Sujosh Nag

Durham, NH | (603) 842-2129 | nagsujosh2004@gmail.com
sujoshnag.com | linkedin.com/in/sujoshnag | github.com/nagsujosh | Google Scholar

Research Interests

Robot learning, dynamical systems, and learning-based control, with an emphasis on structured representations of dynamics and action for compositional, generalizable, and reliable autonomy.

Education

University of New Hampshire Durham, NH
B.S. in Computer Science *Expected May 2027*

- **Relevant Coursework:** Reinforcement Learning, Machine Learning, Artificial Intelligence, Foundations of Neural Networks, Mobile Robotics, Linear Algebra, Applied Regression Analysis, Numerical Methods, Statistics for Engineers.

Publications

- Riad Ahmed, **Sujosh Nag**, Moniruzzaman Akash, Mostafa Hussein, and Momotaz Begum. **"Trajectory-Consistent Flow Matching for Robust Visuomotor Policy Learning."** IEEE/RSJ International Conference on Intelligent Robots and Systems (**IROS 2026**). arXiv:2605.08511.
  Flow-matching policies train on a pointwise velocity objective but integrate that field at inference, and the mismatch compounds along the trajectory. The paper adds rectified-flow regression for uniform temporal coverage, supervises the integrated displacement of trajectory segments directly, regularizes the velocity field for smoothness, and replaces Euler with fourth-order Runge-Kutta integration; the ablation shows none of the four suffices alone.

Workshop Papers

- **Sujosh Nag** and Momotaz Begum. **"Where Should a Skill End? Discovering Skills from How the Robot's Action Moves Objects."** 4th CoRL Workshop on Learning Effective Abstractions for Planning (**LEAP 2026**). openreview.net/pdf?id=cW6YVaGb73.
  A planner composes skills, so a skill boundary should follow what a planner reasons about: the effect of the robot's action on objects, not appearance or arm motion, which place boundaries wherever the arm happens to move. The paper argues a skill is a recurring effect and lets out-of-sample prediction set how many there are, merging two whenever one effect model predicts held-out demonstrations as well as two, and instantiates this label-free with switched linear models of object motion conditioned on the gripper's.

Manuscripts Under Review

- **Sujosh Nag**, Kaito Namatame, Riad Ahmed, Moniruzzaman Akash, Meisam Khorshidi, and Momotaz Begum. **"MoK: A Mixture of Koopman Operators for Learning Visuomotor Policies." Under review**.
  Prior Koopman policies fit one operator to a whole task. MoK trains a bank of them and composes them without segmentation labels: a sparsemax gate reads the latent at each step and applies the selected operators in sequence, so one step can draw on several operators and a regime is a recurring sequence of compositions. Every composition the gate can form is contractive in closed form before training begins, over observables the lift learns through the operator itself, keeping no raw measurement of the robot or the scene.
- Moniruzzaman Akash, **Sujosh Nag**, Sajay Arthanat, and Momotaz Begum. **"Hand2Robot: An End-user Friendly Demonstration Collection Framework to Facilitate Home Deployment of Imitation Learning." Under review**.
  Lets an end user demonstrate a task with their own hands, with no teleoperation interface, and converts egocentric RGB-D recordings into parallel-jaw policy data. Six grasp-specific contact frames set the gripper pose across precision and power grasps, predicted manipulation phases select the task-relevant motion and time grasp and release, replay in Gaussian-splat scenes supplies matching robot images, and calibrated surface electromyography supervises how firmly to hold the object.
- Riad Ahmed, **Sujosh Nag**, Moniruzzaman Akash, and Momotaz Begum. **"Detecting and Explaining Robot Policy Failures from Sparse Internal Features." Under review**. arXiv:2611.10400.
  Trains a sparse autoencoder on the activations of a vision-language-action policy and names each basis feature by a manipulation concept that holds whenever it fires, keeping only names that stay precise on unseen initial states and on an unseen task and that survive a closed-loop clamp confirming the policy acts through the feature. A run-time monitor over those same features flags a failing execution and names the internal quantities responsible.

Research Experience

Cognitive Assistive Robotics Lab (CARL), University of New Hampshire Durham, NH
Undergraduate Researcher | Advisor: Prof. Momotaz Begum *August 2025 - Present*

- Design and run the training and closed-loop evaluation of **visuomotor manipulation policies** (ACT, Diffusion Policy, DP3, flow matching, $\pi_0/\pi_{0.5}$, OpenVLA, NVIDIA Isaac GR00T) in simulation and on hardware, and own the surrounding infrastructure for demonstration collection, training, rollout, and failure analysis.
- Lead the lab's real-robot experiments on **Franka FR3**, **Boston Dynamics Spot**, **Hello Robot Stretch 3/4**, and **ABB YuMi**, spanning single-arm, mobile-manipulation, and bimanual settings.
- Built and maintain the **Franka + Deoxys** stack for teleoperation, demonstration collection, and policy evaluation, with SpaceMouse and Quest interfaces, multi-camera rigs, and **Inspire dexterous-hand** control via Quest hand tracking.
- Supported an **NIH R01-funded assistive-robotics deployment** in a participant's home, outside the controlled conditions of the lab.

Independent Research Projects

Structured Belief-Guided Diffusion Policy

Conditioned a diffusion policy on a belief over task phase and progress predicted from image observations, taking as given the skill boundaries that define those phases. The belief helped in Meta-World and on real Franka tasks, yet under temporal perturbation the policy recovered less reliably than an unconditioned diffusion policy. That negative result moved my attention upstream, to where skill boundaries should come from at all.

Belief-PDDL: Neuro-Symbolic Planning Under Partial Observability

Planning framework that maintains factorized Bayesian beliefs over grounded predicates, projects them through CP-SAT constraints to drop inconsistent states, and plans over Top-K feasible worlds with active sensing; robust to heavy observation noise in Blocksworld and ALFWorld.

Course Research

Semantic Observation Repair for Decentralized Multi-Agent Path Finding GitHub | AI course

Inference-time repair of corrupted observations for a frozen decentralized PPO policy in POGEMA, across eight corruption types (noise, dropout, bit flips, FGSM/PGD). Over 4,000 episodes, semantic projection to valid observation grids raised mean success from **68.8% to 88.2%** and recovered FGSM success from **0% to 97.5%** without retraining.

AAPR: Adversary-Adaptive Privacy-Preserving Representations GitHub | ML course

Suppressed biometric attributes (speaker identity, gender) in learned representations while keeping task utility, via a variational information bottleneck filter trained through a gradient reversal layer against a periodically re-initialized adversary. Gender leakage fell to chance but identity stayed fully recoverable.

Teaching Experience

CS 733/833: Mobile Robotics University of New Hampshire
Teaching Assistant *Fall 2026*

Run the laboratory section for the combined undergraduate and graduate offerings of the course. Teach students **ROS 2** in simulation and on the **Hello Robot Stretch 3**, covering robot programming, mapping, localization, navigation, and AprilTag-based perception.

CS 518: Introduction to Software Engineering University of New Hampshire
Teaching Assistant *Spring 2025*

Supported team software projects and assignments, answered student questions, and graded coursework.

Professional Experience

Sustainability Roundtable Inc. Boston, MA
Software Engineer Intern *June 2025 - August 2025*

Built production features for the SBER 3.0 platform in React and TypeScript, plus internal automation that cut manual workflow time by roughly 75%.

University of New Hampshire InterOperability Laboratory Durham, NH
Senior Developer *October 2023 - March 2025*

Built test automation for the lab's **NVMe conformance test suite**, which qualifies datacenter SSDs and NVMe-oF targets for the NVM Express Integrator's List: contributed 70+ conformance tests across NVMe and its PCIe transport and resolved 120+ issues. Led a team of interns building a test-result comparison tool that cut regression-testing time roughly fourfold, automated Jira ticket creation from test results, and taught Microbit programming and hardware/software integration to 20+ students as a teaching assistant for the lab's TechVenturers program.

Dorm Drop Durham, NH
Chief Technology Officer *November 2023 - July 2024*

Owned the full software development lifecycle at an early-stage campus delivery startup: set the technical direction, led the engineering team, and architected a **scalable full-stack platform** around the ordering and dispatch data model. Shipped the beta mobile app and marketing site end to end, and ran the customer-discovery and feedback cycle, over which internal metrics showed roughly **34%** higher user satisfaction and **30%** higher beta engagement.

Honors, Scholarships, and Awards

- **Dean's Scholar**, University of New Hampshire, 2023-2026.
- **Shaw Innovation Scholar**, University of New Hampshire, 2024-2025.
- **International Engagement and Service Scholarship**, University of New Hampshire.
- **Summer Seed Grant**, University of New Hampshire, 2024. Competitive funding awarded to Dorm Drop.
- **Paul J. Holloway Prize Competition**, Semifinalist, 2025. Reached the semifinal round with Dorm Drop.

Leadership and Activities

- **Founder, UNH Squash Club** (Fall 2024 - Present)**.** Run an unofficial club offering regular coaching sessions and court time to UNH students and Durham community members across experience levels.
- **Founder and Chair, Desi Student Association, UNH** (2023 - Present)**.** Lead the executive team of a South Asian cultural organization serving students from India, Pakistan, Nepal, Bangladesh, and Sri Lanka. Run its annual Diwali and Holi celebrations, street-food and performance nights, and campus outreach, frequently in collaboration with other cultural student organizations.
- **Treasurer, Bangladeshi Student Association, UNH** (2025)**.** Managed the association's budget and funding for its Bangladeshi cultural events on campus.
- **Finance Manager, McLaughlin Hall Council, UNH** (Fall 2023)**.** Managed the hall's budget and event funding.
- **Member, Idea and Innovation Society, UNH** (2023-2024)**.**

Technical Skills

Robot Learning: Diffusion Policy, DP3, flow matching, $\pi_0/\pi_{0.5}$, OpenVLA, NVIDIA Isaac GR00T, Koopman models, sparse autoencoders, PointNet, R3M, DINO, SAM, CoTracker.

Robotics & Simulation: ROS 2, Deoxys, MuJoCo, NVIDIA Isaac Sim, Isaac Lab, Newton, robosuite, robomimic, Meta-World, SpaceMouse/Meta Quest teleoperation.

Perception & Data: RealSense RGB-D, point clouds, multi-camera calibration, AprilTags, iPhone ARKit/LiDAR, IMU, HDF5/Zarr datasets.

Software: Python, PyTorch, C++, C, Java, MATLAB, Bash, SQL, JavaScript/TypeScript, React, React Native, Angular, Node.js, Swift, Docker, Linux, Git, LaTeX.

Selected URLs

Portfolio: https://sujoshnag.com
GitHub: https://github.com/nagsujosh
LinkedIn: https://linkedin.com/in/sujoshnag
Google Scholar: https://scholar.google.com/citations?user=C-DVuIcAAAAJ
Trajectory-Consistent Flow Matching (IROS 2026): https://arxiv.org/abs/2605.08511
Where Should a Skill End? (LEAP 2026): https://openreview.net/pdf?id=cW6YVaGb73
Detecting and Explaining Robot Policy Failures (under review): https://arxiv.org/abs/2611.10400
Semantic Observation Repair for Decentralized MAPF: https://github.com/nagsujosh/Observation-Repair-MAPF
AAPR: Adversary-Adaptive Privacy-Preserving Representations: https://github.com/nagsujosh/APPR-PHOTOS