CV

Independent Consultant in Agentic AI, Scientific Computing, and Quantitative Systems Architecture

Contact: augusto [dot] kiel [at] gmail [dot] com

Executive Profile

Scientific Architect with a Physics background and elite tenure across Tier-1 financial institutions and scientific R&D organizations (Qontigo, SimCorp, Mercado Libre, AstraZeneca). I specialize in **Research Software Engineering (RSE)**, **High-Performance Computing**, and **Production-Grade AI Engineering**. My work focuses on translating complex mathematical models into stable production systems, ranging from **Neural Stochastic Differential Equations (NeuralSDEs)** and **Domain-Specific Languages (DSL)** to **Agentic Orchestration** and **Evaluation Frameworks (Evals)**. I help firms move research-quality models into production without losing mathematical correctness or computational performance.

My Physics degree (Licenciatura, University of Buenos Aires) taught me to find the right level of abstraction before writing a line of code. Seven years across JP Morgan, Qontigo, SimCorp, and Mercado Libre taught me what production actually costs.

Core expertise: High-Performance Computing • Neural Stochastic Differential Equations • Agentic AI platforms • Multi-asset pricing systems • Risk analytics • DSL development • Numerical optimization • Python/C#/Julia • System Architecture

Engagement Focus: Performance Optimization Sprints, Agentic Platform Integration, Knowledge Engine Architecture, Agent Evaluation Frameworks, RSE Consulting, DSL Design, Architectural Design Sprints, Specialized Corporate Training.

Consulting Services

Agentic AI Engineering

- Production-grade **Model Context Protocol (MCP)** server design and implementation
- Multi-tenant enterprise authentication architecture (OAuth 2.0 PKCE, Grant Chains, CCG)
- Competitive Intelligence Knowledge Engines: document ingestion pipelines, hybrid RAG, IR trace optimization
- Automated **Agent Evaluation (Evals)** frameworks for task accuracy, tool-calling reliability, and retrieval stability
- Multi-layer agent memory systems (episodic, semantic, working memory) and personalization engines
- LLM security architecture and threat modeling for advanced reasoning models
- Agentic cost engineering: pre-filtering, caching, and execution graph optimization

Research Software Engineering (RSE)

- Translating research-quality code into stable, reproducible production systems
- **Neural Stochastic Differential Equations (NeuralSDEs)** and physics-informed ML for scientific simulation and quantitative finance
- Standardizing analytical workflows from Jupyter notebooks to CI/CD-managed pipelines
- Reproducibility, testing, and DevOps practices for R&D and Quantitative Finance teams
- Scientific computing architecture for high-mathematical-complexity domains
- Migration and modernization of legacy analytical systems

High-Performance Computing (HPC) Sprints

- Performance audit and optimization for numerical bottlenecks
- Vectorization, caching, and parallelism strategies
- 300% performance gains achieved in production systems
- Real-time computational finance optimization
- Multi-asset class pricing and valuation frameworks

Domain-Specific Language (DSL) Architecture

- Custom language design for complex pricing models
- Interactive UI development for quantitative libraries (ipywidgets, Voilà)
- Simplifying access to sophisticated mathematical models for non-engineer users

Training & Knowledge Transfer

- High-Performance Numerical Computing with Julia workshops
- Research Software Engineering (RSE) best practices for scientific teams
- Corporate workshops on quantitative finance and computational methods
- Technical mentorship for quant and engineering teams
- Interactive educational materials using Jupyter notebooks

Education

Licentiate degree in Physics, University of Buenos Aires (2011-2017)

Thesis: Statistic Analysis and Numerical Modeling of Single Particle Trajectories: Diffusion and Confinement Mechanisms

---

Signature Case Studies

1. The Performance Optimization Sprint (Qontigo)

The Challenge: Critical risk calculations were too slow for real-time reporting due to unoptimized convertible bond pricing engines.

The Solution: Led a forensic performance audit and implemented cache optimization strategies in the core C# numerical library.

Outcome: Achieved a **300% performance gain**, enabling real-time production risk reporting and significantly reducing Azure compute spend.

2. The Scientific AI Implementation (Research/SimCorp)

The Challenge: Traditional Monte Carlo simulations for European Option Pricing were computationally expensive; standard AI lacked mathematical constraints.

The Solution: Managed research into **Neural Stochastic Differential Equations (NeuralSDEs)** using **Julia**, combining deep learning with physical laws. Parallelly developed an LLM-based RAG system for querying complex financial documentation.

Outcome: Demonstrated superior convergence speeds over traditional solvers and established a framework for “Safe AI” in financial contexts.

3. The “Quant Experience” Architecture (SimCorp)

The Challenge: Quants struggled to interact with complex underlying pricing models, leading to errors and slow iteration.

The Solution: Designed and developed a Proof-of-Concept **Domain-Specific Language (DSL)** and integrated interactive Jupyter-based UIs (Voila/ipywidgets).

Outcome: Enabled non-engineers to safely construct and test pricing logic by abstracting underlying complexity.

4. The Enterprise Agentic Platform (AstraZeneca, Oncology R&D)

The Challenge: An R&D agentic platform needed a complete external-integration surface, a measurable quality framework, and reliable knowledge retrieval — none of which existed.

The Solution: Architected production-grade MCP servers with multi-tenant OAuth patterns; built automated Evals frameworks to quantifiably measure agent task accuracy and tool-calling reliability; designed document ingestion pipelines and optimized IR traces for a Competitive Intelligence Knowledge Engine; engineered multi-layer agent memory and personalization systems.

Outcome: Delivered the full external integration layer to production; established reusable OAuth blueprints adopted for all subsequent integrations; reduced LLM token consumption by **13%+** through semantic pre-filtering and cache optimization; established an Evals suite tracking tool-calling accuracy and retrieval precision as ongoing KPIs.

Professional Experience

AstraZeneca

Senior AI Engineer (Independent Contractor) *March 2026 - Present* | Remote

Advancing the R&D Agentic Platform and strengthening internal developer enablement for agent-based workflows.

- **Enterprise Agentic Integrations**: Architected and shipped the external integration surface of the R&D agentic platform, designing production-grade **Model Context Protocol (MCP)** servers with robust multi-tenant authentication patterns (OAuth 2.0 PKCE, Grant Chains, and Client Credentials).

- **Competitive Intelligence & Knowledge Engines**: Built automated document ingestion pipelines to index enterprise sources into the LLM-Wiki/Knowledge Base. Tuned hybrid search (dense + sparse) and analyzed IR traces to cut retrieval latency and improve source hit rate.
- **Automated Agent Evaluations (Evals)**: Designed and ran automated evaluation frameworks measuring agent task-execution accuracy, tool-calling reliability, and retrieval stability across model updates.
- **Multi-Layer Memory & Personalization**: Built user content personalization engines and structured agent memory systems—episodic, semantic, and working memory layers—so agents return context-relevant responses rather than generic ones.
- **LLM Security & Threat Modeling**: Audited tool-use capabilities in advanced reasoning architectures, identifying and neutralizing critical tool-injection vectors and token-abuse vulnerabilities before execution in production.
- **Agentic Cost Engineering**: Optimized execution graphs, achieving a **13%+ reduction in LLM token consumption** through semantic pre-filtering and custom cache optimization layers.

Phorma

Co-founder & Research Software Engineer *February 2026 - Present* | Buenos Aires, Argentina

Co-founded **Phorma** with Agustín Corbat to apply *Research Software Engineering* (RSE) to R&D and Quantitative Finance teams. Scientists keep their research focus; Phorma owns the engineering execution.

- Transforming research-quality models and workflows into stable, reproducible production systems
- Designing RSE architectures for scientific simulation and quantitative finance
- Establishing engineering best practices (testing, CI/CD, reproducibility) in high-mathematical-complexity codebases

Independent Consultant

Quantitative Software & Scientific Computing *January 2026 - Present* | Buenos Aires, Argentina

Providing specialized consulting services to financial institutions and technology companies:

- Quantitative finance system architecture and development

- Scientific computing and numerical software optimization
- Technical leadership and team mentorship
- Developer experience and tooling for analytical platforms

Focus areas: Multi-asset pricing libraries, risk analytics systems, DSL development, performance optimization, Python/C#/Julia consulting.

Mercado Libre

Software Technical Lead, IT Staff / Financial Planning & Analytics *June 2025 - March 2026*

Leading technology strategy and managing 14 engineers across Financial Planning & Analytics for Latin America's largest e-commerce ecosystem.

Key Focus: Strategy, Standardization & AI Workflows

- Architecting scalable financial planning and analytics platforms
- Championing AI-assisted development workflows and clean architecture standards
- Transforming ad-hoc Jupyter analyses into CI/CD-managed production systems

Impact: 90% reduction in forecasting pipeline errors through RSE principles; 15% velocity increase across engineering team

Technologies: Go, TypeScript, Python, BigQuery, Jupyter, CI/CD, distributed systems

---

SimCorp

Lead Software Engineer, Core Analytics *March 2024 - May 2025 | 1 year 3 months*

Key Focus: Core Analytics & Quant UI

- Developed Domain-Specific Language (DSL) POC for pricing model interaction
- Created LLM-based RAG system for financial documentation Q&A
- Integrated Quant UI with Axioma Risk UI for institutional investors
- Redesigned libraries for Automatic Differentiation (AD) support

Qontigo (Axioma Risk)

Associate Principal, Core Analytics *September 2020 - March 2024 | 3 years 7 months*

@akielbowicz-qontigo

Key Focus: Core Quant Libraries & High-Performance Computing

- Designed and developed core Quant Monorepo (C#) for Analytical Libraries
- Managed NeuralSDE research internship for European Option Pricing using Julia
- Built Axioma Pricing Library (APL) from ground up with 100% accuracy
- Developed comprehensive curve construction library (rates, yields, discounts, spreads)
- Led development of interactive UI (ipywidgets/voila) for Quant library access

Impact: 300% performance gain on Convertible Bond Pricing Engine enabling real-time production calculations; "Exceptional Performance" rating (2023)

J.P. Morgan

Technology Analyst, Rates CIB *July 2018 - August 2020 | 2 years 2 months*

- Built production-grade financial reporting systems with zero-downtime requirements for Rates CIB
- Provided critical support to Rates Quant team, enhancing analytical capabilities
- Delivered infrastructure for mission-critical reporting services ensuring compliance and reliability

---

Open Source Contributor & Content Creator

Scientific Software & Educational Tools *February 2016 - Present | 9+ years*

@akielbowicz | YouTube: @SCA314 | GitHub: SCA314

- Creator of **SCA314**, an educational YouTube channel focused on software craftsmanship, scientific computing with Julia, and automated testing practices in Spanish
- Development of interactive educational materials based on Jupyter notebooks
- Contributions to scientific computing and data visualization projects
- Created **charly-vibes** (site, GitHub) as a personal R&D initiative to explore the capabilities and limits of AI-assisted coding and agentic autonomous development across platforms
- Educational content bridging academic knowledge and industry best practices

---

Teaching Experience

University of Buenos Aires

- **Professor** of Calculus and Linear Algebra, CBC Engineering (December 2020 - July 2022)
- **Teaching Assistant** on Summer Course of Optics and Thermodynamics for Biology and Geology (February 2015 - March 2015)
- **Science Communicator** at Physics Department (March 2013 - December 2014)

Southern International School

- High School teacher of Physics, Mathematics and Information Technologies (2016)

Publications

- Shared Memory Semi-Implicit Solver for Hydrodynamical Instability Processes (2023)
- Photon Counting Module based on Avalanche Photo-Diodes (2017)

Speaking & Community

Active participant in technology conferences and meetups as speaker and collaborator:

- **SciPy Latinoamérica 2022** (Argentina): Workshop presenter
- Regular speaker at Python and Julia user groups in Buenos Aires
- Represented Qontigo at ECI UBA (School of Information Sciences, University of Buenos Aires)

All talks available at: talks.saxa.xyz

Technical Skills

Programming Languages Production: Python, C#, Julia Functional/Niche Languages: F#, Clojure

Agentic AI: Model Context Protocol (MCP), LangChain, LangGraph, LlamaIndex, RAG pipelines, vector databases, AWS Bedrock, multi-agent systems, LLM Evals, enterprise OAuth (PKCE, OBO, CCG)

Research Software Engineering: Reproducible workflows, scientific Python stack, numerical methods, automatic

differentiation, stochastic differential equations, NeuralSDEs

Quantitative Finance: Multi-asset pricing, risk analytics, derivatives valuation, curve construction, model calibration

Architecture & Leadership: DSL design, API design, monorepo infrastructure, microservices, Technical Mentorship, Solution Design, Corporate Training

Tools: Jupyter, Git, Docker, Azure, GitHub Actions, Visual Studio

---

Available for consulting engagements globally (remote) and in Buenos Aires

Contact: augusto [dot] kiel [at] gmail [dot] com
