Antonio Guillen-Perez, Ph.D.

antonio_algaida@hotmail.com | +1-510-203-8412 | San Jose, CA

Website | LinkedIn | GitHub (HPE) / Personal | Google Scholar

Professional Summary

Ph.D. AI/ML Engineer with +6 years of experience specializing in the design, implementation, and robust deployment of Reinforcement Learning systems for autonomous systems. Proven expertise in building and training agents (PPO, SAC, TD3) in high-fidelity physics simulations (SUMO, CARLA) and custom-built, open-source environments. Architected and scaled RL training pipelines on distributed infrastructure (Ray/RLlib, >500 workers) and pioneered novel Imitation Learning techniques (LfOD) to accelerate training and improve policy robustness.

Professional Experience

Hewlett-Packard Enterprise (HPE) AI Labs   Milpitas, CA
Applied AI Research Scientist   Sep 2022 – Present

- Architected and implemented scalable ML infrastructure using Python, PyTorch, and Ray (RLlib) for large-scale distributed training of RL/MARL agents, scaling to over 500 parallel workers to accelerate data collection and model iteration cycles.
- Co-led the end-to-end design and open-source implementation of **SustainCluster** and **SustainDC**, complex Gym-compatible physics simulations for training RL agents with real-world data integration (workload traces, network costs, weather, carbon intensity).
- Engineered a complete fine-tuning pipeline using PEFT/LoRA/QLoRA to adapt LLMs (Llama 3.1) for an agentic control task, demonstrating hands-on experience in applying modern models to practical problems.
- Developed 3D CNN surrogate models (U-Net) for complex physical phenomena (CFD), achieving a >2800x inference speedup over traditional simulators and enabling rapid optimization with Genetic Algorithms.
- Contributed to research on LLM refinement (**N-Critics**, NeurIPS'23 Workshop), co-authored 20+ publications, and filed 1 patent.

Polytechnic University of Cartagena (UPCT)   Cartagena, Spain
Ph.D. Researcher in Artificial Intelligence   Sep 2018 – Jun 2022

- Designed and implemented **AIM5LA**, a latency-aware MARL system that achieved zero collisions by training agents to adapt to real-world 5G network delays in a co-simulation environment (SUMO & Simu5G), directly addressing the challenge of robust deployment.
- Implemented and rigorously tested advanced RL algorithms (TD3, LSTMs, Transformers) to create an end-to-end learning system for autonomous vehicle behavioral planning, reducing travel times by 59% in complex SUMO simulations.
- Developed and implemented **Learning from Oracle Demonstrations (LfOD)**, a novel Imitation Learning algorithm that accelerated DRL training convergence by 5-6x and significantly improved policy stability.
- Engineered systems for large-scale synthetic data generation and diverse scenario creation to train and rigorously evaluate the robustness and generalization of autonomous agent policies.

University of California, Davis   Davis, CA
Research Visitor   Jun 2021 – Jan 2022

- Developed and implemented multimodal Convolutional Neural Networks (CNNs) for processing complex biomedical sensor signals, achieving >90% accuracy in early-stage cancer detection (pub. in Nature Scientific Reports).

Education

Polytechnic University of Cartagena (UPCT)   Cartagena, Spain
Ph.D. in Computer Science (Artificial Intelligence) – *CUM LAUDE*   Sep 2018 – Jun 2022
AI for Autonomous Agent Development (Focus on RL, MARL, IL, & Simulation). Dissertation: Link

Polytechnic University of Cartagena (UPCT)	Cartagena, Spain
M.S. in Telecommunication Engineering	Sep 2016 – Dec 2017
Polytechnic University of Cartagena (UPCT)	Cartagena, Spain
B.S. in Telecommunication Systems Engineering	Sep 2012 – Jun 2016

Selected Publications

- **A. Guillen-Perez**, et al. "Hierarchical Multi-Agent Framework for Carbon-Efficient..." (**SustainCluster**) [GitHub] *AAAI, 2025 Demo.* My Role: Co-designed the hierarchical RL framework and implemented the complex, multi-data center simulation environment for spatio-temporal workload optimization.

- **A. Guillen-Perez**, et al. "SustainDC: Benchmarking for Sustainable Data Center Control." [GitHub] *NeurIPS, 2024.* My Role: Co-led the design and open-source implementation of the Gym-compatible MARL benchmark for training and evaluating advanced control agents.

- **A. Guillen-Perez**, et al. "Multi-Agent DRL to Manage Connected Autonomous Vehicles..." (**advRAIM**) *IEEE Trans. on Vehicular Technology, 2022.* My Role: Designed and implemented the end-to-end MARL system with a novel LSTM-based state encoder, reducing simulated travel time by 59%.

- **A. Guillen-Perez**, et al. "AIM5LA: A Latency-Aware DRL-Based Autonomous Intersection Management..." *Sensors, 2022.* My Role: Designed and implemented the latency-aware MARL system that achieved zero collisions by learning to adapt to simulated real-world network delays.

- **A. Guillen-Perez**, et al. "Learning From Oracle Demonstrations—A New Approach..." (**LfOD**) *IEEE Access, 2022.* My Role: Invented and implemented the LfOD framework and TD3fOD algorithm, accelerating DRL training convergence by 5-6x.

Technical Skills

- **Reinforcement Learning:** Deep RL (DRL), Multi-Agent RL (MARL), Hierarchical RL (HRL), Imitation Learning (LfD, LfOD), Model-Free (PPO, SAC, TD3), Model-Based Concepts (via Surrogate Modeling), Multi-Objective Optimization, Policy Optimization, RLHF Concepts.

- **ML Engineering & Infrastructure:** Distributed Deep Learning (**Ray: RLlib, Tune**), Scalable ML Pipelines, PyTorch (Expert), TensorFlow (Familiar), MLOps Concepts, LLM Fine-Tuning (PEFT, LoRA/QLoRA).

- **Simulation & Data:** Environment Design (**SustainCluster, SustainDC**), Gym/Gymnasium, Digital Twins, PettingZoo, Stable Baselines3, **Physics Simulation** (SUMO, CARLA), Synthetic Data Generation, Real-World Data Integration.

- **Core Software & Tools: Python** (Expert), NumPy, Pandas, Scikit-learn, Git, Docker, Linux, HPC Environments.

Awards & Certifications

- **Best ML Innovation Award**, NeurIPS 2023 Workshop on Tackling Climate Change with Machine Learning.

- **Key Certifications:** AI Agents (Hugging Face, 2025), Reinforcement Learning From Human Feedback (DeepLearning.AI, 2025), Self-Driving Cars Specialization (U. of Toronto, 2024), Deep Reinforcement Learning Nanodegree (Udacity, 2020), Generative AI Fundamentals (Google, 2024).