- Jaime Ledesma -

Independent Research Engineer – Reinforcement Learning & AI Safety

portfolio: https://twoquarks.com
Email: mind@consciousnesslabs.com
GitHub: https://github.com/Jaime2pb3
LinkedIn: https://www.linkedin.com/in/jaimedeepmind/
Location: Mexico (Remote).

Profile

Independent research engineer focused on adaptive intelligence, reinforcement learning dynamics, interpretability, and experimental AI safety. I build reproducible RL frameworks, run systematic training experiments on consumer GPUs, analyze failure modes, and design novel algorithmic variants (e.g. Levo-HF, isomeric Q-modulation) to study robustness under non-stationary and corrupted conditions.

Technical Skills

Python, PyTorch, NumPy, (basic) JAX, RL pipelines, custom Q-learning variants, agent evaluation, Linux, Git, data processing, experiment logging and visualization, basic distributed training.

Core Research Work

Levo-HF (High-Frequency Reinforcement Learning): Q-learning variant with frequency-conditioned updates on a single Q-table. Used to study stability and adaptation in a 7×7 “corrupted valley” gridworld with phased reward corruption.

Isomeric Q-Modulation / LevoThinking: Ensemble of policy “isomers” sharing a base policy but differing in HF-phase modulation. Designed to probe resilience to delayed reward corruption and oscillatory perturbations.

Corrupted-Valley Environment: Custom gridworld with three phases (clean, corrupted, reshaped rewards) to evaluate how agents adapt, overfit, or fail under shifting reward landscapes and injected noise.

Projects (TwoQuarks Research Hub)

TwoQuarks.com: Technical portfolio for RL and AI safety experiments, including models, logs, plots, and analysis notebooks.

Exploratory Symbolic & Representation Prototypes: Early-stage projects (e.g. Pneuma, AMUNET, NeuroCIES) exploring symbolic resonance and signal-to-symbol mappings, informing my intuition about representation, stability, and noise.

Experience

Independent Research Engineer — 2024–Present

Built and iterated custom RL agents, logging convergence behavior under multiple noise regimes and reward-corruption phases.

Designed and executed reproducible ML experiments on RTX 3060 hardware, including multi-run aggregation and time-series analysis.

Structured research roadmaps, documentation, and baselines to compare algorithmic variants and ablation results.

Software Developer (Freelance)

Developed production-grade tools (automation, dashboards, small-scale distributed scripts, and UI components) for real clients, integrating reliability, maintainability, and clear documentation.

Education

B.S. in Software Engineering (in progress) — UnADM, Mexico
Self-directed AI/ML education: reinforcement learning, optimization, and interpretability literature.