Nathanael Lara

nathanael.a.lara@gmail.com • (210) 454-5940 • linkedin.com/in/nathanael-lara

EDUCATION

Fordham University, Gabelli School of Business — **New York, New York**
Master of Science, Artificial Intelligence in Business — *May 2026*

New York University, Tandon School of Engineering — **New York, New York**
NYU Tandon Bridge - Computer Science Certificate — *In Progress*

The University of Texas at San Antonio — **San Antonio, Texas**
Bachelor of Business Administration, Finance — *May 2025*

PROFESSIONAL EXPERIENCE

BCG X (Boston Consulting Group) — **New York, New York**
Delivery AI Product Analyst — *Sept. 2026 – Present*
AI Product Builder (Internship) — *June 2026 – Aug. 2026*

- Owned the web product for the CFS Builder Toolkit, a platform of reusable AI components for BCG's Corporate Finance & Strategy practice - moved it from a standalone prototype into the practice's production platform (Next.js/TypeScript) and live via Azure CI/CD, cut the interface from ten tabs to three, and fixed the top stakeholder complaint by saving each user's work across reloads and re-logins
- Secured the platform for launch - found and fixed an authentication-bypass vulnerability in code review and added a regression test so it stays fixed, built an allowlist with a request-access flow for ~170 users, and put 13 API routes behind separate maintainer and admin permissions
- Kept the platform within a market-data contract that permits only human-approved use - one fail-closed, code-level control stops the AI model from calling the vendor on its own across the web app, hosted MCP connector, and desktop client - then found a second agent that could still reach the data and flagged it before credentials went live
- Built the platform's first quality rubric and automated evaluation harness (TypeScript/Vitest, grown from 27 to 120 tests) - its ground-truth benchmark showed retrieval at ~15% on real builds versus 71% on synthetic prompts and caught broken cross-references in 37% of the catalog; then raised retrieval to ~90% recall by replacing keyword matching with a tuned BM25 model
- Replaced data integrations duplicated across internal tools with seven shared bring-your-own-key connectors (SEC EDGAR/XBRL, FRED, Statista, Crunchbase, Inven, and two BCG-proprietary sources built from scratch), each tested against its live API, and showed that keys stored server-side can't be extracted even when the model is told to leak them

OCI Group Consulting, LLC — **San Antonio, Texas (Remote)**
Independent Contractor - AI & Content — *May 2026 – Sept. 2026*

- Designed and led an AI training for the firm's managing partner and five other partners, who adopted Claude, ChatGPT, and Gemini in their work to automate repeatable deliverables such as slide decks and client drafts

PwC / Fordham Capstone — **New York, New York**
AI Agent Developer — *Jan. 2026 – May 2026*

- Automated end-to-end M&A deal monitoring across 51 live sources (SEC EDGAR, RSS, press releases) via a RAG pipeline (LangChain, Chroma, GPT-4o) with 8-K section parsing, MD5 deduplication, and daily scheduling - so the deal feed refreshes every day with no manual sourcing
- Engineered a structured deal extraction and financial enrichment pipeline (Pydantic-typed records, confidence-based SQLite upsert logic, SEC XBRL integration), cutting deal-attribute extraction from hours to under a minute and standardizing 10+ fields across inconsistent source formats

United Nations / Fordham University — **New York, New York**
AI Graduate Consultant — *Jan. 2026 – May 2026*

- Reduced case law research time from hours to minutes for UN legal officers - enabling on-demand precedent lookup across 9,000+ UNDT/UNAT cases spanning 15+ years without manual database querying
- Delivered UN-LEX, a production tool the UN Office of Legal Affairs adopted as its main way to research tribunal case history

PROJECTS

Rho Pulse — *Nov. 2026*
AI Runway Feature for Startup Banking, 1st Place - NYU PMC Case Competition (Rho Track)

- Led product on an AI feature for Rho's startup-banking platform that prices every spending approval in months of runway - modeled the case company and found 2.7 of its 3 lost months came from spending approved weeks before it was paid, so the runway check moved to the moment of approval
- Designed for trust from 120 public reviews of Rho, Mercury, Brex, and Ramp - code calculates every number, the AI only writes explanations that cite their transactions, and invoices that try to give the AI reviewer instructions are held

AML Review Console — *Aug. 2026*
Self-Improving Money-Laundering Detection - Dell x NVIDIA Hackathon (Team Lead)

- Led a four-person team that built, in eight hours, an anti-money-laundering system running entirely on one local machine (Dell Pro Max with NVIDIA GB10) so no transaction data leaves it - a detection rule that rewrites itself against an adversarial transaction generator, and a two-model review panel (Qwen, Nemotron) on every alert
- Sent a case to a human whenever the two models disagreed, decided in code; where both agreed, precision rose 5.6x (0.060 to 0.338) while keeping 95.2% of genuine cases, and the rule's own precision held on 1.44M unseen transactions (0.059 versus 0.060)

AI Blocks — *Apr. 2026*
Visual AI/ML Programming Environment, 1st Place - vibeFORWARD Hackathon

AEGIS — *Mar. 2026*
Multimodal Situational Intelligence System, 1st Place - Grayscale Hackathon

TECHNICAL SKILLS

AI/Programming: Python, TypeScript, Node, SQL, RAG, Information Retrieval (BM25), Evaluation & Benchmarking, Prompt Engineering, ML, NLP, NumPy, APIs, LLM
Frameworks/Tools: LangChain, Next.js, ChromaDB, Pydantic, Vitest, MCP, Anthropic SDK, Agent Skills, Microsoft Copilot Studio, Git
Finance/Economics: Financial Modeling, Statistical Modeling, Algorithmic Fairness, Empirical Research, XBRL/SEC Data, CRM Management
Certifications: AI Fluency Ladder: AI Multiplier (L4), Boston Consulting Group (Oct. 2026)