Priya Goyal

https://prigoyal.github.io/ | priy2201@gmail.com | +1 (512)-965-0682 | US Citizen | ~73K citations

EXECUTIVE SUMMARY

Technical Leader & Researcher in Scalable Frontier Intelligence Systems

AI Research & Engineering Leader with startup founding experience, a track record of driving 0->1 breakthroughs, winning top honors like the ICCV Best Student Paper Award, receiving extensive media coverage, and scaling frontier AI systems across Google DeepMind, Meta MSL/FAIR, Datology AI, LinkedIn.

Deep technical and strategic expertise spanning across agentic search, agentic evaluation foundations, learning loops, recursive self-improvement, multimodal personalization, AI fairness, computer vision, and large-scale self-supervised learning with ~73K+ citations.

Designed and scaled an end-to-end agentic evaluation and learning loop frameworks delivering massive business value and ROI by driving up to 33% gains in KPIs across high-impact enterprise and consumer products. Deployed production AI judges at 800M+ scale while reducing evaluation costs by 10x.

Proven organizational leader in developing and taking frontier research to production, establishing standardized AI evaluation playbooks, building high-performing cross-functional teams, and aligning research initiatives with core business priorities.

EMPLOYMENT

LinkedIn CoreAI Dec 2024 - Present, New York, NY

- Founded LinkedIn Agentic Search, scaling the initiative from concept to a company-priority and 50+ person effort that defined LinkedIn's agentic AI future.
- Designed a novel Agentic Evaluation framework from scratch for a subjective problem in a cold start 0->1 setup: defined metrics, rubrics, agent-as-judge, user feedback hill climbing loop, human annotation, grounded / real synthetic data generation, multi-turn eval. The eval framework is used in every corner of Agentic Search effort and also enables all model training (RL, SFT) and critical launch-readiness assessments for Enterprise and consumer products.
- Built and productionized recursive self-improvement loops for GenAI products across LinkedIn.
- Appointed by executive leadership (CTO and Chief AI Officer) to lead NextGen company-wide AI Evaluation, authoring the standardized AI Eval playbook and driving adoption across multiple product pillars including Hiring Assistant, Feed, News, Sales Navigation Agent, Career Agents.
- Served as a top contributor (highest #PRs, ~500 PR reviews) while coordinating technical alignment and daily execution rhythms across a cross-functional team of 50+ ICs.

Datology AI, Founding Member of Technical Staff Jan 2024 - Oct 2024, New York, NY

Led research on curating high-quality training data to match or exceed baseline model accuracy in roughly half the training time, building the company's data-quality research infrastructure and driving customer collaborations that produced more efficient, higher-quality language models.

Google Deepmind Aug 2022 – Jan 2024, New York, NY

Started and led a moonshot project on multimodal personalization to uplift LLMs like Gemini with preferences derived from Google Photos, spanning fine-grained image segmentation research and bias/mitigation studies for multimodal models such as CLIP (ICLR'24).

Facebook AI Research, Meta MSL Jan 2016 – Aug 2022, New York, NY

Tech-led internet-scale self-supervised learning research, producing the "ImageNet in 1 Hour" distributed training breakthrough, RetinaNet and Focal Loss (ICCV'17 best paper), the open-source VISSL library, and computer-vision fairness benchmarks (FAccT 2022), accumulating 71K+ citations.

EDUCATION

Indian Institute of Technology (IIT) Kanpur India, 2010 - 2015

Bachelor and Masters, Mathematics and Scientific Computing

M.Sc. Advisor: Prof. Harish Karnick

Thesis Title: Study of auto-encodings in the hidden units of deep neural networks.
GPA: 9.1/10 (Ranked 2 of 50 in graduating class).

RESEARCH INTERESTS

Agentic Eval: Agent-as-Judge and Fine-grained dynamic objective rubrics for subjective problems (consumer / enterprise).
Synthetic Agentic Data: High quality grounded synthetic data generation for agentic products.
Multimodal learning: learning with many modalities including text, images.
Computer Vision / Machine Learning: learning with limited supervision (from uncurated data).
Socially Responsible AI: diagnosing and mitigating harms, model privacy and robustness.
Large scale vision models: pushing the limits in terms of data and model scale.

PUBLICATIONS

- *Clip the bias: How useful is balancing data in multimodal learnings?* - Ibrahim Alabdulmohsin, Xiao Wang, Andreas Peter Steiner, Priya Goyal, Alexander D'Amour, Xiaohua Zhai, **ICLR'24**
- *Vision Models Are More Robust And Fair When Pretrained On Uncurated Images Without Supervision* - Priya Goyal, Quentin Duval, Isaac Seessel, Mathilde Caron, Ishan Misra, Levent Sagun, Armand Joulin, Piotr Bojanowski, **arXiv 2022**
- *Fairness Indicators for Systematic Assessments of Visual Feature Extractors* - Priya Goyal, Adriana Romero Soriano, Caner Hazirbas, Levent Sagun, Nicolas Usunier, **FAccT, 2022**
- ***A Self-Supervised Descriptor for Image Copy Detection*** - Ed Pizzi, Sreya Dutta Roy, Sugosh Nagavara Ravindra, Priya Goyal, Matthijs Douze, **CVPR 2022**
- ***Fully Sharded Data Parallel: faster AI training with fewer GPUs*** - Myle Ott, Sam Shleifer, Min Xu, Priya Goyal, Quentin Duval, Vittorio Caggiano, **Facebook Engineering blog, 2021**
- *VISSL: A library for state-of-the-art self-supervised learning from images* - Priya Goyal, Quentin Duval, Jeremy Reizenstein, Matthew Leavitt, Min Xu, Benjamin Lefaudeux, Mannat Singh, Vinicius Reis, Mathilde Caron, Piotr Bojanowski, Armand Joulin, Ishan Misra, **2021**
- *Self-supervised pretraining of visual features in the wild* - Priya Goyal, Mathilde Caron, Benjamin Lefaudeux, Min Xu, Pengchao Wang, Vivek Pai, Mannat Singh, Vitaliy Liptchinsky, Ishan Misra, Armand Joulin, Piotr Bojanowski, **arXiv 2021**
- *Unsupervised Learning of Visual Features by Contrasting Cluster Assignments* - Mathilde Caron, Ishan Misra, Julien Mairal, Priya Goyal, Piotr Bojanowski, Armand Joulin, **NeurIPS 2020**
- *Scaling and Benchmarking Self-Supervised Visual Representation Learning* - Priya Goyal, Dhruv Mahajan, Abhinav Gupta*, Ishan Misra*, **ICCV 2019**
- *The next 700 accelerated layers: From mathematical expressions of network computation graphs to accelerated gpu kernels, automatically* - Nicolas Vasilache, Oleksandr Zinenko, Theodoros Theodoridis, Priya Goyal, Zachary Devito, William S Moses, Sven Verdoolaege, Andrew Adams, Albert Cohen, **ACM TACO 2019**
- ***Focal Loss for Dense Object Detection*** - Tsung-Yi Lin, Priya Goyal, Ross Girshick, Kaiming He, Piotr Dollár, **ICCV 2017 (*best student paper award*)**
- *Accurate, Large Minibatch SGD: Training ImageNet in 1 Hour* - Priya Goyal, Piotr Dollár, Ross Girshick, Pieter Noordhuis, Lukasz Wesolowski, Aapo Kyrola, Andrew Tulloch, Yangqing Jia, Kaiming He, **arXiv 2017**

Tech Reports

- Peer-to-peer insult detection in online communities - Priya Goyal, Gaganpreet Singh Kalra, 2013

WORKSHOP CHALLENGE

Facebook AI Self-Supervision Learning Challenge — ICCV, 2019
Extreme Computer Vision Workshop — *Seoul, Korea*
Organized the first ever self-supervised challenge aimed at benchmarking representation learning on a variety of tasks. Wrote the transfer learning code and set up the challenge evaluation/submission server.

INVITED TALKS

- SEER, a self-supervised billion parameter computer vision model - Facebook AI Innovation Summit, 2021
- Scaling and Benchmarking Self-Supervised learning – Carnegie Mellon University, April 2019.
- Facebook North America Women in AI Leadership Summit, 2019 – 10K attendees.
- ImageNet-1k training in 1 hour – Caffe2 Meetup, Hawaii, CVPR 2017
- ImageNet in 1 hour – Deep Learning at Supercomputer Scale, NeurIPS 2017

MEDIA COVERAGE

- TechCrunch article on ImageNet in 1-Hour.
- CNBC article on SEER (training A.I. to "see").
- NVIDIA Developer on ImageNet on 1-Hour.
- Geekwire on ImageNet in 1-Hour.
- NVIDIA Developer on Self-supervised learning beating SOTA Computer vision models.

- WIRED article on AI Teaching Itself to See With Less Human Help.
- CNET on training computers to learn like humans do.
- ImageNet in 1-Hour at NeurIPS 2017 Supercomputing workshop.

AWARDS

- Best Student Paper Award, ICCV 2017.
- Facebook Grace Hopper Scholarship, 2014 and Google Global Scholarship, 2013.
- Academic Excellence Award 2011-12 and 2012-13 – given to the top 5% of 850 students.
- Best Poster Award (amongst 60 posters) presented in SURGE, IIT Kanpur 2012.
- SURGE undergraduate research fellowship (2012), Indian Institute of Technology (IIT), Kanpur, India.
- National Talent Search Scholarship, 2008 – given to top 0.01% high school students all over India.

ACADEMIC SERVICE

Conference reviewer: Computer Vision and Pattern Recognition (2018 – 2022), International Conference on Computer Vision (2018 – 2022), European Conference on Computer Vision (2018 – 2022), International Conference on Learned Representations (2018 – 2022), International Conference on Machine Learning (2018 – 2022), Neural Information Processing Systems (2018 – 2022)

INTERNSHIPS

Facebook May-August 2015
Advisor: Kay Rottmann, Juan Miguel Pino, Language Technology Team *Menlo Park, CA, USA*
Built profanity detection classifier to detect profanity in the machine translations produced by Facebook. Built classifiers for all the languages and automated the re-training of any classifier.

Goldman Sachs May–July 2014
Advisor: Andrei Bergners, Panna Pavangadkar, Developer Practices Group *New York, US*
Built software for making the firm approved SDLC build cloud dynamic and self-aware. The product analyzes nearly 3000 product builds across the firm on a daily basis.

Yahoo! May–July 2013
Advisor: Bennet Manuel, Yahoo! Right Media Exchange *Bangalore, India*
Built visualization tool for tracking the usage of web bugs in online display ads to detect malicious vendors - helped increase the revenue of the company by curbing ban on minor defaulters.

Carnegie Mellon University Winter School Dec 2012–Jan 2013
Advisor: Prof. Carolyn P. Rosé, Department of Computer Science, CMU
Develop an interactive tutor moderated chat system for automated analysis and dynamic support for an online collaborative learning process. Proposed and implemented 2-pass architecture model for modelling the students' conversation and building an agent prototype capable of detecting distraction/ confusion among students and providing necessary support to ensure effective learning.

Summer Undergraduate Research Grant for Excellence May-July 2012
Advisor: Prof. Akash Anan, Department of Mathematics and Statistics *IIT Kanpur, India*
Developed efficient algorithms for projecting the point cloud on the projection surface. Automated the process of 3-D surface generation.