Prisha Priyadarshini

+1 609-751-6691 | prishapriya31@gmail.com | LinkedIn | Github | Personal Website | Google Summer of Code Work Product

Summary

Rutgers University Computer Science major and Mathematics minor with research experience in interpretable machine learning, multi-agent AI systems, scientific machine learning, and multimodal machine learning. Conducted research at DIMACS, Deep Forest Sciences, and Algoverse, and contributed to DeepChem through Google Summer of Code 2026. Published researcher with first- and second-author work accepted at ICML, NeurIPS, and AAAI workshops. Seeking a research-focused M.S. in Computer Science to investigate reliable and interpretable AI systems and prepare for doctoral study.

Education

Rutgers University – New Brunswick — New Brunswick, NJ
B.S. in Computer Science, Minor in Mathematics; GPA: 3.675/4.00; Dean's List (Spring 2026) — *Expected May 2027*

Research Experience

Machine Learning Research Intern — Sept. 2026 – Present
Deep Forest Sciences — *Remote*

- Developing **ChemBERTa-4**, an OLMo-based chemical language model for molecular representation learning and property prediction; extended **OLMo's vocabulary with ChemBERTa-3's chemical vocabulary** and resized model embeddings while preserving pretrained vocabulary.
- Conducted comparative tokenizer analysis across **OLMo, ChemFM, and ChemBERTa-3**, identifying **9.55%** and **87.97%** vocabulary overlap with OLMo, respectively, to inform chemical tokenization and vocabulary adaptation strategies.
- Benchmarked **OLMo-7B across 9 molecular property datasets** following continued pretraining on **10K–5M PubChem samples**, evaluating toxicity, bioactivity, and physicochemical predictions using ROC-AUC and RMSE.
- Applied **QLoRA fine-tuning** to OLMo-7B, improving performance on **8 of 9 molecular benchmarks** over non-fine-tuned baselines, achieving up to **22.1% relative ROC-AUC improvement** and **16.4% RMSE reduction** in the 100K PubChem configuration.

Machine Learning Researcher — May 2026 – Present
TRACE AI Lab, Rutgers University–New Brunswick — *New Brunswick, NJ*

- Conducting independent research under **Dr. Linda Ness and Dr. Lesia Semenova** on **fairness, interpretability, and threshold stability** across sparse decision-tree Rashomon sets, investigating how feature representation influences predictive accuracy and model behavior.
- Benchmarked **4 experimental configurations** (Standard ArborEnum, Anytime ArborEnum, PRAXIS with TGB, and PRAXIS with Unbiased Fairness Binning) across **3 fairness datasets**, finding that fairness-aware binning generally reduces group disparities at the cost of predictive accuracy.
- Evaluated **3 fairness metrics** (statistical parity, equal opportunity, and equalized odds) across up to **1,000 near-optimal trees per configuration**; developed **leaf-level fairness attribution** methods showing that disparities concentrate in a small subset of influential leaves rather than being explained by tree complexity alone.
- Conducted **1,000-resample bootstrap analyses with 95% confidence intervals** and compared up to **50 most-fair and 50 least-fair trees** using prediction self-consistency, finding that improved group fairness does not consistently imply greater prediction stability.

Artificial Intelligence Researcher — Sept. 2025 – June 2026
Algoverse AI — *Remote*

- Built a multi-agent framework to study consensus formation, running 20-round deliberation experiments across **GPT-4.1, Mistral, LLaMA, and Phi** models on subjective and objective benchmarks.
- Introduced quantitative metrics (inter-agent disagreement rate, model deference rate), identifying hierarchical influence patterns **(up to 80% small-to-large model deference)**, and developed a rotation-based framework separating model identity from response quality.
- Showed that prompting strategies can control multi-agent behavior, increasing disagreement and reducing deference, with implications for system reliability. Accepted to the ICML 2026 Pluralistic Alignment Workshop.

Lead Machine Learning Researcher Via Rutgers MBS Exchange — Jan. 2026 – May 2026
DIMACS (Center for Discrete Mathematics & Theoretical Computer Science) — *New Brunswick, NJ*

- Led a student research team under **Dr. Linda Ness**, investigating whether feature representation can close

predictive-performance gaps between **interpretable sparse decision trees and gradient-boosting models**.

- Benchmarked **7 machine learning algorithms** (GOSDT, SPLIT, LicketySPLIT, LicketyRESPLIT, XGBoost, LightGBM, and CatBoost) across **6 real-world datasets**, including datasets containing up to **284,807 observations**, evaluating predictive accuracy, recall, and Rashomon-set size.
- Conducted controlled experiments comparing **TGB and non-TGB feature representations** across **3 estimator configurations and 3 tree depths**, finding that threshold representation can substantially improve sparse-tree competitiveness with boosting models, while increasing threshold count does not necessarily improve accuracy.

Artificial Intelligence Researcher June 2025 – Dec. 2025
Algoverse AI *Remote*

- Contributed to a hierarchical scene-captioning pipeline integrating dynamic stride window selection and multimodal chain-of-thought reasoning for temporally coherent caption generation.
- Designed evaluation frameworks **(BLEU-4, METEOR, CIDEr, BERTScore)** and ran ablations on frame sampling **(5, 20, 40)** and model aggregators **(Phi, Mistral, Qwen3).**
- Ran large-scale YouCook2 experiments on **30 RTX A6000 GPUs**, achieving **17% higher CIDEr than GPT-4o and 14% over VLLaMA-3**. Accepted for an oral presentation at NeurIPS 2025 (7HVU) and to AAAI 2026 (AI4EDU).

Algorithms Research Shadow Jan. 2025 – May 2025
The College of New Jersey *Ewing, NJ*

- Researched classical and modern folding algorithms under Dr. Dimitris Papamichail (Nussinov, Zuker, LinearFold, MCTS); implemented sparse dynamic programming to prune redundant structures and reduce runtime.
- Conducted SLURM-based large-scale experiments using ViennaRNA across thousands of RNA sequences.

Publications

- **Directional Influence and Consensus Formation in Multi-Agent Systems.**
  Prisha Priyadarshini, Aryan Shrivastava.
  Accepted to ICML 2026 Pluralistic Alignment Workshop.
- **DynaStride: Dynamic Stride Windowing with MMCoT for Instructional Multi-Scene Captioning.**
  Eddison Pham, ***Prisha Priyadarshini***, et al.
  Accepted to NeurIPS 2025 7HVU Workshop (Oral) and AAAI 2026 AI4EDU Workshop.

Industry Experience

AI Engineering Intern Aug 2026 – Present
DrRx.ai *Remote*

- Developing a RAG pipeline leveraging **BM25S** retrieval and **Qwen2.5-7B-Instruct** over **2K+ pharmaceutical documents**, integrating data from PostgreSQL and **AWS S3.**

AI Fellow May 2026 – Oct. 2026
AI4ALL *Remote*

- Developed an end-to-end **Braille-to-speech computer vision pipeline** integrating YOLO11n, EfficientNet-B0, Liblouis, and Piper TTS, achieving **96.15% detection mAP50, 99.41% classification accuracy, and 97.32% character coverage**.
- Trained and evaluated detection and classification models across **4 datasets and 5 model families**, scaling to **10,117 augmented images and 193K+ character crops** and performing end-to-end error analysis to identify false-positive detection and reading-order failures.

Machine Learning Intern May 2026 – Aug 2026
Regeneron *Troy, NY*

- Deployed a **Streamlit** app for the QA Validation and Operations team after evaluating **10+ ML approaches** and selecting **TabPFN Regressor**.
- Engineered **40 trend-based features** from **500K+ GMP records** to predict **remaining useful life (RUL)** and classify equipment-breakdown risk using an **RUL threshold**.
- Implemented **K-fold cross-validation** with **single-class fold rejection** and **data-leakage safeguards**, achieving **67.8% ROC-AUC**.

Google Summer of Code 2026 Contributor (DeepChem) May 2026 – Aug 2026
Google Summer of Code *Remote*

- Selected as 1 of 5 from 200+ applicants to extend DeepChem's support for large-scale molecular language models.
- Extended the **HuggingFaceModel class** by implementing the **generate() API**, enabling batched autoregressive inference, tokenizer padding, and device-aware execution.
- Developed an **OLMo wrapper class** supporting **regression**, **classification**, multitask learning, and causal language modeling with custom prediction heads, pretrained checkpoint loading, mixed-precision inference, **LoRA/QLoRA**, and **multi-GPU support**.
- Benchmarked QLoRA-adapted OLMo across **9 molecular datasets**, achieving **93.1% ROC-AUC** on BBBP and **92.2%** on BACE with random splits, and **91.6%** on ClinTox with scaffold splits.

AI Associate Developer Oct. 2025 – Jan. 2026
Insurity – SpatialKey Team *Remote*

- Built geospatial preprocessing pipelines over **6M+ records,** integrating climate and peril data via Pandas and GeoPandas.
- Trained LightGBM peril classifiers with SMOTE balancing achieving **+10% accuracy and F1 score**.
- Built a **U-Net++ change-detection system** to detect before/after wildfire satellite imagery by adding shadow normalization and shadow masking to reduce false positives from terrain illumination .

Projects

Kalinga — Rutgers HealthHack 2026 Honorable Mention | *Python, PyTorch, Hugging Face, SetFit, Scikit-learn*

- Developed a multilingual, multi-label NLP classifier for hospital request routing across **4 languages and 6 departments**; curated a labeled dataset and benchmarked TF-IDF/logistic regression and BGE embeddings.
- Fine-tuned **multilingual MiniLM** using SetFit and contrastive learning; achieved **0.8963 macro F1 / 83.73% accuracy** on grouped test data and **0.8339 macro F1 / 79.03% accuracy** on an independent holdout set, using StratifiedGroupKFold for group-aware splitting.
- Implemented emergency escalation and out-of-scope abstention rules for safety-aware routing; published the fine-tuned classifier on **Hugging Face**.

Knowledge Distillation for Reasoning in MoE and Dense LLMs | *Python, PyTorch, Hugging Face, LoRA*

- Developed sequence-based and logit-based knowledge distillation pipelines to transfer chain-of-thought reasoning from **GPT-OSS 120B and LLaMA 70B** teacher models to a **3B student** using LoRA-based fine-tuning.
- Implemented **temperature-scaled KL divergence**, top-k logit extraction, and cross-tokenizer alignment to match teacher-student output distributions despite vocabulary mismatches.
- Achieved up to **4.5% accuracy improvement** and **96%+ BERTScore** on SciBench and TheoremQA; demonstrated that logit distillation consistently outperformed sequence-based distillation.

GenreBlender | *Python, PyTorch, MusicGen, Streamlit, Scikit-learn*

- Architected an end-to-end generative AI system using **Meta's MusicGen** to synthesize 10-second hybrid audio clips with controllable genre blending via a continuous weighting parameter ($\alpha \in [0, 1]$).
- Engineered a GTZAN preprocessing pipeline with feature selection, normalization, and **GroupShuffleSplit** validation; trained a custom **4-layer PyTorch MLP** achieving **92% validation accuracy**.
- Developed a probability-based evaluation framework comparing predicted genre distributions against weighted target mixtures; deployed an interactive **Streamlit application** for audio generation and quantitative evaluation.

Professional Development & Extracurriculars

- **Workshop Reviewer**, NeurIPS 2026 AI4GOOD Workshop (Sept 2026):
  Reviewed research submissions through the OpenReview peer review process.
- **Workshop Reviewer**, ICML 2026 Pluralistic Alignment Workshop (May 2026):
  Reviewed research submissions through the OpenReview peer review process.
- **Mathematics Grader**, Rutgers University (Dec. 2025 – May 2026):
  Graded homework and exams for Multivariable Calculus across 3 sections.
- **Peer Tutor**, TCNJ (Sept. 2024 – May 2025):
  Tutored Data Structures, Algorithms, Linear Algebra, and Statistics.
- **Recruitment Chair**, Kappa Theta Pi (Dec. 2024 – May 2025):
  Led recruitment and onboarded 10 new members.

- **Student Athlete** , NCAA Division III Tennis, TCNJ (Aug. 2023 – May 2024):
  Balanced 20+ hrs/week training with academics; earned All-NJAC recognition.

Technical Skills

AI & ML: PyTorch, Hugging Face, Transformers, scikit-learn, OpenCV
Research Domains: Multi-Agent Systems, Interpretable ML, Multimodal ML
Languages & Libraries: Python, C, C++, Java, JavaScript; NumPy, Pandas, GeoPandas, Matplotlib, Seaborn
Cloud, Systems & MLOps: AWS, Azure, Docker, CI/CD, GitHub, pytest, Jira, Anaconda, Vercel
Certifications: AWS Certified Machine Learning – Specialty; AWS Certified Cloud Practitioner