

NIKHIL P. S. BISHT

Open to relocation • Available Summer 2026 (Intern) • U.S. Permanent Resident

Tallahassee, FL

nikhilbisht8@gmail.com • linkedin.com/in/nikhilpsbisht • astrodnerd.github.io • github.com/AstroDnerd

EDUCATION

Florida State University (FSU), Tallahassee, FL (*Expected Graduation Summer 2027*) 2022–Present
PhD Physics, Doctoral Candidate; GPA: 3.54/4.0

Thesis: Forecasting Molecular Cloud Collapse with ENZO using Deep Learning

Birla Institute of Technology and Science (BITS), Pilani, India 2017–2022
MSc. (Hons.) Physics + B.E. (Hons.) Computer Science; GPA: 9.1/10

SKILLS

Data Science & Analytics: Python, SQL, Pandas, NumPy, SciPy, scikit-learn, XGBoost, statistical modeling, experiment design, feature engineering, model evaluation

Machine Learning: Time-series and spatiotemporal modeling, computer vision (CNNs, UNet), deep learning (PyTorch), model interpretability (SHAP)

Data Engineering & Visualization: ETL pipelines, data preprocessing, Parquet/HDF5, Matplotlib, Plotly, Dash

Systems & Software: C/C++, Linux, Git, HPC workflows, reproducible research pipelines

PROFESSIONAL EXPERIENCE

Graduate Research Fellow, *Florida State University, FL* Nov'22 - Present

- **Built and evaluated predictive ML models** on 100M+ spatiotemporal records, improving detection of rare events and forecasting system evolution under noisy conditions.
- **Designed end-to-end data pipelines processing 84+ TB** of raw simulation data, cutting preprocessing and iteration time by 67%, enabling faster experimentation cycles.
- Ran **systematic ablation studies and hyperparameter sweeps** to quantify performance–complexity tradeoffs and guide model selection.
- Delivered insights through **dashboards, visualizations, and presentations** to cross-disciplinary stakeholders, directly influencing modeling and research decisions.

Undergraduate Applied Physics & Machine Learning Researcher, *BITS Goa* Jul'20 - Aug'22

- **Developed probabilistic regression models (Bayesian Neural Networks)** for high-dimensional, noisy datasets, producing calibrated uncertainty estimates to support risk-aware decision-making.
- Evaluated model reliability using **cross-validation, error decomposition, and uncertainty diagnostics**, informing tradeoffs between accuracy, interpretability, and robustness.

PROJECTS

Patient Journey & Access Analytics Project (PHN), *XGBoost, Feature Engineering, SHAP, Streamlit* May'25 - Aug'25

- Analyzed **101K+ patient appointment records** and built a no-show risk classifier (XGBoost), achieving **AUC 0.74**, enabling targeted interventions with **projected \$193K annual savings**.
- Engineered features capturing **scheduling behavior, payer mix, and access constraints**; validated model stability across demographic and payer segments.
- Built an **interactive dashboard** to visualize patient funnels and operational KPIs; used **SHAP** for model interpretability and fairness checks to support responsible deployment.

Lead Engineer (Computer Science Vertical), Project Radio Telescope, *BITS Goa* Mar'18 - Jul'22

- Led a **10+ person cross-functional team** to design and deploy a **real-time data acquisition and signal processing system**, integrating hardware, firmware, and analytics pipelines.
- Built automated pipelines for **streaming ingestion, cleaning, feature extraction, and noise reduction** of high-frequency sensor data, improving signal quality validation and system reliability.
- Defined **system-level performance metrics** to guide iterative improvements across software and hardware components; project presented at national venues and received departmental awards.

SELECTED PUBLICATIONS

PreStCo: Predicting Prestellar Core Formation in Turbulent Molecular Clouds using spatiotemporal neural networks

— **Nikhil Bisht***, David Collins - The Astrophysics Journal (ApJ), 2026 (*in prep*)

- Built a **2.9M parameter physics-informed spatiotemporal neural surrogate** for turbulent 3D simulations, trained on 32 high-resolution MHD runs, enabling future-state prediction up to 1.45 Myr ahead while preserving physical constraints.

- **Developed composite physics-aware losses**, reducing global mass error and spectral mismatch relative to pixel-only baselines on held-out simulations and a training dataset guided by physical constraints, including causally ordered temporal inputs, SO(3)-like rotational and periodic translational symmetries, and multi-scale sub-volume representations.

ML in Astrophysical Turbulence I: Predicting Prestellar Cores in Magnetized Molecular Clouds using eXtreme Gradient Boosting

— **Nikhil Bisht***, David Collins - The Open Journal of Astrophysics (OJA), 2026 (*in review*)

- **Engineered a novel pipeline** that leveraged ~ 2.1 million Lagrangian tracer particles to map complex turbulent gas flows, **transforming unstructured 3D MHD simulation snapshots into a structured, tabular time-series dataset** for predictive modeling.
- Determined the **optimal temporal prediction window for supervised learning** by calculating the integral time scale of turbulence (~ 0.44 Myr) via density autocorrelation functions, ensuring the dataset captured physically relevant evolutionary timescales.
- Demonstrated that complex, non-linear turbulent flows can be effectively modeled as tabular supervised learning problems, **successfully forecasting the 3D trajectories of star-forming cores to within 0.9% error** compared to ground-truth physics simulations.

OTHER TECHNICAL PROJECTS

ENZO Codebase Contribution (High-Performance Computing)

Nov'22 - Mar'24

- Contributed to the **development of ENZO**, an open-source adaptive mesh refinement code used for petascale fluid dynamics simulations.
- Diagnosed and patched **critical memory leaks** and race conditions by utilizing GDB and Valgrind to trace segfaults in multi-node distributed environments. Implemented **dynamic load-balancing algorithms** to optimize grid refinement. Engineered custom linked-list data structures to manage tracer particles in a distributed memory space, reducing overhead by 15%.

Non-Linear Regression & Inference Modeling, *Python, Scipy, Statistical Analysis*

Jun'18 - Dec'18

- Applied **Non-Linear Least Squares regression** to fit multi-component mass models to the observed data to minimize χ^2 **error**, statistically confirming the necessity of Dark Matter halos to explain velocity dispersions in local cluster galaxies.

High-Performance N-Body Simulation Engine

Aug'17 - Apr'18

- Developed a **physics engine** to simulate gravitational interactions between thousands of bodies, addressing the computational bottleneck of direct summation.
- Implemented the Barnes-Hut algorithm using Recursive Octrees (spatial partitioning trees) to approximate long-range forces, successfully **reducing algorithmic time complexity** from $O(N^2)$ to $O(N \log N)$.
- Achieved a **100x** speedup for large particle counts ($N > 10^5$) compared to brute-force methods, enabling real-time visualization of complex dynamical systems.

AWARDS

First Place, 3 Minute Thesis Competition (3MT), Florida State University

Mar'25

State Topper (Goa), National Graduate Physics Examination (NGPE)

Jan'19

Silver Medal, University Physics Competition (UPC)

Nov'18