

Mohan Krishnamoorthy

RESEARCH & DEVELOPMENT · SCIENTIFIC ENGINEER & PROGRAMMER

✉ mkrishnamoorthy2425@gmail.com

📱 mohan-krishnamoorthy-46058bb

Areas of Interest

AI Algorithm Engineering; ML & Optimization Modelling; Interdisciplinary Research; Big Data Analytics; High Performance 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

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

06/2026–Present

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

11/2025–Present

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, GAITHERSBURGH, 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

Software packages

MÆSTRO

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

MÆSTRO stands for Multi-fidelity Adaptive Ensemble Stochastic Trust Region Optimization and it is an open source plug-and-play derivative-free stochastic optimization solver. The problem being considered in MÆSTRO involves fitting Monte Carlo simulations that describe complex phenomena to experiments by finding parameters of the resource intensive and noisy simulation that yield the least squares objective function value to a noisy experimental data. This problem is solved using an active machine learning algorithm where in each iteration, a local approximation of the simulation signal and of the simulation noise is constructed over data. This methodology supports verifiable evidence over blind trust by ensuring every simulation component of the objective is replaced by a transparent analytical approximation before the next iterate is found within the trust-region. Then the trust region is moved and the iterations continue until a satisfactory convergence criterion is met.

Apprentice

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

An open source package for construction of multivariate analytic surrogate model for computationally expensive Monte-Carlo predictions. The surrogate model is used for numerical optimization of a prediction function since it can be prohibitively expensive to perform optimization over functions with the Monte-Carlo predictions.

Outer optimization

[CODE](#)

An open source package to assign weights and solve the tuning problem of finding optimal parameters that minimizes a least-squares function between approximations of noisy simulations and experimental data or data observed in nature. Instead of setting weights manually based on experience and intuition, the weights are automatically adjusted using a bilevel optimization or a single level robust optimization formulation, thus yielding results efficiently that are less subjective.

SODA

[CODE](#)

An open source package containing the algorithm to perform stochastic optimization based on deterministic approximations to efficiently solve the problem of finding control settings for stochastic processes in a large manufacturing service network subject to the satisfaction of stochastic feasibility constraints.

Factory optima

[PAPER](#)

Web-based prototype system that allows manufacturing process engineers to compose, optimize and perform trade-off analysis of manufacturing service networks based on a reusable repository of performance models.

Experience

Synod IntelliCare

Toronto, ON

CHIEF TECHNOLOGY OFFICER & FOUNDING ENGINEER

05/2026–Present

- Architected and led development of DDFA, a fairness auditing platform for healthcare AI, building a FastAPI/Python and React/TypeScript system with OAuth2/PKCE authentication, RBAC, audit logging, and Canadian census-based bias benchmarking, deployed on AWS (EC2, RDS, S3, IAM).
- Built and directed a 5-person engineering team across backend, infrastructure, and frontend functions, while establishing structured processes (SRS gap analysis reviews, task tracking, cross-functional syncs) to align a fast-moving startup team around a shared technical spec.
- Partner closely with executive leadership to shape company strategy, from securing key research and accelerator partnerships to defining product positioning and go-to-market priorities, translating high-level vision into an actionable execution roadmap for the team.

Self-employed

Greater Toronto

Area, ON

CONSULTANT

11/2025–Present

- Continuing consultation on a AI modeling paradigm to perform efficient asset allocation and portfolio optimization using operations research and machine learning techniques.

University of Ottawa

Ottawa, ON

RESEARCHER

09/2024–09/2025

- Designed a context-sensitive Human-AI orchestration framework using a multi-stage methodology to systematically assign optimal collaboration modes (HITL, HOTL, HITLFE, HOOTL) based on decision criticality and latitude.
- Designed Human-AI responsibility allocation using 4C model (Communication, Coordination, Cooperation, Collaboration) to distribute tasks across interaction layers, ensuring accountability and reducing risks like automation bias in high-stakes environments.

Self-employed

Greater Toronto

Area, ON

CONSULTANT

10/2023–09/2024

- Developed mathematical programming models to solve the staffing optimization problem.
- Consulted on a project to perform efficient asset allocation and portfolio optimization using operations research and machine learning techniques.

Argonne National Laboratory (ANL)

Greater Toronto

Area, ON

RESEARCH CONSULTANT

01/2022–06/2023

- Developed mathematical and algorithmic techniques for directly fitting Monte Carlo simulations to experimental data or data observed in nature using a stochastic trust-region optimization algorithm.
- Developed a workflow package to efficiently solve derivative-free stochastic optimization problems in a high performance computing environment.

Argonne National Laboratory (ANL)

Lemont, IL

POSTDOCTORAL APPOINTEE

10/2018–09/2021

- Developed and implemented mathematical and algorithmic techniques for approximating expensive functions in High Energy Physics (HEP).
- Developed robust optimization and design of experiment formulations to decide weights of importance with the goal of tuning a HEP Monte Carlo simulator.
- Developed, maintained, and published HEP analysis packages called *apprentice* and *outer optimization* for efficiently constructing polynomial/rational approximations and for performing χ^2 minimizations.

George Mason University (GMU)

Fairfax, VA

GRADUATE RESEARCH ASSISTANT

01/2013–08/2018

- Designed and developed reusable mathematical models for non-linear, stochastic, hierarchical, and temporal manufacturing processes from real-world data.
- Designed and developed one-stage stochastic optimization algorithms based on deterministic approximation heuristics.
- Developed and published the stochastic optimization algorithm based on deterministic approximations (SODA) to efficiently solve the problem of finding controls for stochastic processes in a large manufacturing service network.

National Institute of Standards and Technology (NIST)

Gaithersburg, MD

GUEST RESEARCHER

01/2015–12/2017

- Designed and populated a repository of reusable mathematical models that were sourced from real-world data, publications, and crowdsourced data.
- Developed a software framework and prototype (*FactoryOptima*) to perform composition, analysis, and optimization on reusable models.

Los Alamos National Laboratory (LANL)

RESEARCH TECHNOLOGIST

[Los Alamos, NM](#)

05/2010–06/2012

- Designed and developed scientific algorithms for highly variable and large scale bioinformatics tools.
- Developed and debugged multiple backend modules of the HIV project.
- Redesigned the HIV website using Model-View-Controller (MVC) framework and web services.

Los Alamos National Laboratory (LANL)

RESEARCH ASSISTANT

[Los Alamos, NM](#)

07/2008–04/2010

- Designed and developed Tree Viewer, Pruner and Decorator tools to perform selection and annotation of Influenza sequences.
- Designed and developed a schema for a large Influenza Sequence Database.

Publications

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