

TRIASHA SARKAR

Atlanta, GA · +1 (404) 948-8761 · tsarkar34@gatech.edu

[linkedin.com/in/triasha-sarkar](https://www.linkedin.com/in/triasha-sarkar) · github.com/triasha72 · triasha72.github.io/Portfolio

MS in Aerospace Engineering (Georgia Tech) with a background in scientific machine learning and numerical methods. I build and own data-generation and training pipelines end-to-end, geometry ingest through GPU training, and my strongest work is finding the failures that pass every standard check: a 15,120-case simulation campaign where 45% of runs terminated with no output and no error, and a GPU training run that spent 11 hours frozen behind a loss curve that looked converged. Prior industry experience validating safety-critical engineering models at Rolls-Royce.

EDUCATION

MS, Aerospace Engineering | Georgia Institute of Technology, Atlanta, GA

Aug 2026 (08/26)

- Graduate Research Assistant, Aerospace Systems Design Laboratory (ASDL), 2025 - 2026.
- Coursework: ISYE 6739 Statistical Methods (A); ISYE 6414 Regression Analysis.

MSc, Aerospace Vehicle Design | Cranfield University, United Kingdom

Feb 2023

First Class Honours · Thesis: geometry optimization under FAA FAR 91.817 constraints in a data-scarce, simulation-expensive design space (SENECA consortium: Rolls-Royce, DLR, ONERA).

B.Tech, Aerospace Engineering | SRM Institute of Science and Technology, India

Jun 2021

First Class Honours.

TECHNICAL SKILLS

Languages: Python (primary), MATLAB (proficient), C++ (working knowledge)

ML & scientific computing: PyTorch (autograd behaviour, custom loss functions, Dataset/DataLoader with variable-length collation), NumPy, SciPy, pandas, scikit-learn, h5py/HDF5, OpenCASCADE, NetworkX, OSMnx, pyDOE2, Jupyter

Engineering practice: Git/GitHub, GPU training on Kaggle and Colab, checkpointing and resumable long-running jobs, numerical precision trade-offs (fp32/fp64), leakage-aware evaluation design, regression and sanity tests against closed-form solutions

Statistics & modeling: Regression and statistical inference; Gaussian process regression; design of experiments; uncertainty quantification and calibration; time-series anomaly detection; clustering; neural networks and CNNs; surrogate and reduced-order modeling; proper orthogonal decomposition; Operator Inference; multifidelity and Monte Carlo methods

ENGINEERING & RESEARCH PROJECTS

Silent Failure Detection in Large-Scale Simulation Campaigns

May - Jul 2026

Georgia Tech ASDL (AE8900) · Advised by Dr. Adam Cox

- Audited a **15,120-case design-of-experiments campaign** and found that 6,804 runs (45%) terminated with no output and no error message; failures that passed every standard validity check and would have silently biased every model built downstream.
- Validated labels against the full burn archive before modeling: every case wrote all ten-output metrics or none, with zero exceptions, establishing the failure label as deterministic rather than noisy.
- Built a leakage-free classifier predicting failure from the **five pre-run geometry inputs alone**, with no post-run data admitted as a feature. A naive random split scored **99.8%** by memorizing geometries duplicated nine times over; collapsing to the 1,680 unique combinations and testing only on geometry never seen in training gave the honest figure of **96.8% accuracy at 0.986 ROC AUC**, against a 55.0% majority-class baseline. The more impressive result was the wrong one.
- Replicated the result on an independent 1,050-case Latin hypercube study under a different sampling scheme (AUC 1.0) and identified a second failure mode at small absolute outer diameter.
- Traced the root cause to an uncapped pressure-balance loop in the solver source; a timeout counter incremented but never checked and wrote and tested a patch converting silent hangs into logged convergence errors with full diagnostics.
- Stage 2 profile classification over the 8,316 clean thrust-time histories; it has been scoped, not yet implemented.

NURBS-BEM Solver and Operator-Learning Pipeline

Oct 2025 - Present

Independent · github.com/triasha72/NURBS_BEM_EMSolver

- Own the full pipeline end to end: STEP geometry ingest through OpenCASCADE, NURBS fitting, a mesh-free boundary-element field solver, parametric dataset generation across 500 geometries at two fidelity levels, and GPU surrogate training on Kaggle.
- Debugged a training run that was silently doing nothing. Loss was bit-identical to five significant figures for 25+ consecutive epochs, including across a scheduler warm restart that raised the learning rate 200x. Root cause: nan_to_num guards in the forward pass, whose registered derivative is $\text{grad} \times \text{isfinite}(x)$; a numerical overflow was rewritten into a finite loss carrying zero gradient. **11 GPU-hours ran with a frozen network and a healthy-looking loss curve.**
- Confirmed the diagnosis from the checkpoint rather than by inference: parsed the serialised tensor storages directly and found AdamW state for only **45 of 61 parameter tensors**. The 16 without state mapped exactly onto a graph structure the code built and never used, so an entire message-passing sub-network had never received a gradient.

- Replaced the physics loss after finding it contradicted the solver's own constitutive law by a factor of the material permeability and added a closed-form regression test the previous implementation fails at **8.0e-2** relative error and the replacement passes at **3.9e-5** with second-order convergence.
- Instrumented the training loop so the same failure class is loud rather than silent: per-parameter gradient finiteness checked before clipping (an Inf gradient makes the clip coefficient zero, and $\text{Inf} \times 0 = \text{NaN}$ poisons every parameter in one step), an alarm on loss unchanged to machine precision, and no NaN suppression anywhere in the forward path.
- Retraining under the corrected pipeline is implemented, not yet run. No surrogate accuracy figure is claimed.

Cross-Sectional Equity Signal: Walk-Forward Backtest

Jul 2026

Independent · github.com/triasha72/Equity-Backtest

- Built a monthly, dollar-neutral cross-sectional equity backtest over **11,304 stock-months** of US large caps (2014–2024): expanding-window validation rather than k-fold, features standardised within each cross-section, and no post-formation data admitted as a model input. A shuffled-target test in the suite asserts the pipeline earns nothing when labels are permuted.
- Found that the momentum signal does not survive transaction costs and was not significant before them: **net Sharpe 0.25 (t = 0.81)**, rising only to $t = 1.07$ with costs set to zero, against an equal-weight universe Sharpe of 1.00.
- Logged and published **all six specifications tested** rather than the best one. The strongest ($t = 1.58$) is reported as noise against a $t = 3.0$ hurdle stated before the result (Harvey, Liu & Zhu, 2016) and traced to portfolio diversification rather than predictive power.

Network Disruption Modeling: Atlanta Mobility Digital Twin

2026 - Present

Independent · github.com/triasha72/atlanta-mobility-resilience-digital-twin

- Graph model of Atlanta's road network (OSMnx, NetworkX, pandas) simulating disruption scenarios across origin-destination pairs and measuring travel-time and accessibility outcomes under stress.

Surrogate Model Comparison: Gaussian Process, Response Surface, RBF

2025 - Present

Independent · github.com/triasha72/Surrogate-model-learning

- Compared Gaussian process, response-surface, and radial-basis-function regressors across benchmark functions.
- Characterized how model assumptions drive error, uncertainty behavior, and failure modes in nonlinear regimes where each estimator breaks down and why.

Cultural and Attitudinal Predictors of Life Expectancy

Summer 2026

Georgia Tech, ISYE 6414 Regression Analysis · Team of four

- Cross-national regression study over 79,143 World Values Survey responses across 53 countries, combining WHO and World Bank structural indicators with attitudinal measures; final model adjusted R^2 0.8097, test RMSE 2.245 years.
- Contributed the literature synthesis and cross-model integration, reconciling observation counts, country definitions, and variable specifications across three independently fitted models and authored the abstract, introduction, and conclusions.
- With $n = 79,000$, conventional p-values flag nearly everything; variables were instead selected by **consensus across 100 subsamples under six criteria** (stepwise AIC/BIC, LASSO, elastic net). Documented pseudo-replication and ecological-fallacy limits constraining what the attitudinal coefficients can support.

Demand and Price-Sensitivity Modeling under Uncertainty (Project EAGLE)

Sep 2025 - Apr 2026

Vehicle Grand Challenge: RIC Enterprise / Georgia Tech ASDL

- Demand modeling and fare-sensitivity analysis for a 750+ passenger transport concept, including fleet-sizing across uncertainty scenarios.

Techno-Economic Modeling under Parametric Uncertainty (Project GREEN TEA)

Sep 2025 - Apr 2026

Georgia Tech ASDL / U.S. Endowment for Forestry

- Techno-economic and well-to-wake life-cycle assessment of alternative fuel pathways under parametric uncertainty; benchmarked model outputs against GREET, CORSIA, and 40BSAF-GREET reference frameworks.
- The resulting model is **in active use by the sponsor**, the U.S. Endowment for Forestry.

INDUSTRY EXPERIENCE

Systems Engineer | Alten India Pvt Ltd, embedded at Rolls-Royce, Bangalore

Oct 2023 - Apr 2025

- Supported airworthiness certification of the **Trent XWB-84 EP** engine; validated engineering models against certification standards under regulatory scrutiny.
- Worked directly with the consequences of model failure in deployed systems; the experience motivating subsequent research on silent failure detection.

LEADERSHIP

Graduate Chair, Women of Aeronautics and Astronautics (WoAA), Georgia Tech Chapter

2025 - 2026