

Jordy Kieto

Los Angeles, CA, 347-983-3736, jordy.kieto@gmail.com

-
- <https://scholar.google.com/citations?user=syjN7xYAAAAJ&hl=en>
 - <https://github.com/JordyKieto>
 - <https://jordykieto.tech/>
-

WORK EXPERIENCE

Sept 2026 — Present	Senior Research Scientist, Riot Games	Los Angeles, CA
	Researching world models for interactive environments, training agents with PPO over millions of imagined transitions in learned dynamics models to study scalable reinforcement learning, long-horizon behavior, and transfer to real environments. Training beyond the throughput limits of simulators.	
Nov 2025 — Sept 2026	Senior Machine Learning Engineer, Riot Games	Los Angeles, CA
	Engineered cross-game ML infrastructure and a scalable evaluation service integrated into training loops for asynchronous, per-checkpoint validation and artifact tracking. Partnered with researchers to standardize metrics and cut experimentation latency across game AI systems.	
Nov 2024 — Sept 2025	Software Engineer, Meta	Los Angeles, CA
	Developed large-scale ML infrastructure for Ads personalization, enabling rapid experimentation and deployment of ranking models. Built distributed training and serving pipelines (embeddings, two-tower, ANN, feature stores) supporting billions of impressions and thousands of concurrent experiments.	
Mar 2022 — Oct 2024	Software Engineer, Google	San Francisco, CA
	Built scalable ML pipelines for anomaly detection in YouTube revenue systems using TFX. Focused on reproducibility, monitoring, and analytics, enabling faster debugging, experiment evaluation.	
Nov 2021 — Mar 2022	Software Engineer, Uber	New York, NY
	Productionized ML models for routing and matching in a real-time marketplace. Integrated models into live systems, improving latency, utilization, and overall efficiency across millions of daily trips.	
Aug 2019 — Dec 2019	Data Scientist, PepsiCo	New York, NY
	Led an ML initiative optimizing ad spend across Amazon and Instacart, training TensorFlow models to predict campaign uplift. Built interactive dashboards to surface ML-driven budget recommendations.	

RESEARCH

Oct 2026	Learning in Dreams, Winning in Reality: A Continuous Dyna Loop for a MOBA
	Investigated transfer in long-horizon model-based RL by training PPO across 2.5B steps inside a 0.66M-param world model from 10.7k games (234 GB); using an 8-GPU A40 fleet, Dyna anchoring, and an entity debugger, demonstrated policies inherit mechanic fidelity to hit a 70.2% holdout win rate. https://arxiv.org/abs/2610.08033
Sept 2026	Neural Simulators for Interactive Environments
	Asked whether a learned world model can replace the simulator itself, not just predict its next frame. Built compact Dreamer/Genie-style world models (1.6M and 3.8M params) that run CARLA driving and a browser flight game fully autoregressively at 30 fps on a laptop CPU. https://github.com/JordyKieto/neural-simulators/
Sept 2026	Imitation Learning for Autonomous Driving in CARLA
	Studied closed-loop driving limits from offline data, developing a 1.36M-param multimodal policy (temporal RGB, LiDAR, telemetry) trained on 3.3 h of expert demonstrations that achieves hours of zero-intervention driving, cross-town transfer, and robust error recovery. http://arxiv.org/abs/2609.17757

TEACHING

Feb 2019 — Dec 2021	Columbia University	New York, NY
	Led rigorous and intensive data science programs for 200+ students across multiple cohorts.	