Paris, France · shipaleks@gmail.com · +33 6 58 16 28 66 · shipulin.md/cv · linkedin.com/in/alexeishipulin · github.com/shipaleks

Aleksei Shipulin

Head of Product Research · AI Ambassador — Yandex Search & AI

Product research lead for consumer search and AI. Finding what drives satisfaction and trust — product tests and interviews alongside trackers and market studies — and rebuilding the lab on agentic AI. 11+ years across consumer search, conversational UX and human—AI interaction.

EXPERIENCE

YANDEX

Sep 2015 — Present · Paris (since 2022)

One of the world's largest search engines · ~110M MAU.

→ Head of Product Research, Search & AI — UX + Market Research · Aug 2025 — Present

Lead a 6-person lab (3 UX, 3 market research): usability tests, interviews, surveys, trackers and ad-hoc market studies.

- **Leading the lab's move onto agentic AI, UX and market research alike.** A copilot runs qualitative studies from brief to final deck; a quant agent built by a researcher on my team, now in testing, sets quotas, briefs recruitment, writes questionnaires and builds dashboards. Next: one view joining A/B tests, analytics, trackers and UX research. ↗
- **Turned the lab's recordings into a searchable Atlas with an MCP server** for PMs and designers to test hypotheses fast. ↗
- **Showed that "clear" isn't "understood":** a regression found usage, not interface, drives clarity ratings, yet nearly a third of finance-section users rated clarity top while expecting the wrong thing — comprehension needs testing, not just satisfaction.
- **Led a 100-interview side-by-side programme against category leaders:** in finance, people moved to whoever had less friction at entry — forced sign-up and app downloads put them off more than missing features.
- **Mapping the Search ecosystem as answers move to AI:** speed for users, leads for publishers and advertisers.

→ AI Ambassador — parallel role · Jun 2026 — Present

CEO-initiated programme: AI adoption for 244 non-technical specialists in 10+ teams — marketing, design, research, support.

- Allocate frontier-model quotas and tokens and decide which projects get them; design pipelines around what teams need; help build adoption measurement toward a goal of agentic AI on 80% of working days.
- **Host a fortnightly AI demo for several hundred non-engineers**: colleagues show what they built with AI and how to try it.

→ Lead UX Researcher, Search & AI · May 2019 — Aug 2025

From 2022, embedded sole researcher in the 0-to-1 group that shipped Yandex's LLM answer mode (Alice in Search) before Google: pilot Apr 2024, main search Oct 2024. The launch contributed to ~+7% YoY growth in total search queries in 2024.

- **Studied how people choose a credit card** (12 interviews): transparent terms and an official tariff document outweigh grace-period length; first-timers want safety, some a way to build credit history; seasoned users optimise terms.
- **Studied why users distrust AI answers to medical and financial questions:** trust rests on the answer text and official sites, not on source chips; findings fed the work on choosing expert sources. ↗
- **Turned answer preferences into evaluation criteria:** blind side-by-side tests of anonymised answers (64 participants), ours vs a competitor's, showed what drives preference, from formatting and emoji to a closing follow-up question. The ML team built the criteria into model evaluation, and its scores came to match user judgments far more closely. ↗
- **Caught a metric win that wasn't a user win:** answers matching or beating DeepSeek on accuracy were too complex for teens.
- **Ad load:** irritation came from intrusiveness, not volume; frequency caps cut promo impressions, installs barely moved. ↗

Earlier: Yandex UX Lab (2015—2019) — interviews, usability and behavioural research; contributed to launching the Alice voice assistant and smart speaker line. · **Usethics** (2015) — True Intent Study. · **USABILITYLAB** (2014—2015) — Internet Banking Rank.

APPLIED AI & DATA SCIENCE

— **RAISE Summit Hackathon, Paris (Jul 2026) — team 1st place, Google DeepMind track.** FlowTwin: agentic ER capacity management on live open data from 18 Hong Kong hospitals — Gemini agents keep each patient's journey, NVIDIA Nemotron forecasts waits. Led the build end to end: architecture, my repo, the interface, the demo. Demo · GitHub

— **Grounded** — LLM coder for open-ends, checked by a second model: **7.6M responses**, **62,000+ hours** saved.

— **Telegram TON Hackathon (2024)** — team 1st place.

— **Yandex ART (2025)** — designed an implicit reaction-time test (affective priming) for gut reactions to generated images, alongside crowd ratings: **10,000+ responses**, used to train and compare model versions; built the prototype.

— **Mistral AI Hackathon, Paris (2026) — shortlisted.** Eidetic: a voice interviewer on Mistral models that rewrites its guide after every call; fine-tuned a small real-time Mistral model on several hundred of my interviews to probe like a live moderator.

Education **Lomonosov Moscow State University** · M.A. Social Psychology (cognitive science) · full-tuition scholarship · 2011—2018

Teaching Prompt Engineering programme for **1,000+ employees** · agent-coding workshops for non-technical specialists
"AI in UX Research" at Yandex Practicum, 4 cohorts

Fieldwork Russia · Qazaqstan · Türkiye · Saudi Arabia · US

Co-founder Papper OÜ (2022—2025): handwritten to-do list scanner — conceived it, led design, hired ML and iOS engineers, App Store 2023. Ad:res (2014—2016): social-network analysis for brand and civic projects, incl. Monchegorsk's revitalisation (2017).

Skills **Methods** usability testing, interviews, JTBD, CJM, Grounded Theory, ethnography, surveys & trackers, sampling & quotas, conjoint & MaxDiff, ANOVA, psychometrics, LLM-as-judge · **Data & ML** R, Python, SQL, fine-tuning, RAG & MCP · **Languages** Russian (native), English (fluent), French (A2)