

SAMUEL KAHN

(760) 914-1717 ◇ sam@tickr.com ◇ <https://github.com/samueltkahn>

EXPERIENCE

Tickr, Inc

Chief Technology Officer & Head of AI

Dec 2025–Present

San Francisco, CA

- Own end-to-end technical direction across product, engineering, and AI/Data Science—defining roadmap, architecture, and go-to-market while personally contributing to vertical-specific frontier AI research.
- Forge strategic partnerships with data providers, labeling vendors, and AI platform partners (Anthropic, OpenAI); lead migration from proprietary to open-source models.
- Serve as the company's technical face to enterprise clients and partners, leading live demos, defending our technology, and scoping AI integrations alongside C-suite counterparts.
- Architected and launched *RiskWise*, an AI-native risk intelligence platform being licensed by some of the largest corporate insurers and Big 4 advisory firms.
- Built and led a team of nearly 20 AI scientists, ML engineers, data scientists, and software engineers, managing 10+ concurrent AI features including Risk Deep Research, Risk Query Agent, and Litigation Prediction.
- Led technical overhaul of platform infrastructure, reducing processing costs by more than 20x through coordinated engineering, R&D, and data science initiatives.
- Developed an agentic ML research system using Monte Carlo Tree Search guided experiment planning to autonomously generate hypotheses, run ML experiments, and improve model performance across evaluation benchmarks for predicting future litigation activity.

Tickr, Inc

VP, AI & Data Science

2021–2025

San Francisco, CA

- Led R&D and technical strategy for Tickr's AI platform, automating client analytics workflows and driving cost and time savings at scale.
- Built partnerships and co-sold with some of the largest risk, retail, and consumer packaged goods data providers including Circana, Nielsen, and Numerator.
- Spearheaded Tickr's pivot to AI-first products, scaling the technical organization from founding data scientist to a team of nearly 20.
- Built data labeling and evaluation pipelines for fine-tuning domain-specific RAG and chain-of-thought reasoning.
- Designed blinded experiments benchmarking LLM quality against human and competitor baselines.
- Fine-tuned Qwen 3 using PEFT on multi-GPU H100 instances, achieving superior task-specific performance vs. GPT-5 on vertical-specific internal benchmarks.
- Designed and optimized retrieval-augmented generation (RAG) pipelines using LlamaIndex, LangChain, Weaviate, OpenAI, and Weights & Biases for production inference.

Subpoena Solutions, SpeedLegal, Joinder (acq. Brightflag)

Senior AI Advisor

2019–Present

San Francisco, CA

- Serve as fractional AI lead for legal-tech startups—built the initial AI proof-of-concept at Subpoena Solutions that directly contributed to private equity acquisition and formation of a dedicated AI and engineering team; advised SpeedLegal and Joinder on AI strategy through product launches and acquisition.
- Lead ongoing AI strategy for fine-tuning domain-specific models and error analysis pipelines leading to an over 50% reduction in error over frontier labs on over 20 tasks.
- Build legal-task benchmarks for automating high-volume legal processes.
- Trained legal process classification and signature detection deep learning models.

University of California, Santa Cruz*Research Fellow*

2022–Present

Santa Cruz, CA

- Developed *NEO*, a conditional GAN for photometric super-resolution improving galaxy morphological measurements by 2–10x for ground-based telescopes. Research featured in NVIDIA Blog (April 2026).
- Built an agentic scientific discovery system for climate research—an autonomous AI agent that formulates hypotheses, designs and executes statistical experiments, performs adversarial self-checks, and iteratively refines a unified theory of Earth’s albedo stability using MCTS-guided experiment planning.
- Trained deep learning and explainable AI models on NASA Pleiades and NERSC supercomputers using Slurm.

Tickr, Inc*Lead Data Scientist → Director, AI*

2016–2021

San Francisco, CA

- Built and deployed transformer architectures, Bayesian time series models, and scalable ML infrastructure (Kubernetes, MLFlow, Snowflake).
- Engineered NLP pipelines for entity extraction, sentiment analysis, and unsupervised clustering; built data labeling systems using Amazon GroundTruth.

Market.Space (acquired by Tickr)*Data Scientist*

2015–2016

San Francisco, CA

- First data scientist; built and deployed customer-facing NLP models and production ML systems using Spacy, Scikit-learn, Flask, Docker, and AWS.

Salesforce.com*Data Scientist, Graduate Intern*

2014–2015

San Francisco, CA

- Built churn prediction models and delivered statistical analysis correlating heterogeneous data sources to revenue.

NASA Ames Advanced Studies Laboratory*Research Assistant*

2013–2014

Moffett Field, CA

- Trained recurrent neural networks to predict energy microgrid load and demand using Theano.
- Integrated microgrid sensor data with a Raspberry Pi and pushed data to a MySQL database for inference.

PUBLICATIONS, PATENTS & FEATURES

- *NEO: Photometric Super-Resolution for Improving Galaxy Morphological Measurements using Conditional Generative Adversarial Networks*. Kahn, S., Hausen, R., Bretonnière, H., Drakos, N., Robertson, B. arXiv:2604.20195 (2026).
- Featured in NVIDIA Blog: “Making Sense of the Early Universe: How AI and GPUs are helping astronomers work through unprecedented volumes of cosmic data.” (April 2026).
- *An Approach to Advance the Predictive Understanding of the Earth’s Top-of-Atmosphere Albedo Stability and Symmetry in the Anthropocene* (Accepted, Nature Environmental Communications).
- U.S. Patent: *Approaches to Predicting the Impact of Marketing Campaigns with Artificial Intelligence and Computer Programs for Implementing the Same* (110565.8005.US01).

EDUCATION

University of California, Berkeley

Masters of Information & Data Science

*June 2016***University of California, Santa Cruz**

B.S. in Physics

June 2013