Carlos Abreu

New York, NY | 917-651-7293 | carlos.m.abreu5@gmail.com | linkedin.com/in/carlosabrx | github.com/carlosabrx

SUMMARY

Product Manager with 5+ years building AI and data products for financial institutions, taking them from customer discovery and prototypes through launch, adoption, and iteration. Own product strategy, roadmap, and execution for LLM-powered credit document analysis used by private credit analysts, attorneys, and institutional investors.

PROFESSIONAL EXPERIENCE

Product Manager, Covenant Review | CreditSights (Fitch Group) | New York, NY | 01/2025 – Present

- Own product strategy, roadmap and execution for Covenant Studio, an AI platform taking private-credit agreements from document intake through publication; test working prototypes using Codex, and translate financial workflows into requirements.
- Deferred narrative drafting in favor of clause-anchored extraction after a four-week evaluation with three analysts on 10 live agreements found 40% narrative rework and a 10-minute median verification per section, exceeding drafting time.
- Ran demos and staged onboarding with client services; two private credit analysts produced 3x report throughput in six weeks at 65% output acceptance (extracted terms published without edits); iterated on prompts and UI from their feedback.
- Translated asset managers' trust and privacy requirements for AI-generated research into product decisions with Legal: required attorney signoff before publication and displayed clause citations beside each extracted term for source traceability.
- Designed a document-driven entitlement model for private reports: analysts map source documents at publication; the report inherits security permissions that are governable and auditable, and source-permission changes recalculate report access.
- Traced a production-only integration failure to the publication system's service account lacking the entitlement to grant report access; corrected the deployment config and validated the end-to-end permission change with an analyst before rollout.
- Bridged research and engineering, translating specialist judgement into product logic, structured data contracts, workflows, and implementation-ready requirements; drove dependencies across ingestion, AI extraction, scoring, review, and publication.
- Reduced median receipt-to-publication time from eight hours to two in a 20-report pilot on new $50M+ private-credit deals; used pilot data evaluations, and analyst feedback to set the next-quarter roadmap and extend to high-yield bonds and leveraged loans.

Associate Director, Product Manager | Fitch Ratings | New York, NY | 06/2022 – 01/2025

- Launched the Ratings Application for 92 global ratings analysts, owning requirements, workflow design, and rollout with an 8-engineer squad; consolidated fragmented rating review and approval workflows for daily credit ratings publications.
- Translated findings from 25 analyst interviews into workflow improvements, reducing manual handoffs by 30% and achieving 85% adoption within six months.
- Turned recurring analyst tasks into reusable, versioned workflows; introduced discovery, customization, and sharing, increasing four-week repeat usage from 28% to 46%.

Marketing Data Analyst | Anteriad | Rye Brook, NY | 06/2020 – 04/2022

- Built SQL model-evaluation datasets across four quarterly snapshots, linking weighted score deciles and purchase recency to subsequent 30-day revenue; traced coverage gaps to account exclusions, channels, and identity joins.

PROJECTS

AgentTrace - LLM Evaluation & Agent Observability | github.com/carlosabrx/AgentTrace

- Built a step-level evaluation framework for a research agent across 20 10-K tasks and 60 human-labeled steps; reached 70% end-to-end task completion and turned trace failures into targeted retrieval and extraction completeness experiments.

Contract Clause Extraction Eval - CUAD | github.com/carlosabrx/contract-eval

- Benchmarked LLM clause extraction across 40 attorney-annotated contracts; an audit of 33 first-run failures found 26 were measurement defects vs. 7 true model misreads, preventing misleading accuracy claims and guiding field-specific fixes.

SKILLS

Product: Product strategy, customer discovery, rapid prototyping, product requirements, roadmap prioritization, launch and adoption, success metrics (task completion, output acceptance, time saved, repeat usage), cross-functional leadership (Engineering, Design, Legal, GTM), human-in-the-loop review, data access and permissions.

AI and technical: LLM evaluations and datasets, model feedback, precision/recall, regression testing, source traceability, structured extraction, agent observability, quality/latency/reliability/cost tradeoffs, APIs and integrations, Python, SQL, Gemini API, Cursor, Codex, Figma, Git, Supabase/Postgres.

Financial domain: Private credit, leveraged finance, credit agreements and covenants, credit ratings, investment research and diligence, 10-K and filings analysis; customers across asset management, private equity sponsors, and hedge funds.

EDUCATION & CERTIFICATIONS

B.S. Chemistry, SUNY Geneseo 2018 | Certified Scrum Product Owner (CSPO), Scrum Alliance