SIDISHA SHYAM BARIK

sidishabarik@gmail.com | LinkedIn | Website

EDUCATION

Master of Science (MS) in Computer Science
Sep 2024 - May 2026
University of Massachusetts, Amherst
CGPA: 4.0
Relevant Coursework: Neural Networks, Machine Learning, Probability and Statistics, Robotics, Advanced Algorithms, Natural Language Processing

Bachelor of Science (BSc) in Data Science and Applications
Jan 2021 - May 2024
Indian Institute of Technology Madras (IIT Madras), India
CGPA: 8.6/10
Relevant Coursework: Deep Learning, Large Language Models, Data Structures and Algorithms

Bachelor of Engineering (BE) in Mechanical Engineering
Aug 2016 - May 2020
Birla Institute of Technology and Science, Pilani (BITS Pilani), India
CGPA: 8.6/10

RESEARCH INTERESTS

My research interests broadly lie in developing mathematically grounded and scalable AI systems by combining deep learning with structured inference, probabilistic reasoning, and classical algorithmic frameworks. I am interested in understanding and improving modern learning architectures, and in integrating domain structure and physical constraints into learned systems.

PUBLICATIONS

[1] P. Balakrishnan*, S. Barik*, S. M. O'Rourke, and B. M. Marlin, "GPU-GLMB: Assessing the Scalability of GPU-Accelerated Multi-Hypothesis Tracking," in 5th International Workshop on the Internet of Things for Adversarial Environments, IEEE Military Communications Conference (MILCOM), 2025

[2] B. M. Marlin, S. M. O'Rourke, S. Barik, P. Balakrishnan, and C. Samplawski, "Improving target tracking in stone soup via state-dependent process noise models," Under review at International Conference on Information Fusion, 2026

[3] S. Barik*, P. Balakrishnan*, S. M. O'Rourke, and B. M. Marlin, "GPU-Accelerated Generalized Labeled Multi-Bernoulli Filter for Multi-Detection Point Targets," Manuscript in preparation for submission to IEEE Transactions on Signal Processing.

* - equal contribution

RESEARCH EXPERIENCE

UNCERTAINTY-AWARE PERCEPTION AND TRACKING

Differentiable Probabilistic Perception and Multi-Object Tracking
Advisor: Prof. Benjamin Marlin

- Designed an end-to-end **multi-sensor perception and tracking system** integrating **vision-based object detection with probabilistic inference** for real-time state estimation. *[MILCOM, 2025]*
- Developed a scalable **Bayesian multi-object tracking framework** that **replaces traditional sequential data-association** procedures with **massively parallel GPU execution**, enabling real-time inference while preserving exact probabilistic reasoning.
- Formulated the tracker as a **differentiable PyTorch architecture**, bridging **deep perception models** and **structured Bayesian inference** within a unified learning framework.
- Integrated **physics-based motion models** and **uncertainty-aware inference** to improve robustness under clutter, missed detections, and complex vehicle trajectories. *[Information Fusion, 2026 under review]*
- Achieved **sub-100 ms tracking latency** across **Dockerized** edge (**Jetson Orin**), consumer (**RTX 4090**),

and server-class (**L40S**) GPU environments, enabling real-time deployment in challenging tracking scenarios. *[IEEE TSP manuscript in preparation]*

Differentiable Particle Filter (DPF) for Geospatial Tracking
Advisor: Prof. Benjamin Marlin

- Designed and implemented a **Differentiable Particle Filter** (DPF) in PyTorch, utilizing a Constant Turn Rate and Velocity (CTRV) motion model to capture non-linear vehicle trajectories.
- Solved the gradient flow challenge through stochastic sampling steps by employing the re-parameterization trick, enabling the model to learn process noise parameters directly via backpropagation.
- Addressed blind spot failure modes where standard motion models produced physically impossible predictions (e.g., driving through walls).
- Introduced **map-based constraints** into the prediction step, forcing probability distributions to split into multiple valid hypotheses (e.g., bimodal distribution at a T-intersection) rather than diffusing through obstacles.
- Demonstrated that differentiable probabilistic layers serve as an interface for **injecting symbolic priors into learnable systems**, ensuring physical robustness and interpretability.

Autonomous Human Following Robot
Advisor: Prof. Hao Zhang

- Built an end-to-end **robotic perception and control system** integrating object detection, multi-object tracking, re-identification, and closed-loop navigation for autonomous human following on a ROS-enabled omnidirectional robot.
- Developed a custom **appearance-based re-identification** framework, extending SORT with lightweight visual features to maintain target identity through occlusions and operate under embedded compute constraints.
- Designed and deployed a modular **ROS-based perception-to-control architecture** on an NVIDIA Jetson platform, achieving 10–15 FPS real-time perception and 1.2 s response time for stable target following.

FOUNDATION MODELS AND REPRESENTATION LEARNING

Efficient Residual Stream Scaling for Transformers
Collaboration: Google DeepMind

- Investigated **efficient transformer scaling techniques** to increase model representational capacity without quadratic compute growth, addressing a fundamental bottleneck in large language model design.
- Built a **unified PyTorch-based pretraining pipeline** for systematic evaluation of residual-stream scaling architectures (mHC, VWN, AltUp), training **135M-parameter language models on 13.5B tokens** under controlled experimental settings.
- Developed an evaluation framework benchmarking model performance across **training/inference FLOPs, runtime overhead, throughput, and downstream language understanding tasks**.
- Analyzed the tradeoff between theoretical compute efficiency and practical training behavior, identifying major runtime bottlenecks arising from memory access patterns and CUDA kernel launch overhead.
- Observed up to **30.4% improvement on downstream language benchmarks** while maintaining near-constant parameter and FLOP growth, highlighting how architectural modifications to residual pathways can improve representational efficiency without proportional compute growth.

Improving Spatial Reasoning in Vision-Language Models via JEPA World Models
Advisor: Prof. Brendan O'Connor

- Designed a novel **Vision Language Model (VLM)** replacing standard vision encoders with a Vision Transformer (ViT)-based **self-supervised world model (I-JEPA)** to extract highly robust spatial and structural representations.
- Solved the vision language alignment problem by developing a feature translation module that maps geometric visual representations into the language embedding space of a large language model (Qwen2-7B).
- Built a two-stage training pipeline (feature alignment + instruction tuning) to integrate non-semantic visual features with language reasoning.
- Engineered an efficient training setup using **LoRA** and **4-bit quantization**, enabling large-scale multimodal experiments on a single GPU.

Investigation of Noise Schedulers for Generative Diffusion Models
Advisor: Prof. Subhransu Maji

- Studied the **impact of noise scheduling** on diffusion model training and sampling behavior across Variance Preserving (VP) and Sub-VP formulations.
- Built and **trained DDPMs end-to-end** with a 100M-parameter UNet backbone, incorporating and evaluating Cosine, Sine, and Laplace noise schedulers within the diffusion process.
- Formulated a **unified logSNR-based framework** for comparing scheduler distributions and analyzed generative performance using FID evaluation on CIFAR-10.

TEACHING EXPERIENCE

Graduate Teaching Assistant

Jan 2026 – May 2026
University of Massachusetts Amherst

- Course: **COMPSCI 690S – AI Alignment**
- Assessed student reading assignments and written analyses of approximately 20 foundational and contemporary AI alignment research papers.
- Evaluated students' understanding of key concepts, technical contributions, and broader implications discussed in AI alignment literature, and addressed grading-related queries.

Graduate Teaching Assistant

Sep 2025 – Dec 2025
University of Massachusetts Amherst

- Course: **COMPSCI 689 – Advanced Machine Learning**
- Assessed student assignments, reports, and exams for graduate-level machine learning coursework.
- Evaluated students' understanding of advanced machine learning concepts and provided clarification on grading and underlying technical concepts during discussions with students.

WORK EXPERIENCE

Senior Analyst, Credit & Fraud Risk | Data & ML Engineering

Aug 2020 – May 2024
American Express

- Optimized a production XGBoost credit risk model through feature engineering and training-data optimization, contributing to $14.7M in recoveries through improved default-risk identification.
- Built end-to-end ETL and feature-engineering pipelines using PySpark, Hive, and SQL, integrating data from multiple financial sources to create high-quality training datasets for credit risk modeling.
- Led migration of merchant exposure analytics from legacy systems to a distributed big-data platform (Hadoop), reducing infrastructure costs by $45M while improving data accessibility and scalability for credit-risk modeling.

AWARDS AND HIGHLIGHTS

1. **Q2 Star Award, American Express (2022)**: Awarded by the Risk Products and Data Strategy Department for exceptional contributions to the Global Merchant Services Exposure Management Project.
2. **Top 0.1% Merit Certificate, CBSE India (2016)**: Recognized for excellence in Mathematics and Physics, placing in the top 0.1 percentile of 1,000,000 students in the All India Senior School Certificate Examination (AISSCE).
3. **JEE Mains All India Rank 285 (2016)**: Achieved a rank of 285 out of 1.3 million candidates in the Joint Entrance Examination (JEE) Mains.
4. **National Talent Search Examination (NTSE) Scholar (2014)**: Recipient of the prestigious national scholarship awarded to the top 1,000 students out of 900,000 applicants across India.