

Sergey Zhitansky

AI Solutions Architect · Solo Founder

I build AI products, custom models, and business integrations.

25+ years in software · Kaggle Competitions Master · Owner & Solo Developer of Scovant

zhitansky.com LinkedIn Kaggle

Selected work

Scovant

AI Agent Readiness Platform

I own and independently developed Scovant, a platform for testing how websites work with AI agents. It brings together technical compatibility checks, agent task testing, and ongoing monitoring to help teams identify issues and assess readiness using different levels of evidence.

My role: Owner & Solo Developer.

Connected AI Agents

Business & Engineering Integrations

I independently delivered AI integrations connecting company databases, knowledge bases, Jira, GitHub, and AWS. Across these implementations, agents analyze infrastructure and fix bottlenecks, assess the value and complexity of proposed features, audit security, and identify opportunities to strengthen a company's products.

My role: Sole developer.

Selected capabilities across several client implementations.

Linguist

Enterprise Translation & Document Processing

Within a team-built platform, I independently developed the entire translation core: neural models, language-specific preprocessing, postprocessing, and document-format handling. My work spans recurrent and Transformer-based translation systems, connecting model development with enterprise document workflows across 64 translation directions.

My role: Sole developer of the translation core within a team-built platform.

AI-assisted engineering for teams

I introduced Claude Code into a product engineering team of eight developers working on eight services. Within four months, every change in every service was submitted as a reviewed change request: an automated reviewer gives written feedback within minutes, a build check runs before merge, and each service carries a guidance file and a shared library of standard commands for people and AI alike. The automated review caught two critical security weaknesses before release.

What changed in four months

- **Development time per ticket (median):** from 3.1 days to 0.9 days, the same ticket sizes, held for four months.
- **Bug fix turnaround (median):** from 11.6 days to 4.9 days.
- **First review feedback:** from hours or days to under three minutes.
- **Review coverage:** from 2 of 8 services to all 8; 290 change requests in seven weeks.
- **Adoption and cost:** 80% of developers use the assistant weekly; under $1,000 for the whole program.

Figures read from the team's Jira and GitHub, February–September 2026. The gain is correlated with the program, not yet attributed to it: weekly output stayed flat and QA became the next bottleneck.

Methods and tools: Claude Code, automated code review, CI build checks, per-service guidance files, shared command libraries, adoption surveys.

Specialized AI, built around the problem

My experience combines neural model development with the engineering needed to make models useful. Across production work and competition projects, I have worked with language, images, speech, and structured data—adapting established methods to specific tasks and constraints.

Neural machine translation

I independently developed the translation core within the team-built Linguist platform. This work includes recurrent and Transformer-based models, language-specific preprocessing, postprocessing, and document-format handling across 64 translation directions. The engineering challenge extends from preparing language data to preserving the structure of translated business documents.

Methods and tools: RNN-based NMT, Transformers, OpenNMT, Marian, language processing pipelines.

Computer vision and segmentation

My project and competition experience includes medical images, product classification, camera identification, wildlife monitoring, and image segmentation. I have worked with CNN families including VGG, ResNet, DenseNet, Inception, Xception, and InceptionResNetV2, as well as U-Net segmentation models. The work combines model adaptation, data preparation, augmentation, ensembles, and classical image processing.

Methods and tools: CNNs, U-Net, transfer learning, ensembles, TensorFlow, Keras, OpenCV.

Speech and sequence modelling

My speech recognition work has combined Mel-spectrogram and MFCC features with LSTM models. Broader competition experience includes biological signals, video understanding, demand forecasting, and other prediction tasks where the structure of the data matters as much as the model.

Methods and tools: LSTM, recurrent models, signal preprocessing, feature engineering.

Connected AI systems

I also build agents that work with company databases, knowledge bases, Jira, GitHub, and AWS. Selected client implementations support infrastructure analysis, bottleneck identification and remediation, feature assessment, and security review. I connect these capabilities to the systems and decisions they are intended to support.

Tested across domains

I'm a Kaggle Competitions Master with nearly 100 competitions and 18 medals: one gold, two silver, and fifteen bronze. This experience has exposed me to different data types, evaluation methods, and domain constraints, helping me get productive quickly when entering a new field.

Selected leadership experience

Founder
Rational Systems
A five-person team over five years building custom neural networks for healthcare, translation, retail, and security monitoring.

Head of AI Team
TRANSFOTECH
A ten-person team over three years building a retail analytics platform: unique-visitor counting and movement tracking based on appearance-based person re-identification.

VP R&D, Machine Learning Solutions
Spilna Sprava
Teams of up to 20 people over fifteen years delivering CRM systems, business portals, and turnkey cloud solutions.