shlomo Zilberstein - shlomo@cs.umass.edu
Dr. Marek Petrik - mpetrik@cs.unh.edu