

# Daniyl Kovaliov

## Senior Full-Stack Engineer · AI & Systems

Zagreb, Croatia · daniil.kovaliov@gmail.com · github.com/LeGrin · @LeGrin

---

### SUMMARY

Senior Software Engineer with 13+ years building production systems, including AI agents, LLM pipelines, and multi-agent orchestration. Built and deployed a personal AI assistant with 30+ skills, Telegram integration, and automated task orchestration (Python + Node.js). Co-founded AI startup with DSPy self-learning pipelines and RAG. Proven product mindset: crisis turnarounds, 85% latency reduction, 12x workflow automation. Rapid tech stack transitions (2 weeks Go→Python→Kotlin).

---

### TECHNICAL SKILLS

- **AI & LLM:** Multi-Agent Orchestration, LLM API Integration (OpenAI, Anthropic, Gemini), RAG, DSPy, PydanticAI, Prompt Engineering, MCP Protocol, FAISS, pgvector, Hallucination Testing
- **Languages:** Python, TypeScript, Node.js, Go, C#/.NET, Kotlin
- **Backend:** FastAPI, asyncio, PostgreSQL, MongoDB, Redis, Kafka, Microservices, Event-Driven
- **Frontend:** React, Next.js, Chrome Extensions, Telegram Bot API
- **Cloud & DevOps:** AWS, DigitalOcean, Docker, Kubernetes, Cloud Functions, CI/CD, Trunk-Based Development
- **Practices:** Clean Architecture, TDD, E2E Testing (Playwright), Observability, Product-First Development

---

### EXPERIENCE

#### Contract Software Engineer | Nov 2025 – Present

*Self-Employed · Remote*

- AI Personal Assistant setups for clients — LLM integration, workflow automation, Telegram bots
- CRM systems development and integration (Python, React)
- Programming mentorship and technical consulting
- Stack: Python, FastAPI, React, DigitalOcean, Claude Code, LLM APIs (OpenAI, Anthropic, Gemini)

#### Co-founder & AI Engineer | Nov 2025 – Present

*Tenderium · Remote*

- Built DSPy self-learning pipeline for AI-powered tender analysis with automated prompt optimization
- Designed RAG system with FAISS/pgvector for knowledge retrieval across 10K+ government documents
- Integrated governmental procurement APIs for real-time tender monitoring
- Implemented hallucination detection and output validation for critical business decisions
- Stack: Python, FastAPI, TypeScript, React, DSPy, RAG, FAISS, pgvector, PostgreSQL

**Senior Software Engineer | Aug 2023 – Nov 2025***LetyShops · Remote*

- Architected anti-fraud detection system: 100% known fraud + 7% new suspicious accounts (12% total detection in cashback domain)
- Delivered 7 CI/CD pipelines across teams (Kotlin, Python, Android), 50-70% faster through caching and optimized Docker builds
- Owned end-to-end delivery: trunk-based development with dozens of daily deploys
- Configured Amplitude and Segment event tracking across client applications, enabling data-driven product decisions
- Deployed Cloud Functions for event processing and automation
- Established SRE practices: Datadog monitoring, alerting, and full-stack observability (metrics, tracing, logs) across the platform
- Rapid tech stack mastery: Go→Python→Kotlin, each in 2 weeks

**Senior Software Engineer / Tech Lead | Aug 2020 – May 2023***PragmaTech · Remote*

- Built universal jQuery extension for Instagram-style stories in any web store (Shopify, custom) — optimized to single lightweight JS file working across all platforms (4 months)
- Developed reputation platform with microsites and multi-source review aggregation including Google Maps (10 months)
- Designed centralized Anti-fraud & Compliance SaaS architecture for EU/US certification automation (6 months)
- Mentored 3 engineers: 2 mid→senior, 1 junior→mid
- Automated daily researcher workflow from 2 hours to 10 minutes (12x improvement)
- Crisis rescue: Analyzed failing product over 1.5 months, identified critical technical and process issues, implemented Scrum rituals + Chrome extension architecture. Result: company survived

**Senior Software Developer | Aug 2019 – Aug 2020***Reply.io (B2B SaaS — AI-powered sales engagement platform) · Kyiv, Ukraine*

- Reduced core workflow latency from 10 minutes to 1.5 minutes (85% improvement) through smart throttling and parallelization
- Developed smart Chrome extension (UI to backend), reducing average session from 15 minutes to 5 minutes while doubling work quality
- Refactored unit test suite — 2x faster and more stable in CI
- Navigated complex legacy codebase independently — production-ready in 3 months

**Software Developer → Acting Team Lead | Aug 2018 – Aug 2019***Ciklum · Kyiv, Ukraine*

- Delivered first customer demo in 48 hours, directly securing additional project funding
- Bridged communication gaps between 3 development teams, reducing cross-team blockers
- Stepped up as Acting Team Lead for 4-person team, maintaining velocity through transition

## Software Engineer → Senior | May 2013 – Aug 2018

EPAM Systems · Kyiv, Ukraine

- Won internal hackathon — built observability system for 22-microservice infrastructure in 24 hours (user flow tracking + performance monitoring)
- Contributed to enterprise migration of startup product with 17-person team over 6 months (frontend + backend), served as architecture advisor on event-driven design
- Engineered flagship credit-origination platform: services, CI/CD, Web UI — greenfield to production (2016–2018)
- Extracted microservices from monolith, automated infrastructure checks (2016)

## PERSONAL PROJECTS

---

### S.A.G.E. Framework — Schema-Aware Guided Evolution

- AI-agent based development system with multi-agent orchestration and hierarchical agent teams — pioneered agent-to-agent collaboration before it was implemented in Claude Code natively
- Schema-aware approach with safety gates that guide AI development direction, reducing hallucinations through structured constraints
- Auto-learning from production incidents — system improves its own guardrails based on past mistakes
- Leveraging latest AI coding features: E2E-first TDD methodology, persistent memory (Serena MCP), and human approval gates at critical decision points

### KITT — AI Partner System (Production, Daily Use)

- Built production AI assistant with custom 4-layer architecture: Letta memory core, LettaBot channel fork, deterministic workflow engine, tiered human-approval tool system
- Fixed critical OAuth bugs in open-source Letta framework; engineered 780-line idempotent bootstrap with 1,100+ memory chunks migration
- Stack: Python, TypeScript, Letta, MCP, Docker, Telegram Bot API

## EDUCATION

---

Taras Shevchenko National University of Kyiv · Kyiv, Ukraine

MSc in Cybernetics (2014), Specialist (2012) | 2007 – 2014

## INTERESTS

---

- Improv theatre facilitator: 900+ hours, 200 participants, team of four facilitators
- Active in AI/LLM developer community — following latest developments and engaging with practitioners