

# HASSAM GANI

+1 (929) 823-5143 | gani.hassamabdul@gmail.com | linkedin.com/in/hassamgani | github.com/HassamGani | hassamgani.com

## EDUCATION

### Columbia University

Bachelor's, Computer Science

September 2024 – May 2028

- **Relevant Coursework:** Advanced Reinforcement Learning, Machine Learning, Natural Language Processing, Artificial Intelligence, Data Structures & Algorithms in Java, C++ Deep Dive, and Advanced Programming in C

## SKILLS

**Programming:** Python, C++, C, Swift, SQL, Java, TypeScript, Bash, Git, Docker, Linux

**Machine Learning:** PyTorch, veRL, CUDA, vLLM, LORA/PEFT, Modal, FSDP, Ray, WandB, Hugging Face, Pandas, OpenCV, Numpy, Matplotlib, ChromaDB, Scikit-learn, LangChain, Jupyter Notebooks, RL, DL, HPC

**Languages:** Spanish, Portuguese, Arabic, English

## PROFESSIONAL EXPERIENCE

### MorningsideMunch: Columbia-Barnard Dining Halls - [morningsidemunch.com](https://morningsidemunch.com)

Co-Founder & CEO

New York, NY

Jan 2026 – Present

- Scaled the dining hall app to **1,000+ users**, over 30% of Columbia's dining-plan holders.
- Leading partnership negotiations with Columbia IT, Marketing & Dining to become the official campus dining app, then expand to more universities and into student-run deliveries.
- Built a Cloudflare-bypass pipeline, headless Chromium in Vercel lambdas behind a Redis lock feeding a multi-tier cache and hand-rolled HTTP/2 APNs, cutting per-refresh load **from 20+ requests to 1** across 21 halls.

### Columbia Data, Agents, and Processes Lab (DAPLab)

NL2SQL Researcher

New York, NY

May 2026 – Present

- Authoring an **ICLR submission** showing that a graph-based semantic layer over raw schemas makes LLM database agents beat SQL-agent, Spider-Agent, and ReFoRCE on BIRD and LiveSQLBench.
- Wrote 9K+ lines of Python spanning auto-generated PuppyGraph schemas over 58 Postgres databases, a Cypher ReAct agent with a matched SQL twin, and a GPT-5 harness over 1,350 tasks; ported ReFoRCE and Spider-Agent to Postgres.
- Raised graph-agent accuracy **8x (3.9% to 32.2%)** through engine root-cause fixes, then audited 240+ failures to surface corrupted ground truths and 3 artifact classes behind the paper's exclusion protocol.

Reinforcement Learning Researcher & Engineer

March 2026 – Present

- Building a self-improvement framework (targeting AAI or ICLR) that turns an LLM's own errors into training data via challenger-solver-critic GRPO self-play, lifting Qwen3-8B math **+3.8 avg across 4 benchmarks**, including OlympiadBench (+9.1).
- Reproduced R-Zero (ICLR 2026), **beat its published gain (+4.82 vs +2.79)** across 4 models and 7 benchmarks (2,866 problems), and showed that its reported Qwen3-4B +6.49 drops to +0.89 under a proper baseline.
- Implemented veRL GRPO self-play with per-role LoRA adapters, a co-evolving critic scored on predicting frontier-band difficulty, and preemption-safe Modal H100 pipelines over 30+ runs and ablations.

### Columbia Software Systems Laboratory (SSL)

Machine Learning Engineer & Researcher

New York, NY

May 2025 – Aug 2025

- Built cpyAnalyzer, mapping class-hierarchy dependencies to catch malicious transformations in pickle-based Python/CPython models, driving policies that loaded 79.8% of benign models while **rejecting 100% of malicious ones**.
- Shared 10+ known-vulnerable Hugging Face model implementations with 7 PhD students via Docker and GitHub.
- Trained 10+ custom pickle-based classifiers across tabular, image, audio, and text for vulnerability testing.

## PROJECTS & OUTSIDE EXPERIENCE

### xPredict—Live Prediction Markets on X - [github.com/HassamGani/xai](https://github.com/HassamGani/xai)

Palo Alto, CA

- Won **Honorable Mention at the xAI Hackathon** (Grok x X API track) for a live prediction market pricing X-post outcomes via a temperature-scaled softmax over log-odds with exponential decay and EMA smoothing.
- Trained 2 LightGBM models on resolved-market truth (47+30 features, leakage-free TimeSeriesSplit/GroupKFold CV) to correct engine prices in logit space, served via Dockerized FastAPI with a versioned registry.
- Shipped 12K lines of TypeScript/Python across 4 services **in under 24 hours**, from X API v2 filtered-stream ingestion with Grok-generated market rules to 5-D post scoring and Supabase Realtime updates.

### FaceGuard—FaceID for MacOS - [github.com/HassamGani/FaceGuard](https://github.com/HassamGani/FaceGuard)

New York, NY

- Built a lightweight OpenCV face-auth tool that trains on the owner's face **in 6 seconds without a GPU**, storing it as a persistent averaged 128-column numpy array.
- Auto-runs on wake via sleepwatcher to scan for intruders, emailing the 3 most-confident non-matches plus a clip when the owner isn't detected.
- Survives 100% of forced terminations by saving each frame mid-run and sending the alert on next boot.
- Identifies the owner within 2 feet 100% of the time (0.75 threshold), **65% faster** after tuning top-k from 5000 to 50. Demo on LinkedIn; open-source.

### SmartSnake—a Self-Playing PyTorch Snake Game - [hassamgani.com/projects/snake](https://hassamgani.com/projects/snake)

New York, NY

- Built a deep Q-learning agent in PyTorch that plays Snake, an 11-256-3 network reading hazards, heading, and food to pick relative moves in a live Pygame training loop.
- Stabilized training with experience replay (100K buffer, 1K minibatches), one-step temporal-difference targets, and a decaying epsilon-greedy policy.
- Added best-model checkpointing and live score plots, hitting a **733% gain by game 100** before plateauing near an average score of 31.