# NAVEEN MYSORE

nmysore.work@gmail.com
San Francisco Bay Area
H-1B visa holder with approved I-140

**Senior Software and Research Engineer | Member of Technical Staff**

linkedin.com/in/mysore-naveen | github.com/NAVEENMN | naveenmn.github.io | Google Scholar | huggingface.co/nmysore

## Professional Summary

Software and research engineer with 10+ years across Salesforce (Tableau), Dell EMC, UCSB, and Dyssonance AI. Deliver distributed systems, enterprise features, and data infrastructure through requirements analysis, cross-team delivery, and systematic performance optimization: 5% better Tableau extraction performance and 10x UCSB experiment throughput. Build machine learning and deep learning systems for model training, fine-tuning, inference pipelines, Agentic systems, cloud deployment, and large language model (LLM) memory and reasoning. Created InKAN, an open-source PyTorch library for interpretable neural networks. Research spans interpretable and explainable AI, recursive self improvement, reinforcement learning, and causal discovery, with an accepted RLC 2026 paper and scientific computing preprints.

## Technical Skills

| Area | Skills and tools |
|---|---|
| **AI, ML & Research** | Machine learning, deep learning, large language models (LLMs), reinforcement learning, fine-tuning (SFT, RLHF, PPO, GRPO), reasoning, reward modeling, graph neural networks, explainable AI, causal inference, retrieval-augmented generation (RAG) |
| **Software Engineering** | Distributed systems, system design, software architecture, requirements analysis, API development, performance optimization, testing, observability, root cause analysis |
| **Scientific Computing** | Numerical methods, statistical modeling, hypothesis testing, time series analysis, uncertainty quantification, physical parameter estimation, KANs, B-splines |
| **Languages and Libraries** | Python, PyTorch, Hugging Face Transformers, TRL, DeepSpeed, NumPy, SciPy, scikit-learn, Weights & Biases (W&B) |
| **ML Systems** | Distributed training, multi-GPU training, model optimization, model deployment, inference optimization, ML pipelines, experiment tracking, hyperparameter tuning, model evaluation |
| **Cloud and Data** | AWS, Google Cloud (GCP), Terraform, Apache Kafka, PostgreSQL, vector databases |

## Professional Experience

**Dyssonance AI (stealth)** **Feb 2026 - Present**

**Senior Research Engineer**

Develop memory and reasoning capabilities in Llama- and Gemma-based open-weight models through joint training with graph neural network outputs. Train reasoning models for long-horizon planning and convert research prototypes into product features.

**University of California, Santa Barbara** **Oct 2024 - Feb 8, 2026**

**Research Specialist | Prof. Yao Qin's team, REAL AI**

Built AI-assisted experiment loops to explore hypotheses, run hyperparameter scans, and review results, reducing setup from hours to minutes and increasing experiment throughput **10x**. Automated Weights & Biases (W&B) setup and result alerts, stored results in PostgreSQL, and provided scripts to retrieve results and plot curves for weekly research reviews. Built SFT-to-PPO workflows, fine-tuned LLaMA-3.1 8B for nutrition reasoning model with streaming human feedback for continual RLHF, and deployed inference on AWS. Supported ML engineering across research teams and developed Nutribench to collect health data for AI training.

Live demo: SMS +1 805 918 8182 to receive a response from the UCSB-trained nutrition model.

Code: github.com/naveen-mysore/Reinft | Platform demo | Nutribench app

## Professional Experience Continued

**Salesforce (Tableau) | San Francisco, CA** **Feb 2019 - Oct 2024**

**Senior Member of Technical Staff (SMTS) | Nishi Bhonsle's Core Build and Tableau Extract team**

Improved Tableau's data extraction pipeline performance by **5%** through bottleneck analysis and optimization under heavy loads. Led cross-team delivery of Customer Managed Encryption Keys (CMEK), enabling secure onboarding of McKinsey. Built Kafka-based adoption tracking, A/B testing dashboards, a statistical hypothesis testing platform, and AI-assisted intelligent metrics used by cross-functional teams.

**Dell EMC | Santa Clara, CA** **Jun 2015 - Feb 2019**

**Software Engineer**

Designed a distributed data-server synchronization algorithm using balanced N-ary tree updates to reduce inconsistencies. Benchmarked server workloads on AWS and used regression modeling to match workloads to instance types for cost-performance efficiency.

## Open Source Projects

**InKAN | Author and Maintainer**

Developed and published InKAN (1.1k/month downloads), a PyTorch library for B-spline Kolmogorov-Arnold Networks supporting CPU, CUDA, and Apple Silicon, with reusable layers and visualization tools. Optimized B-spline evaluation with torch.compile kernel fusion, achieving 2.8-3.5x faster forward passes than Cox-de Boor recursion in tested NVIDIA H100 benchmarks.

pip install inkan : pypi.org/project/inkan

**Policy Manifolds | Reinforcement Learning and Interpretable Policies**

Built PPO/SAC training and evaluation workflows to identify compact B-spline KAN policies across 14 Classic Control, Box2D, and MuJoCo environments, including CartPole, LunarLander, and Humanoid. Developed an interactive lab where users can edit learned edge functions and compare changes in agent behavior against the original policy without retraining.

Try the interactive demo: policymanifolds.com

## Selected Publications and Preprints

**Prediction-Based Markov Violation Scores for Detecting Non-Markovian Observations in Reinforcement Learning**

Naveen Mysore | Reinforcement Learning Journal, vol. 7, 2026; accepted at RLC 2026.

rlj.cs.umass.edu/2026/papers/Paper70.html

**Recovering Physical Parameters from Fragmented Observations via Exact Distributed Spline Merging**

Naveen Mysore | arXiv preprint, 2026.

arxiv.org/abs/2609.16579

* Recovered diffusion coefficients and wave speed with 0.11% and 0.12% error in synthetic experiments; validated distributed reconstruction on 41 years of NOAA data.

**InKAN: B-Spline KANs via Truncated Power Form**

Naveen Mysore | arXiv preprint, 2026.

arxiv.org/abs/2609.01956

## Patent Application

**System and Method for Distributed Spatiotemporal Field Reconstruction and Physical Parameter Recovery Using Composable Spline Sufficient Statistics**

Named inventor | U.S. provisional application 64/157,134

## Education

**Master's in Computer Science (Machine Learning)** **2013 - 2015**

University of North Carolina at Charlotte

Thesis: Elastic group recommendation systems for multivariate dynamic attributes. IEEE ICIS 2016.

**Bachelor's in Electrical and Electronics Engineering** **2009 - 2013**

PES University (PESIT)