

Mohan Krishnamoorthy

TECHNOLOGY LEADER · DECISION GUIDANCE · STRATEGIC ADVISOR · AI FLUENT BUILDER

☎ (416) 839-2425 | ✉ mkrishnamoorthy2425@gmail.com | 🔗 linkedin.com/in/mohan-krishnamoorthy

Areas of Interest

Algorithm Engineering; Machine Learning & Optimization; Decision Science; Responsible AI; Human-AI Orchestration; Scientific Computing

Education

George Mason University

PH.D. IN COMPUTER SCIENCE (ADVISORS: PROF. ALEXANDER BRODSKY, PROF. DANIEL MENASCÉ)

Fairfax, VA, USA

2012–2018

Rochester Institute of Technology

MASTER OF SCIENCE IN COMPUTER SCIENCE

Rochester, NY, USA

2007–2010

Mumbai University

BACHELOR OF ENGINEERING IN COMPUTER ENGINEERING

Mumbai, INDIA

2003–2007

Positions Held

CONSULTANT & STRATEGIC ADVISOR, SELF-EMPLOYED, GREATER TORONTO AREA, ON, CANADA

11/2025–Present

CHIEF TECHNOLOGY OFFICER & FOUNDING ENGINEER, SYNOD INTELLICARE, TORONTO, ON, CANADA

04/2026–09/2026

RESEARCHER, UNIVERSITY OF OTTAWA, OTTAWA, ON, CANADA

09/2024–09/2025

CONSULTANT, SELF-EMPLOYED, GREATER TORONTO AREA, ON, CANADA

10/2023–09/2024

RESEARCH CONSULTANT, ARGONNE NATIONAL LABORATORY, GREATER TORONTO AREA, ON, CANADA

01/2022–06/2023

POSTDOCTORAL APPOINTEE, ARGONNE NATIONAL LABORATORY, LEMONT, IL, USA

10/2018–09/2021

GRADUATE RESEARCH ASSISTANT, GEORGE MASON UNIVERSITY, FAIRFAX, VA, USA

01/2013–08/2018

GUEST RESEARCHER, NATIONAL INSTITUTE OF STANDARDS AND TECHNOLOGY, GAITHERSBURG, MD, USA

01/2015–12/2017

GRADUATE TEACHING ASSISTANT, GEORGE MASON UNIVERSITY, FAIRFAX, VA, USA

08/2012–12/2012

RESEARCH TECHNOLOGIST, LOS ALAMOS NATIONAL LABORATORY, LOS ALAMOS, NM, USA

05/2010–06/2012

RESEARCH ASSISTANT, LOS ALAMOS NATIONAL LABORATORY, LOS ALAMOS, NM, USA

07/2008–04/2010

Experience

Self-employed

Greater Toronto Area, ON

CONSULTANT & STRATEGIC ADVISOR

11/2025–Present

Independent advisor helping startups make better decisions with AI and data

- Advise an early-stage startup and build AI models that recommend how to split money across investments, balancing return against risk, tested on 10 years of market data.
- Built budgeting models that show where to spend and where to cut, reducing planned costs by about 15%.
- Guide decisions by combining decision math (operations research) with machine learning, so the founders get clear recommendations they can act on with confidence.
- Use AI assistants to capture the nuance in a founder's question and build a working model in days instead of weeks, then pressure test it on real-world scenarios before advising.

Synod IntelliCare

Toronto, ON

CTO & FOUNDING ENGINEER

04/2026–09/2026

Startup building tools that check healthcare AI for bias

- Took a healthcare AI fairness-checking platform from first design to a live cloud product in a few weeks, testing whether AI treats patient groups fairly against Canadian population data.
- Built and led a team across backend, infrastructure and frontend, setting one shared plan and a steady routine of spec reviews, task tracking and team syncs so everyone moved in the same direction.
- Made the key technical calls on design, security and build priorities, protecting sensitive health data with secure login, role-based access and full activity logs.
- Worked with the executive team on company strategy, product positioning and go-to-market priorities, turning them into a roadmap the team could execute.
- Secured research and accelerator partnerships with the Vector Institute and Sheridan College, strengthening the company's credibility with investors and academic groups.
- Used AI assistants across design, coding and review to ship faster, about 30% quicker, keeping quality high through peer review and automated checks.

University of Ottawa

RESEARCHER

Ottawa, ON

09/2024–09/2025

Research on how people and AI should share important decisions

- Created a decision guide that shows when AI should act on its own, when a person should review, and when a person should stay in charge, based on how risky the decision is.
- Defined who is responsible for what between people and AI, reducing overreliance on AI in high-stakes work.
- Gave leaders a practical way to decide where to trust AI and where to keep human judgment.

Self-employed

CONSULTANT

Greater Toronto Area, ON

10/2023–09/2024

Independent practice building decision models for businesses

- Built math models that recommend the best staffing plans, reducing overtime and idle time for a services client.

Argonne National Laboratory

POSTDOCTORAL APPOINTEE, THEN RESEARCH CONSULTANT

Lemont, IL

10/2018–06/2023

R&D software for high-energy physics in a U.S. national research lab

- Built MÆSTRO (Multi-fidelity Adaptive Ensemble Stochastic Trust Region Optimization), an open-source tool that tunes slow, noisy computer simulations to match real experiment data, checking each step so results can be verified. Also built Apprentice, which creates fast stand-ins for slow simulations.
- Replaced personal judgment in weighing competing goals with a repeatable method, so tuning decisions could be explained and defended.
- Packaged the methods into a workflow that runs on supercomputers, so other scientists could reuse them.

George Mason University

GRADUATE RESEARCH ASSISTANT, PH.D. CANDIDATE

Fairfax, VA

2013–2018

Doctoral research on better decisions in manufacturing

- Developed SODA (Stochastic Optimization based on Deterministic Approximations), an open-source method that finds the best settings for large manufacturing networks with uncertain outcomes, by swapping the uncertain parts for simple estimates.
- Built reusable models of real manufacturing processes so engineers could apply them to new problems.

National Institute of Standards and Technology (NIST)

GUEST RESEARCHER

Gaithersburg, MD

2015–2017

Manufacturing modeling R&D in a U.S. government standards lab

- Built FactoryOptima, a prototype that lets engineers compare factory designs and trade-offs using a library of reusable models.
- Created that library from real-world data, published research and crowdsourced data.

Los Alamos National Laboratory

RESEARCH ASSISTANT, THEN RESEARCH TECHNOLOGIST

Los Alamos, NM

07/2008–06/2012

Bioinformatics R&D in a U.S. national research lab

- Built software that helped scientists study HIV and influenza, which taught me to turn messy data into usable tools.
- Redesigning the HIV project website so researchers could find and use data more easily.

Software Packages**

MÆSTRO (Multi-fidelity Adaptive Ensemble Stochastic Trust Region Optimization)

 [DOC](#)  [CODE](#)

Open-source solver that tunes slow, noisy computer simulations until they match real experiment data. It builds a simple stand-in for the simulation as it goes and checks each step against it, so results can be verified rather than trusted blindly. It repeats until the fit is good enough.

Apprentice

 [DOC](#)  [CODE](#)

Open-source tool that builds fast, simple formulas to stand in for slow, costly simulations. Because the formulas are quick to run, they can be optimized quickly.

Outer Optimization

 [CODE](#)

Open-source tool that decides how much weight each piece of data deserves when tuning a model. It sets the weights automatically instead of by experience and intuition, so results are reached efficiently and are less subjective.

SODA (Stochastic Optimization based on Deterministic Approximations)

 [CODE](#)

Open-source method that finds the best control settings for large manufacturing networks where outcomes are uncertain. It swaps the uncertain parts for simple estimates so answers arrive quickly, while staying within required limits.

FactoryOptima

 [PAPER](#)

Web prototype where manufacturing engineers build, optimize and compare factory designs. It draws on a reusable library of performance models and shows the trade-offs between options.

Publications**

 [Click here to see my Publications on Google Scholar](#)

** Clickable links and the latest version of this resume are available at ksmohan.github.io/cv