

Nir Ben Zvi

[Email](#)[Website](#)[LinkedIn](#)[GitHub](#)[X](#)[Phone](#)

Professional Experience

ML/AI Consultant and Fractional Head of Research (2021-present)

- Freelance **applied ML scientist and research lead** spanning **vision, video, language and multimodal systems**.
- Owning the **full machine learning lifecycle**, from problem definition, data and modeling through evaluation, production code and deployment.
- Working with companies ranging from early-stage startups to large public corporations, from **hands-on research to fractional research leadership**.
- Training and deploying **state-of-the-art open-source models**, including on **multi-node, multi-GPU systems**.
- Production work includes:
 - Computer Vision: **diffusion/flow models, VLMs and video models, segmentation, detection and classification**.
 - Language/LLMs: **complex agentic systems, retrieval/RAG, harness design, LLM evaluation suites and benchmarks**, and **on-prem/remote model serving**.
- **Partial list of clients:**
[Keep](#) (Acq. by Elasticsearch), [Mixtiles](#), [Simply](#), [Deci.AI](#) (Acq. by NVIDIA), [Lemonade](#) (\$LMND), [CommonGround](#) (Acq. by Apple), [DeepChecks](#), [Oddity](#) (\$ODD), [Papaya-Global](#), [Honeybook](#), [Clover Security](#), and many more.

Computer Vision Research Manager, Trigo (2018-2021)

- **Founding engineer** and the **first computer vision hire**.
- Built Trigo's first deep learning models from scratch, **running in real-time in live supermarkets**.
- Led development of **object detection and real-time tracking models** optimized for deployment.
- Scaled machine learning R&D from garage phase (<10 employees) to 100+ employees, **hiring and growing** top-tier ML researchers and engineers.
- Built **annotation pipelines, ML infrastructure, and model evaluation tools** to support large-scale CV development.

Computer Vision Research Scientist, Amazon Lab126 (2015-2018)

- Researched, trained and deployed **computer vision models for Amazon's consumer devices**, including the Echo Look and the Echo Show.
- Led research efforts for the group's **first deep-learning-based object detection models** used in production.
- Worked **end-to-end** on ML research, from whiteboard concepts to **production-grade implementations** in Amazon products.
- **Central technical reference point for computer vision research and infrastructure** within the org, evaluating emerging research, methods and tooling.
- Contributed to Amazon's (discontinued) **MXNet deep learning framework**.

Computer Vision Algorithm Engineer, Donde Search (Acquired by Shopify) (2015)

- Coded Donde's first visual search engine using **Caffe**.

Imaginer, Disney Research (2013-2014)

- Research associate under supervision of Prof. Arik Shamir.
- Researched style-transfer algorithms based on real artist data (pre-deep learning).
- Developed image contour tracking models for artistic style transfer.

Education

- **M.Sc., Computer Science, The Hebrew University of Jerusalem** (2013 - 2015)
- **B.Sc., Computer Science, The Hebrew University of Jerusalem** (2010 - 2013)
- **Minor in Visual Communications, Bezalel Academy of Art and Design** (2010-2014)

Publications

Alexander Lorbert, **Nir Ben-Zvi**, Arridhana Ciptadi, Eduard Oks, Ambrish Tyagi

Toward Better Reconstruction of Style Images with GANs

[KDD2017 Workshop on Fashion AI](#)

Nir Ben-Zvi, Jose Bento Ayres Pereira, Moshe Mahler, Jessica Hodgins, Arik Shamir

Line-Drawing Video Stylization

[Eurographics 2016](#)

Other

Teaching

- *Intro to Object Oriented Programming, The Hebrew University of Jerusalem* (2014)
- [MDLI](#) Deep Learning CS231 Community Course, Google Campus (2017-2020)

Military Service

Full military service in the elite unit of the IDF artillery corps, “Moran” (2006-2009)

- Squad-commanders course graduate.
- Reserves duty: In charge of the unit’s assessments of new recruits (“gibush”).

Code/Tools

- **10+ years of experience** writing Python research and production code on top of POSIX systems.
- Data processing; **postgresql** (inc. **pgvector**), **bigquery**, **snowflake** as well as current vector-dbs (currently in love with **Qdrant**).
- Thorough knowledge of **deep learning hardware**, compute/data pipelines and model throughput optimization (**ONNX**, **TRT**, etc.).
- Production experience on **AWS** and **GCP** for ML workloads.
- Full proficiency with common 3D, video and image editing tools (Blender, Photoshop, Lightroom, Illustrator).
- Hardware geek; spec’d (and built!) enterprise-grade GPU rackmounts for ML training/inference.

Inter/Personal

- Full English fluency; Native Hebrew speaker.
- Ex-professional Photographer.
- A generally nice guy.