Kevin E. Stone

Applied AI & Machine Learning Leader | Drug Development

New Jersey | 215-801-5684 | kevin.stone38@gmail.com

linkedin.com/in/kevinedwardstone | Google Scholar

SUMMARY

I am a scientific leader working at the intersection of life sciences and AI. Eleven years at Merck: contributions to eight approved medicines from the lab bench, four years leading the company's push into AI and machine learning for drug development, and today 13 machine learning scientists, data engineers, and chemical engineers reporting to me. I write my own algorithms, I have founded programs that went from zero to hundreds of users, and I coach scientists from new hires to 30-year veterans. I want to help get medicines to people faster.

CORE COMPETENCIES

AI & LLM Systems LLM agents and tool-use harnesses, dynamic tool discovery, adversarial grading loops, retrieval over scientific corpora, evaluation frameworks and hallucination scoring, agent-driven R&D workflows

Machine Learning & Modeling Neural ODEs and hybrid mechanistic models (PyTorch, torchdiffeq), Bayesian optimization and active learning (GPyTorch, BoTorch, deep kernel learning), chemical representation learning (SMILES transformers, graph neural networks), open-source authorship

Engineering & Data Python, PyTorch, SQL, Spark, Databricks, React, electronic lab notebook and LIMS API integration, structured scientific data products, CI/CD and developer tooling

Drug Development CMC authoring (S.2.2, S.2.6) and NDA/BLA support, process validation and PPQ, model verification and validation for regulatory use (V&V40), tech transfer and scale-up, reaction kinetics and mechanism, biocatalysis, crystallization, chromatography, upstream cell culture

Leadership Hiring, coaching, and developing machine learning, data, and chemical engineers, technology strategy and build-versus-buy, vendor and academic partnerships, executive and cross-divisional communication

PROFESSIONAL EXPERIENCE

Director | Process Modeling & Analytics | 2026 - Present

Merck & Co., Inc.

- Lead 13 machine learning scientists, deep learning researchers, data engineers, and chemical engineers building models, agent systems, and data products for biologics development
- Direct the biologics data product, foundational to a departmental target of 20% more project capacity without added headcount, and to 95% of experimental data becoming usable within a day of run
- Direct the enterprise API connector into our electronic lab notebook and LIMS, where deterministic parsers and LLMs together feed dozens of downstream applications from a mix of structured tables and free-text experiment records
- Direct five mechanistic chromatography models (MMCEX, flow-through AEX, HIC), cutting wet-lab resources up to 30% and widening qualified operating ranges up to 50%
- Author Merck's strategy for using validated models to support BLA decisions under V&V40, and lead the first regulatory-agency engagement on in-silico process characterization

Director | Data-Rich Experimentation | 2025 - 2026

Merck & Co., Inc.

- Wrote GARNET, a neural-ODE architecture in PyTorch and torchdiffeq that models mammalian cell culture from as few as thirty experiments; now used across bioprocess development to cut process-characterization workload up to 50%
- Used GARNET to design a feed strategy that raised harvest-day viability 20% and titer 15% on a pipeline program, and to model dynamic pH control, cutting bioreactor control-strategy experiments from roughly 200 to 50
- Proposed and lead chemical representation learning for Bayesian optimization: SMILES transformer and graph embeddings compressed through deep-kernel learning into the GP, so optimization searches catalyst, solvent, and base identity jointly with numeric process parameters. 2x better ligand selection than one-hot across 32 ligands
- Supervised the team that built an LLM system which reads lab instrument protocol manuals, generates driver code, and lets bench scientists compose automation workflows in natural language, with an execution-based evaluation framework
- Deployed a hybrid-model digital shadow at a biologics manufacturing site, now used for deviation analysis and real-time plant decisions

Associate Principal Scientist -> Principal Scientist | Data-Rich Experimentation | 2022 - 2025

Merck & Co., Inc.

- Wrote and open-sourced obsidian, Merck PR&D's first public codebase: a full GPyTorch and BoTorch implementation of Bayesian optimization and experiment design, now behind 160 active users across the company
- Established a software strategy my department has followed ever since: keep the advanced algorithms in-house, and pair them with a vendor product for scale. Now applied to hybrid modeling, lab automation control, and high-throughput chemistry plate design
- Founded Algorithmic Process Optimization and grew it from three pilots to 57 active users through a department hackathon and SME office hours, delivering 25-75% wet-lab resource reductions per campaign
- Improved the multi-objective hypervolume of an mRNA in-vitro transcription process by more than 100% across 12 variables in 43 experiments

Senior Scientist | Reaction Engineering | 2019 - 2022

Merck & Co., Inc.

- Paired machine learning with ODE-based process modeling to design molnupiravir's distillative crystallization, one of the first places I brought data science into process development
- Built in-silico nitrosamine risk assessment across 12 commercial products, avoiding more than $1.5M in contract testing
- Redesigned the molnupiravir API-forming reaction in four weeks under COVID lab restrictions, achieving a 12x stability window, higher yield, and an order-of-magnitude lower impurity load, then transferred it to an international supply site
- Elucidated a critical yield-loss mechanism in the islatravir kinase step through full kinetic characterization, informing a process redesign that reached kilogram scale

Associate Scientist -> Scientist | Chemical Engineering Research & Development | 2015 - 2019
Merck & Co., Inc.

- Technical lead for multiple doravirine steps from development through tech transfer, filing, and launch as Delstrigo and Pifeltro
- Authored 150+ pages of process development reports and CMC sections S.2.2 and S.2.6 for the doravirine NDA
- Served as commercial operations engineer during doravirine tech transfer, including on-site support in Ballydine, Ireland during a validation campaign
- Joined with no laboratory experience and won the department's Technical Achievement Award within six months, then again the next year

EDUCATION

M.S., Analytics | Georgia Institute of Technology | 2023

B.S., Chemical Engineering | University of Delaware | 2015

SELECTED PUBLICATIONS

- 14 peer-reviewed publications, 521 citations, h-index 12 (Google Scholar). Five manuscripts in preparation or submission.
- Joint Acceptable Region Exploration: Bayesian optimization for pharmaceutical process characterization. Digital Discovery, in submission.
- A Kinase-cGAS Cascade to Synthesize a Therapeutic STING Activator. Nature, 2022.
- High-Throughput Algorithmic Optimization of In Vitro Transcription for SARS-CoV-2 mRNA Vaccine Production. Biochemistry, 2024.
- Kilo-Scale Electrochemical Oxidation of a Thioether to a Sulfone. Org. Process Res. Dev., 2022. Outstanding Publication of the Year.

HONORS & AWARDS

- ACS GCI Pharmaceutical Roundtable, Data Science & Modeling for Green Chemistry (2025)
- ACS Division of Organic Chemistry, Technical Achievements in Organic Chemistry (2024)
- Organic Process Research & Development Outstanding Publication of the Year (2022)
- Merck MS&T Innovation Award runner-up (2021)