

Prisha Priyadarshini

609-751-6691 | prishapriya31@gmail.com | [LinkedIn](#) | [Github](#) | [Personal Website](#)

EDUCATION

Rutgers University – New Brunswick

B.S. in Computer Science, Minor in Mathematics; GPA: 3.675/4.00; Dean's List (Spring 2026)

New Brunswick, NJ

Expected May 2027

RESEARCH EXPERIENCE

Lead Machine Learning Researcher

Jan. 2026 – Present

DIMACS (Center for Discrete Mathematics & Theoretical Computer Science)

New Brunswick, NJ

- Leading an empirical study under Dr. Linda Ness comparing interpretable decision-tree methods (**SPLIT, LicketyRESPLIT, GOSDT, PRAXIS**) with SOTA boosting models (**XGBoost, CatBoost, LightGBM**) across diverse tabular datasets and extensive hyperparameter sweeps.
- Demonstrated that **ThresholdGuessBinarizer (TGB)** identifies compact, stable sets of predictive thresholds, with additional threshold guesses and tree depth exhibiting **diminishing accuracy gains while substantially expanding the Rashomon set**.
- Characterized **Rashomon-set diversity, stability, feature importance, and algorithmic fairness**, showing that similarly accurate trees can differ substantially in structure and fairness and that **dataset characteristics—rather than model complexity alone—strongly govern solution diversity and 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.

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, Adrian Maliackel, Kanishk Bandi, Cristian Meo, Kevin Zhu.
Accepted to NeurIPS 2025 7HVU Workshop (Oral); AAAI 2026 AI4EDU Workshop.

INDUSTRY EXPERIENCE

AI Engineering Intern

Aug. 2026

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**.

Machine Learning Intern

May 2026 – Aug 2026

Regeneron

Troy, NY

- Deployed a **Streamlit** application for the QA Validation and Operations Team using **TabPFN Regressor and Classifier models**, engineering **40 trend-based features** from **500K+ GMP manufacturing records** to predict equipment breakdowns, achieving a **67.8% ROC-AUC**.

Open Source Developer (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**, implementing 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 Fellow

May 2026 – Present

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.

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

- Knowledge Distillation for Reasoning in MoE and Dense LLMs** | *Python, PyTorch, HuggingFace* (Apr. 2026): Built an offline distillation pipeline using LoRA (<1% trainable parameters) to transfer chain-of-thought reasoning and token-level distributions from large LLMs to a 3B student via sequence and logit distillation (KL divergence + temperature scaling), achieving 4.5% accuracy gains and 96%+ BERTScore on SciBench and TheoremQA.
- GenreBlender** | *Python, PyTorch, Streamlit* (Feb. 2026): Engineered a generative AI system combining Meta's MusicGen with a 4-layer MLP (92% validation accuracy) to create and quantitatively evaluate controllable music genre blends via a weighted probability framework.
- PocketRAG** | *Python, Flask, FAISS, Gemini, AWS, Docker* (Sept. 2025): Engineered a production RAG pipeline with document chunking, FAISS vector retrieval, and Gemini inference; containerized and deployed on AWS EC2 (Docker + Gunicorn) with sub-second latency.

PROFESSIONAL DEVELOPMENT & EXTRACURRICULARS

- 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