**Felix A. Brener**
New York, NY
felixbrener.com • github.com/felix-ab • huggingface.co/VisualInference

## EDUCATION

**University of Rochester** **Rochester, NY**
*B.S. in Biological Sciences: Neuroscience; B.A. in Economics* *Expected May 2027*

- **Coursework:** Advanced Computational Neuroscience, Advanced Statistical Methodology, Microeconomics, Genetics; in progress: Brain-Machine Interfaces, Data Methods for Economists, Developmental Biology.

## EXPERIENCE

**Quadrant Capital Advisors, Inc.** **New York, NY**
*Summer AI Analyst* *May–August 2026*

- Built four Claude-based research tools for equity and credit analysts, including a batch Porter's Five Forces pipeline that automated manual reviews at 99% concordance with senior analyst baselines, cutting turnaround by weeks.
- Automated quarterly research updates with a pipeline that diffs new SEC filings and earnings materials against prior snapshots and records what changed for public-equity and credit review.
- Shipped an AlphaSense-fed updater that refreshes existing analyst workbooks with source-attribution links and validation checks, and a screener that narrows the full ticker universe to a short list for deeper research.

**DriverCourseOnline** **New York, NY**
*Co-Founder / Developer* *April 2025–present*

- Built a New York State online pre-licensing course to DMV DTP-402 rules (Next.js, TypeScript, Firebase): server-enforced 270-minute seat time, identity checks down to keystroke dynamics, a three-strike lockout, and electronic completion reporting to the DMV. Awaiting DMV approval.
- Deploys gated on 70+ compliance, end-to-end, and security test suites (Docker, GitHub Actions, AWS App Runner); idempotent Stripe payments; S3 Object Lock retention; daily Firestore backups to S3.

**Mount Sinai Hospital, Abilities Research Center** **New York, NY**
*Research Assistant* *October 2022–August 2023, May–August 2024*

- Co-authored a paper on smartphone heart-rate-variability baselines for concussion in 752 adolescent athletes (*Med. Sci. Sports Exerc.*; ACSM 2025); built ML pose tracking from video and quaternion-based iPhone motion tracking.

## PROJECTS AND OPEN SOURCE

**Open-Weight Models for Apple Silicon** | *Hugging Face (10,700+ downloads), GitHub* **July–October 2026**

- Built MOSS-Image Turbo: 6× faster than the official 50 steps (12.4 min to 119 s, M4 Pro) without the stamped textures of community few-step (DMD) adapters. Traced them to repeated distilled steps mid-trajectory; fixed it with trajectory-regression training on teacher-relabelled targets; passed a pre-registered test on 32 held-out prompts.
- Ported OpenMOSS's 5B bilingual text-rendering DiT, MOSS-Image, to pure MLX, exact against official PyTorch (69.7 dB end to end in fp32; official seeds reproduce official images). The 8-bit tier peaks at 9.3 GB vs. 38.7 GB for mlx-vlm's port, with OCR recall identical to bf16. Also ported FLUX 3 Image Klein (cosine 0.9999 vs. BFL).
- Shipped both ports as pip-installable packages with command-line tools (moss-image-mlx, flux3-mlx).
- Extracted Qwen3.8-27B's multi-token-prediction head as a speculative-decoding drafter (up to 1.89× faster on coding prompts) and released 6-bit MLX weights with a full-precision vision tower; 7,700+ downloads across both.
- Distilled Ideogram 4 for MLX by folding classifier-free guidance into one branch and cutting steps: 2.9× faster at 800×992 on an M4 Pro (250 s to 86 s), quantized to 18.6 GB to run on a 48 GB Mac.

**Filmstock Responsiveness Benchmark** | *GitHub* **May–September 2026**

- Designed a full-factorial eval of how image models respond to filmstock, lighting, and scan language (288 images, 195 features; PCA, MANOVA, PERMANOVA). A September replication found the newer model 24% less responsive to filmstock, 14% to lighting, and 37% to scan language (Holm-adjusted $p < .001$). 31-page write-up.

**Gesso Club** | *Founder, gessoclub.com* **June 2026–present**

- Curate an AI-image archive that keeps each image's model, prompt, and workflow; built its Builder, which compiles camera, lens, film, and lighting choices into prompts for five model formats.

## SKILLS

- **Languages & Data:** Python (pandas, NumPy, SciPy, scikit-learn, Matplotlib, Plotly), TypeScript, SQL
- **AI/ML:** LLM evals, Claude API, PyTorch, ComfyUI, MLX (ports, quantization, distillation, speculative decoding)
- **Infrastructure:** Next.js, Node.js, Firebase, AWS (App Runner, S3, SES), Stripe, Redis, Docker, GitHub Actions