

# Arseny Kravchenko

## Staff AI/ML Engineer

[me@arseny.info](mailto:me@arseny.info) | <https://github.com/arsenyinfo> | +48 503961992

15 years in tech, 10 years of practicing applied AI/ML.  
Co-author of "[Machine Learning System Design](#)" book.  
Love building core tech and dev tools.

### Job Experience:

**Neon (acquired by Databricks, remote)** 2025 - now  
Member of Technical Staff, AI Agents team

Core contributor behind [app.build](#), AI prompt-to-app generator ([report at arXiv](#))

**Instrumental (<https://instrumental.com/>, SF Bay Area / remote)** 2023 - 2025  
Staff ML Engineer → Senior Staff ML Engineer

Tech leadership and implementation for high uncertainty ML initiatives.

Key achievements:

- 30-40% latency reduction for edge inference system;
- remediated all known vulnerabilities in the ML components;
- built ML core of unannounced product from scratch.

**Ntropy (<https://ntropy.com/>, London / remote)** 2022 - 2023  
Senior ML Engineer

Tech leadership for a core ML component, including training pipelines, language models inference, and coordinating with other teams. Used Python, Rust, AWS, GCP, k8s.

Key achievements:

- scaled internal clustering pipeline for 10x throughput with the same hardware;
- built a merchant category code estimation system from scratch;
- optimized transaction categorization system (-15% error rate, -20% inference time, -70% training time);
- introduced a new role of Data Quality Engineer from scoping the role to hiring (finding candidates, designing a take home exercise, interviewing) and onboarding

**Instrumental (<https://instrumental.com/>, SF Bay Area / remote)** 2020 – 2022  
Senior ML Engineer

Developing and improving ML/CV algorithms (optimized for small imbalanced datasets) and related infrastructure using Python, Scala, TypeScript, AWS.

Key achievements:

- established the groundwork for the ML inference architecture implemented throughout the product suite, achieving a latency reduction of 30% and model flexibility;

- reduced required number of labeled samples for production-ready models from 12 to 4 in a few-shot setup;
- enhanced image alignment pipeline for 86% of samples with a custom DL-based keypoint detector and RANSAC;
- built a “physical size from photo” measurement system granting subpixel precision;

### **WANNABY (acquired by Farfetch) (<https://wanna.fashion/>) 2017 – 2019**

R&D Team Lead / Founding Engineer

Led a team of six engineers to apply deep learning, 3d geometry, and classic computer vision algorithms for real-time applications on mobile devices. Built core of required ML infra and product improvements such as ensuring robust AR performance in the context of reflecting surfaces.

The flagship product is **Wanna Kicks** - a mobile app doing 3d reconstruction on video for AR purposes, featured by Techcrunch ([1](#), [2](#)), Apple and other major platforms and media.

### **Juno (acquired by Gett) 2016 – 2017**

Data Science Team Lead

Led a team of four engineers building GIS flavored data related microservices and running analysis of product features.

### **Yandex (<https://yandex.com/company/>) 2015 – 2016**

Software Engineer (Search Distribution Tools)

Built a data-focused toolset for user acquisition needs.

### **Wargaming (<https://wargaming.com>) 2011 - 2015**

Marketing R&D Team Lead → Product Manager, User Acquisition Platform

Solving tech problems related to product growth and user acquisition.

## **Other Related Experience:**

### **Advising**

In 2017 - 2020, I've been the first advisor for Onesoil.ai - currently a series A agtech startup. As an advisor, I helped founders to hire core tech team and establish their first steps.

In 2021-2023, I acted as an AI advisor for **Smart Partnership Capital**, a VC company. As an advisor, I helped the partners to analyze companies from tech/AI perspectives.

### **Competitive machine learning**

Kaggle Master <https://www.kaggle.com/arsenyinfo>.

### **ML Community events**

During 2017-2019 I've been organizing multiple data-related community events in Minsk. The biggest one was **DataFest 2019** - a free one day conference with 21 speakers. The event was attended by 700+ researchers, engineers and other data folks from multiple countries.