Roman Prigodskii

Moscow, Russia • roman@prigodskii.dev • **prigodskii.dev** • github.com/romanprigodskii

Education

Letovo School Moscow, Russia | Expected Jun 2027
Grade 11 (final year): Russian national curriculum, advanced (profile) mathematics track; IB MYP in middle school
SAT Mathematics 790 (Aug 2026) • IELTS 7.0

Honours

Moscow Olympiad in Economics – Ranked 1st of 201 finalists in Grade 10 2025/26
- Nationwide olympiad: 10,000+ participants in the qualifying round; final held at 12 venues across Russia. Officially recognized on Russia's federal olympiad list.

All-Russian Olympiad in Economics – National finalist, two consecutive years 2025, 2026
- Qualified for the national final in Grade 9 (one of 91 finalists) and again in Grade 10 (one of 118); placed in the top half of the final both years.
- Russia's top national olympiad: in 2025/26, 530,000+ students entered, 9,330 advanced to the regional round (Grades 9–11), and 390 reached the final.

RANEPA Olympiad in Economics – Final-round prize-winner, Grade 10 2025, 2026
- National olympiad run by RANEPA, one of Russia's largest universities; officially recognized on Russia's federal olympiad list.

International Mathematical Modeling Challenge (IM²C) – Commendation, 7th place in the Russian national round 2026
- International team contest in mathematical modeling of real-world problems; competed with Letovo's team "Elbrus".

All-Russian Olympiad in Mathematics – Regional-round prize-winner (Moscow), Grade 9 2024/25

Leonhard Euler Mathematical Olympiad – Regional-round prize-winner, Grade 8 2023/24
- Russia's national math olympiad for Grade 8 (the All-Russian Olympiad's regional and final rounds begin at Grade 9).

All-Siberian Open Olympiad in Mathematics – Third-prize diploma, Grade 7 2022/23

Letovo School – Best MYP Personal Project in Mathematics Mar 2025
- Grade 9 project on the Collatz conjecture, written in English: derived congruence conditions (mod 4 and mod 3) that a minimal counterexample must satisfy.

Research (papers at prigodskii.dev/research)

An e-value audit of 84 pre-registered hypotheses about a deployed forecaster – sole author **Accepted at two NeurIPS 2026 workshops | Sep 2026**
- Re-tested my deployed UFC forecaster against market-implied probabilities on 1,787 bouts, re-analysing my earlier study with e-values: the model does not beat the market. Poster at **New in ML**; shorter version at **E-Values: From Statistics to ML** (both non-archival).
- Priced three routine pitfalls of ML evaluation: Benjamini–Hochberg on overlapping slices (its one favorable discovery fails e-BH), a "failed replication" that is still accumulating evidence rather than refuted, and a post-hoc threshold that survives its own search but not transaction costs.

Measure the instrument first: detection floors for model selection and for LLM judges – sole author **Working paper | 2026**
- Measured detection floors with a null lever, the identical pipeline re-run under a new seed: on my deployed forecaster it produced a bigger effect than any of 11 real candidate changes and **80%** of the effect of the only improvement ever shipped.
- On my production LLM classifier, two identical runs disagree on 5.6% of items: prompt gains under 2.9 points are indistinguishable from noise, and no number of reruns lowers that floor below 2.5.

Vertex Boxing – Vertex MMA's models ported to professional boxing **Open-source research | Jul–Sep 2026**
- Built a fundamentals model on **413,279 professional bouts** (1950–2026) to test the founding thesis that market probabilities are least accurate in regional boxing; the evidence says the opposite. Blended with the market's opening probabilities, the model passed a pre-registered 2026 holdout at $e = \mathbf{41.8}$ against a bar of 20.
- Caught target leakage in my own data (judges' IDs appear only when a bout ends on the scorecards) and made the benchmark reproducible to 16 significant digits. A prospective test, registered in public, runs to 2029.

Experience

Gluline (Apri Software L.L.C-FZ) – Co-Founder, Engineer and Head of AI UAE, remote | 2026–Present

- Co-founded the company behind **Gluline**, a messenger with a built-in AI assistant, released in ten languages on the **App Store** (19 August 2026), **Google Play** and the web; Android (Jetpack Compose), iOS (SwiftUI), web and Python/FastAPI in one repository.
- As Head of AI, responsible for everything the product does with a model, including an assistant that takes part in chats and stays off until the user turns it on. Chat content is encrypted on the device with a key for each chat.

Alfa-Bank – Intern Moscow | Jun 2026

- At Russia's largest private bank, built **Alfa-Romeo** solo, a working prototype of the bank of 2035: a SwiftUI app spanning retail, business and investing on a self-deployed NestJS and PostgreSQL backend, with live quotes from the Bank of Russia, the Moscow Exchange and CoinGecko.
- Designed its LLM copilot to turn plain-language requests into signed, single-use transaction drafts; it can't move a ruble until the user approves with Face ID.

Zacks Rank stock-recommendation system – commissioned by a private investor Remote | Jul 2026–present

- Built a pipeline that archives Zacks' full ranked screen (about 4,400 stocks) every weekday in a SHA-256 hash chain, so no past day can be quietly edited, and issues daily buy and sell recommendations; nothing goes to a broker.
- Showed that 44 sessions cannot resolve an edge (Rank 1: +8.7%/yr, 95% CI $[-48, +65]$), so the rules were frozen at their defaults before the first recommendation (1 Oct 2026); a $10,000 paper book fills every call at the next close, and a mutation test catches 23 of 24 planted look-ahead bugs.

Projects (more at prigodskii.dev)

Vertex MMA – UFC analytics platform, vertexmma.com **Next.js, Python, LightGBM | 2026–Present**

- A nine-model forecasting stack: a 118-feature winner ensemble over **LightGBM, CatBoost** and logistic regression, method and finish models, and a 10,000-run Monte Carlo simulator. English and Russian; 4,581 fighters, 8,906 bouts.
- **AUC 0.7244** on 664 bouts held out from January 2025, measured against the market's closing price, which stays ahead: the model is better calibrated and less sharp. Market prices are never a feature, and about two dozen rejected ideas are documented with their reasons.

indelible – open-source self-study tutor for Claude Code **Python | Sep 2026**

- Plans back from the learner's date, closes the theory before the drills, re-tests every skill cold two days later and brings each miss back after 1 day, 3 days, 1 week and 3 weeks. 1,000+ downloads.

Clipwell – native clipboard history for macOS and Windows **Swift, C# | Jun 2026**

- Written natively twice, in SwiftUI and in WPF, rather than once in Electron; MIT licensed, installed with one Homebrew command, no account and no network calls.

Grade 9 mathematics projects **2024–2025**

- Proved that a sausage wrapped in newspaper and cut at an angle unrolls into a sine wave, its period the circumference; on a real one, the traced edge follows a sine to within 1.8% of its height.
- Found the volume of a Turkish tea glass by fitting its profile with four parabolas and summing thin discs with the formulas for sums of powers: 216.9 cm^3, against 200 to 225 ml in a water test.

Skills

Statistics: e-values and test martingales, anytime-valid inference, walk-forward validation, calibration, power analysis
Modelling: Python, LightGBM, CatBoost, scikit-learn, NumPy, pandas, statsmodels, Monte Carlo simulation
Engineering: TypeScript, React, Next.js, Postgres, Supabase; Swift and SwiftUI, Kotlin and Jetpack Compose, C# and WPF; FastAPI, NestJS, Docker, Linux; LaTeX
Languages: Russian (native), English (working)