

ANTON TRESKUNOV

Machine Learning Engineer (Computer Vision & Robotics) | PhD

Santa Clara, CA | +1-310-844-6745 | anton@treskunov.net | github.com/snusnumrick | linkedin.com/in/antontreskunov | anton.treskunov.net

SUMMARY

Innovative PhD-level engineer with 15+ years of delivering production-grade computer-vision and 3-D perception systems for autonomous vehicles, construction robotics, and consumer electronics. Expert in LiDAR & RGB-D pipelines, point-cloud registration, deep learning (PyTorch), and real-time C++/Python. Named inventor on 8+ U.S. patents and co-author of 26+ peer-reviewed publications. Proven record of turning research into shipped products—from identifying user pain points and rapid prototyping to user-study validation and large-scale launch (e.g., Samsung Smart-TV handwriting UI). Active open-source contributor.

CORE EXPERTISE

• 3D Computer Vision (SLAM, ICP, multi-sensor fusion) • Deep Learning (CNN, YOLOX, Transformers) • Point-Cloud Processing (PCL, Open3D) • Lidar perception • Real-time C++11/20 & Python • PyTorch, Triton Inference Server • ROS • MLOps (Docker, CI/CD, AWS/GCP) • Data-driven algorithm evaluation & visualization • Leadership & mentoring

PROFESSIONAL EXPERIENCE

Kelly Services (Nissan ADAS R&D) — LiDAR Perception Consultant | 2023 – 2024

Brought on as a specialized consultant to develop perception debugging and visualization tools for Nissan's autonomous driving team:

- **Developed interactive toolkit for visualizing multi-object tracking** on Velodyne LiDAR point clouds, enabling engineers to debug perception pipeline failures in real-world driving scenarios.
- **Created Python-based LiDAR decoder module** (released as open source) that processes raw Velodyne packet data, used by team for rapid algorithm prototyping and validation.
- **Provided expert consultation** on point cloud processing techniques and 3D perception algorithms to support core ADAS development team.

NavTrac — Computer Vision Engineer Remote | 2021 – 2022

- **Architected and deployed the “Activity Detector,”** a real-time truck-motion extraction pipeline for 24 / 7 gate cameras. Combined **YOLOX** detection, Kalman-filter tracking, visual-similarity re-ID, and deep matching in a Python/PyTorch stack served via **Triton** on AWS EKS.
- **Built a key-frame extractor that tracks moving container text**, selecting representative frames, and trimming manual review sets by an order of magnitude.

Doxel.ai — Senior 3-D Vision Engineer Redwood City, CA | 2017 – 2021

- **Invented a fiducial-marker scan-to-BIM alignment system** that detects circular and rectangular markers in 2-D projections of 3-D LiDAR scans to match them with surveyed control points; granted as **US11348322B1**. Cut manual alignment time from hours to minutes and unlocked near real-time progress monitoring on active construction sites.
- **Evolved the pipeline into HALT (Horizontal Alignment using Level-Tops)**, exploiting the natural flatness of concrete floors to register mobile-LiDAR sweeps to CAD models without markers, removing most hand-tweaking on typical floors.
- **Led the deviation-detection module** that compares aligned scans against design models to surface clashes and automate progress billing.

Google ATAP — Computer Vision Engineer Mountain View, CA | 2016 – 2017

- **Prototyped a real-time multi-image-stitching pipeline for a confidential Snapdragon-based camera device** (C++, OpenCV), informing sensor-tuning decisions for the hardware team.
- **Implemented on-device visual search using a Bag-of-Visual-Words index** to reconnect sessions after camera interruption; validated in live field tests and prepared for hand-off to the productization group.

Samsung Research America — Senior Research Engineer Mountain View, CA | 2009 – 2015

- **Conceived the handwriting-based channel-entry UI for Samsung’s 2013 Smart-TV touch remote.** Recruited a visual designer and front-end engineer to build a high-fidelity prototype and partnered with a UX researcher who ran the user study that won executive approval.
- **Led core software for the touch-remote concept end-to-end—porting prototypes to production C++, integrating with TV firmware, and coordinating launch on 30 million-plus sets worldwide.** This broader work on the touch-remote platform earned **3 U.S. patents** and Samsung’s **“Outstanding Achievement” award**.

Additional Experience

Prior roles include Computer Scientist at ICT/USC (2003-2009), where I led VR software development for exposure therapy systems deployed to 20+ VA medical centers, earning IEEE VR Best Paper award.

EDUCATION

PhD, Computer Science — Keldysh Institute for Applied Mathematics, Moscow. Dissertation: Algorithms for Automatic Visual Inspection.

Patents & Publications (most cited)

- “Dynamic Text Input via On-/Above-Surface Finger Sensing,” first inventor, US 9,430,145 (2016) — 43 citations
- “VR Exposure Therapy for Iraq-War PTSD,” co-author (lead VR software), IEEE VR 2006 — 77 citations; deployed in 20+ VA clinics

SELECTED OPEN SOURCE

- Cubie: AI voice assistant for Raspberry Pi (Python) with multimodal interaction, multi-LLM backends, and persistent memory
- LiDAR Point Cloud Decoder: Open-source Python module for Velodyne data processing (github.com/snusnumrick/lidar-decoder)