Mateo Robles

Software Engineer | Machine Learning / Research

Santa Clara, CA · syberlabs.software@gmail.com · 925-949-3105 · syberlabs.space · github.com/SyberLabs · linkedin.com/in/mateo-robles-71260b189

SUMMARY

Recent Computer Science graduate (Santa Clara University, June 2026, emphasis in Data Science) building empirical software across applied ML, LLM generation, and simulation systems. Through SyberLabs, my one-person research studio, I turn ambitious, hard-to-test ideas into working software artifacts: evaluation discipline, falsifiable claim boundaries, and clear reporting that keeps results, prototypes, and art distinctly labeled.

EDUCATION

Santa Clara University: B.S. Computer Science (emphasis in Data Science), June 2026

Relevant coursework: Machine Learning · Artificial Intelligence · Algorithms · Data Structures · Programming Languages · Automata & Theory of Computation

TECHNICAL SKILLS

Languages: Python, TypeScript, JavaScript, SQL, HTML/CSS

ML / Data: LightGBM, scikit-learn, pandas, NumPy: leakage-aware feature engineering, staged ablation, controlled evaluation, LLM generation & activation steering, experiment tracking

Software: React, Next.js, FastAPI, WebSockets, Streamlit, Canvas, REST APIs, dashboards

Practices: controlled experiments, falsification & negative-result analysis, bootstrap / uncertainty reporting, reproducible pipelines, technical writing (LaTeX)

EXPERIENCE

SyberLabs: Founder & Independent Researcher · Santa Clara, CA · 2026-present

Self-directed research studio spanning applied ML, complex-systems simulation, LLM generation, and generative art. Responsible for the full lifecycle of each studio project below: problem framing, system design, implementation, evaluation, and documentation.

SELECTED PROJECTS

SkyPredict V2: Leakage-Aware Flight Delay Prediction *(Santa Clara University · CSEN 140)*

Applied ML pipeline predicting U.S. domestic flight delays before departure over 6.96M BTS 2024 records, built under strict pre-departure (gate-close) observability constraints.

- Engineered a three-stage feature pipeline (flight metadata → backward-joined hourly weather → 3-hour rolling airport network-pressure) with whitelist/blacklist filtering and runtime leakage-guard assertions.
- Trained LightGBM models on chronological splits; on a ~1.16M-flight held-out Nov-Dec 2024 window, reached ROC-AUC 0.713, PR-AUC 0.349 (at validation prevalence), and 19.23-minute regression MAE.
- Identified network arrival-pressure features as the largest incremental lift over metadata and weather (the project's strongest defensible claim), and shipped a Streamlit dashboard reporting both results and limitations.

Grokking & Scaling: Empirical ML Research Package *(SyberLabs)*

Empirical study of delayed generalization ("grokking") in modular-arithmetic learning dynamics, built to test scaling-law and effective-theory claims rather than confirm a single flattering curve.

- Ran fit competition across candidate scaling laws over 21 scaling points (19 measured), with bootstrap intervals and deliberately harsh pooled fits to test whether one shared exponent pair survives a mixed dataset.
- Found regime dependence rather than architecture-independent universality: published MLP anchors behaved differently from local MLP and residual ladders.
- Treated negative and boundary-setting results as part of the output; wrote up findings in a LaTeX paper separating a promising latent-variable theory from unsupported universal claims.

Dreamer v2: Semantic-Dynamics Generation Engine *(SyberLabs)*

Experimental text-generation engine that inserts a testable intermediate representation between a prompt and generated language, keeping the decoder as an explicit, swappable boundary.

- Built a pure engine path that runs without API keys, plus optional prompted and activation-steering decoders.
- Ran steering diagnostics that surfaced shared-direction collapse, then mitigated it with decorrelation and query-dependent seeding.

Dynamic Kernel: Non-Stationary Routing Simulation Platform *(SyberLabs)*

Full-stack platform for routing over weighted graphs whose costs change mid-traversal, with experiment manifests and guardrails around claim scope.

- Built a NumPy-vectorized transition kernel and population simulator with domain presets (mall, airport, museum, supply chain).
- Exposed diagnostics and live simulation streams through a FastAPI + WebSocket backend feeding a React visualizer.
- Backed by 41 passing tests plus falsification reports that actively try to break the routing logic, separating validated behavior from exploratory theory.

Vital Language: Controlled LLM Generation Study *(SyberLabs)*

Controlled study testing whether chaotic logit modulation could make language-model prose feel more "alive," designed so that a negative result would still be informative.

- Matched chaotic modulation against OU-noise, white-noise, sampling, and prompt-agency controls.
- Found that coherence scaffolding mattered more than token-level turbulence: preserving the negative result as the finding rather than discarding it.