# Ed Pizzi

edpizzi@gmail.com · Google Scholar · edpizzi.com

## SUMMARY

Applied AI researcher and engineer. My work has focused on computer vision: representation learning, large-scale retrieval, and image generation. Author of SSCD (CVPR 2022), now standard for reducing memorization in foundation model training, including multimodal LLMs. My strongest work has taken overlooked problems and driven them to industry-leading solutions. Returning in 2026, open to new areas of research.

## EXPERIENCE

**Facebook / Meta AI** (AML / FAIAR CV, AI Integrity, Ads GenAI) 2019–2024
*Research Engineer / Tech Lead* *Menlo Park, CA*

- Primary researcher for **SSN++**, an **image copy detection** fingerprint trained using a novel self-supervised objective. Deployed to identify misinformation in advance of the 2020 US election, now used extensively for content moderation across FB+IG. Published as **SSCD**. Rebuilt on open dataset; models released under a permissive license.
- Research lead for **VCE**, a large-scale **video copy detection** system for content moderation and copyright enforcement. Developed a lightweight fingerprint based on SSCD. Resolved project-risking blockers: salience metric for "stopword" frames that overwhelmed retrieval pipeline; designed large-scale retroactive search strategy for unusually uniform embeddings.
- Modeling lead for **text-to-image diffusion** for ads. Optimized to match the target domain of photographic ad images. Fine-tuning objective design, training pipeline optimization, negotiating approval for broader dataset use.
- **ShopNet**: a multi-task model for fine-grained **product recognition**. Co-developed training and eval framework; assembled v1 training data via partner catalog mining; created annotation pipeline. Adapted methods from related retrieval problems (GeM pooling, ArcFace) for significant accuracy gains. Extended system later published as **GrokNet**.

**GrokStyle** (acquired by Facebook) 2017–2019
*Software Engineer* *San Francisco, CA*

- Extended a simple KNN search to a flexible retrieval engine that could filter rows during the search based on flexible criteria. This resolved a common problem, where a raw KNN search does not retrieve candidates matching filter conditions.
- Implemented and launched a product recommendation system incorporating visual and interaction data.

**Leanplum** (mobile analytics and marketing automation startup) 2014-2017
*Software engineer* *San Francisco, CA*

**Google** 2005–2014
*Software Engineer* *Mountain View, CA*

- Site reliability engineering for search. Extensive work on automation, eliminating common causes of human-error outages.

## SELECTED RESEARCH CONTRIBUTIONS

Pizzi, Roy, Ravindra, Goyal, Douze. **A Self-Supervised Descriptor for Image Copy Detection**. CVPR 2022.

- SSCD is widely used to remove near-duplicates from training data for foundation models, including DINOv2 (cites this work directly; DINOv3 inherits via DINOv2), Stable Diffusion 3, and Llama 3 (both cite directly; SD3 ablations demonstrate a ~5× reduction in memorization). Deployed in production at Meta for content moderation, copyright enforcement, and identifying original content (for ranking).

**Establishing copy detection as a research area**. 2021–2024.

- Worked to engage the academic CV research community on this topic, through papers, datasets, challenges, and workshops.
- Papakipos et al. *Results and findings of the 2021 image similarity challenge.* NeurIPS 2021.
- Pizzi et al. *The 2023 video similarity dataset and challenge.* CVIU 2024.
- Lead organizer for the *Video copy detection workshop.* CVPR 2023.

Bell et al. **GrokNet: Unified Computer Vision Model Trunk and Embeddings For Commerce**. KDD 2020.

## EDUCATION

**Rice University**, B.A. in Computer Science, 2005. Houston, TX