

# Felipe Felix Arias

felipefelixarias.github.io

felipefelixarias@gmail.com

**Background:** Software/ML Engineer specializing in GenAI and large-scale ML infrastructure. At Uber's Training team, works on job controller and infrastructure for scheduling training jobs on GPUs. Led migration of core ML job scheduler to regional federated compute, migrating thousands of pipelines with zero downtime. Built GenAI safety systems and RAG-based internal AI assistants. Main point of contact for third-party LLM fine-tuning partnerships. Research background in robotics (NSF GRF) and weak supervision. Skilled in Python, Go, Kubernetes, Ray, PyTorch, and cloud platforms. Fluent in English and Spanish.

## EDUCATION

### University of Illinois at Urbana-Champaign

2023

#### Master of Science in Computer Science

GPA: 4.0/4.0

Advisor: Nancy M. Amato

Fellowships: NSF Graduate Research (2021), Ford (Honorable Mention) (2021), GEM (2020)

### University of Illinois at Urbana-Champaign

2019

#### Bachelor of Science in Computer Science (with Honors)

GPA: 3.8/4.0

## EXPERIENCE

### Uber Technologies - ML Training Team

San Francisco, CA

#### Software Engineering

March 2024 – Present

- Led architectural migration of core ML job scheduler from zonal to regional federated compute, migrating thousands of production pipelines with zero downtime. Co-designed scheduler to enable Ray job submission on Kubernetes across federated clusters.
- Developed multi-signal distracted driving detection system combining GenAI and ML models, achieving 3.7% reduction in unsafe behaviors. Built RAG-based internal AI assistants for incident management and employee information lookup.
- Spearheading mTLS security rollout for Kubernetes batch workloads, resolving service mesh incompatibilities with PyTorch/DeepSpeed. Main point of contact for third-party LLM fine-tuning partnerships, educating engineers and optimizing costs.

### Marovi AI

San Francisco, CA

#### Founder

2024 – Present

- Built and launched Marovi Wiki, an AWS-hosted multilingual platform for AI-powered translation. Architected provider-agnostic API, agentic translation pipeline with multi-step verification, and structured document model using MediaWiki, Python, and Pydantic.
- Achieved top 10% selection twice at Y Combinator. Designed correction loop system enabling crowdsourced quality improvement through user edits, reviewer voting, and canonical translation merging.
- Implemented cost controls via caching, batching, and intelligent provider selection. Designed for enterprise deployment with on-premise capability, no vendor lock-in, and open-source glossaries.

### University of Illinois at Urbana-Champaign - Parasol Lab

Urbana, IL

#### NSF Graduate Research Fellowship - Advisor: Nancy M. Amato

2019 - 2023

- Designed and implemented deep learning models for spatial regression, enabling 5x faster multi-agent motion planning and improving collision avoidance.
- Developed and released open-source tools for simulating and analyzing indoor multi-agent navigation with real-world datasets.
- Advanced self-supervised learning frameworks to generate training data for spatial prediction tasks, supporting the development of robust navigation heuristics.

### Uber Technologies - Search Team

San Francisco, CA

#### PhD ML Engineering Internship

June 2023 - August 2023

- Led the development and optimization of XGBoost and deep neural networks for ETA prediction in driver-rider matching, with a focus on low-latency, compact model size, and high-quality candidate ranking
- Conducted ablation and feasibility studies on multiple models and devised a custom evaluation framework using accuracy metrics and Spearman rank correlation to assess both regression and ranking performance

### Stanford University - Hazy Research Lab

Stanford, CA

#### Research - Advisor: Christopher Ré

2018 - 2019

- Collaborated with Alex Ratner (now CEO of Snorkel AI) on extending Snorkel, a system for rapidly creating, modeling, and managing training data, for multi-sentence weak supervision in NLP via LSTMs
- Adapted baseline heuristics for single-sentence extraction, and achieved a 12% F1 score improvement using novel multi-sentence strategies and multi-task learning

## PUBLICATIONS

- Motion Pattern Prediction in Dynamic Environments (Master's Thesis)

**F. F. Arias**, Advisor: N. M. Amato

University of Illinois at Urbana-Champaign, 2023.

- Avoidance Critical Probabilistic Roadmaps for Motion Planning in Dynamic Environments

**F. F. Arias**, B. Ichter, A. Faust, N. M. Amato

Proc. IEEE Int. Conf. Robot. Automat., 2021, pp. 10264–10270