TATHAGATA DEBNATH

PhD Candidate in Computer Science

(575) 621-1224 · tirtha.debnath@gmail.com · GitHub · LinkedIn · Google Scholar · ORCID · Website

SUMMARY

PhD candidate with hands-on experience training and fine-tuning LLMs (DPO/RLHF on multi-GPU H100 clusters), building multi-agent orchestration systems, and designing retrieval-augmented generation pipelines. Published algorithm designer (IEEE TPAMI, IEEE/ACM TCBB; 81+ citations) with 2 CRAN R packages for genomic data analysis. Research spans agentic AI systems, reinforcement learning from human feedback, information retrieval, and computational biology.

TECHNICAL SKILLS

- **Languages:** Python, R, C/C++, JavaScript, SQL, Bash
- **ML/DL:** PyTorch, DPO/RLHF Fine-Tuning, Hugging Face Transformers & TRL, vLLM, scikit-learn, XGBoost, SHAP, Nested Cross-Validation, Feature Engineering
- **LLM/Agents:** LangGraph, LangChain, Claude API, OpenAI API, Prompt Engineering, Retrieval-Augmented Generation (RAG), FAISS, Vector Databases, Multi-Agent Systems, MCTS for Code Generation
- **Infrastructure:** Docker, Git, Weights & Biases, LangSmith, FastAPI, Streamlit, React, Multi-GPU Training (H100)
- **Statistical Methods:** Compositional Data Analysis (CLR), Hypothesis Testing, Cross-Validation, Dimensionality Reduction (PCA, UMAP)
- **Research Software:** CRAN R Package Development, Bioconductor, LaTeX

PROJECTS

CodeQ: Self-Improving Code Debugging Agent *2025–Present*

[PyTorch, DPO, MCTS, vLLM, Hugging Face TRL, Weights & Biases, H100]

- Built an autonomous bug-repair agent combining Monte Carlo Tree Search (MCTS) exploration with Direct Preference Optimization (DPO) self-improvement on Qwen2.5-Coder-7B-Instruct, operating across two shared H100 machines (Machine A: inference/MCTS rollouts; Machine B: DPO training/batch workloads).
- Achieved +2 percentage-point improvement to 38.0% solve rate on SWE-bench Lite after Round 1 DPO training, with Round 2 incorporating step-150 checkpoint selection, early stopping, and few-shot examples for further gains.
- Implemented full training pipeline: MCTS rollout collection generating pass/fail preference pairs, DPO fine-tuning with Weights & Biases logging, automated evaluation harness, and shared-filesystem coordination between inference and training machines.
- Designed evaluation framework including LLM-as-judge scoring with Cohen's kappa inter-rater reliability and ablation studies isolating contributions of MCTS depth, DPO rounds, and checkpoint selection.

Parallel Multi-Agent Code Generation System *2024–Present*

[Python, LangGraph, Claude API, Pydantic, Docker, Streamlit, LangSmith]

- Engineered a DAG-based multi-agent code generation system with concurrent coder workers orchestrated via LangGraph, shared Pydantic state management, and configurable revision loops.
- Implemented Docker-sandboxed test execution with real pytest validation, achieving 79/79 test pass rate on generated code with automatic failure-driven iterative revision up to a configurable max iteration limit.
- Built real-time Streamlit UI with per-agent status indicators and dynamic model selection (Haiku/Sonnet/Opus) enabling users to tune cost-performance tradeoffs per pipeline run.
- Integrated LangSmith observability for agent reasoning traces, token usage tracking, latency profiling, and cost analysis across all pipeline stages.

Self-Evolving Code Generator *2024*

[Python, LangGraph, Claude API, LLM-as-Judge]

- Extended the multi-agent codegen V1 architecture with an automated LLM-as-Judge prompt evolution loop: generated code is scored on correctness, efficiency, and style, and scoring feedback iteratively refines the generation prompts without human intervention.
- Demonstrated measurable prompt quality improvement over successive evolution rounds, validating the self-improvement paradigm later scaled up in the CodeQ project with DPO training.

NutriBot RAG System *2024*

[FAISS, BM25, Reciprocal Rank Fusion, FastAPI, Claude API, Caching]

- Built a domain-specific Retrieval-Augmented Generation system over nutrition research literature with hybrid search combining FAISS semantic similarity, BM25 lexical matching, and Reciprocal Rank Fusion reranking, outperforming either method alone across conceptual and terminology-heavy queries.
- Reduced per-query LLM API costs by approximately 50% through four optimization layers: MD5-based answer caching, semantic relevance score threshold gating, low-relevance chunk filtering, and strict 512-token output budgeting.
- Designed a 30-question evaluation pipeline measuring Keyword Hit Rate, Mean Reciprocal Rank (MRR), and Answer Keyword Coverage with free retrieval-only tuning mode for iterative refinement before incurring generation costs.

Gut Microbiome Disease Classifier *2023–2024*

[Python, scikit-learn, XGBoost, SHAP, CLR Transformation, Nested CV]

- Developed a multi-disease classifier achieving 0.907 weighted AUC across 3 disease classes (Crohn's Disease, Type 2 Diabetes, Healthy controls) from 236 clinical 16S rRNA amplicon samples using CLR compositional log-ratio preprocessing.
- Implemented a 4-method feature selection pipeline (Mutual Information, ANOVA F-score, RFE, Boruta) reducing 4,203 taxa to 52 discriminative biomarkers with cross-method consensus analysis.
- Compared 4 candidate models using nested 5×5 stratified cross-validation; XGBoost achieved tightest variance (AUC 0.896 ± 0.033).
- Conducted SHAP-based biomarker discovery identifying clinically validated microbial signatures; performed external validation on 759-sample independent cohort with documented cross-platform batch effect analysis.

Federated Learning for Health Prediction *2024*

[Python, Flower Framework, PyTorch, Federated Averaging]

- Implemented a federated learning system using the Flower framework on the Diabetes Health Indicators dataset, training models across distributed clients without centralizing sensitive health data.
- Evaluated convergence behavior and accuracy trade-offs between federated and centralized training baselines.

OptCirClust: Fast Circular Clustering Algorithm *2021–2023*

[R, C++, CRAN Package Development, Computational Geometry]

- Designed and published an $O(KN \log^2 N)$ circular clustering algorithm processing round genomes (bacterial, mitochondrial, chloroplast) orders of magnitude faster than brute-force $O(KN^2)$ approaches.
- Developed CRAN R package *OptCirClust* with C++ backend for computational efficiency, comprehensive documentation, and unit test coverage; package continues to be downloaded and cited by bioinformatics researchers.

CircularSilhouette: Linear-Time Quality Metric *2022–2023*

[R, C++, Computational Statistics, CRAN]

- Developed a linear-time O(N) silhouette coefficient for circular cluster quality assessment, 3000× faster than standard $O(N^2)$ computation on 1M-point datasets.
- Published as IEEE TPAMI paper and companion CRAN R package *CircularSilhouette*.

PUBLICATIONS

[1] H. Bendzus-Mendoza, A. Rodriguez, **T. Debnath**, C.D. Bailey, H.A. Luker, I.A. Hansen. "Radiation exposure induces genome-wide alternative splicing events in *Aedes aegypti* mosquitoes." *Scientific Reports*, 2025.

[2] Y. Chen, **T. Debnath**, A. Cai, M. Song. "Circular Silhouette and a Fast Algorithm." *IEEE Transactions on Pattern Analysis and Machine Intelligence (TPAMI)*, 45(11), 2023. DOI: 10.1109/TPAMI.2023.3310495. **[CRAN: CircularSilhouette]**

[3] **T. Debnath**, M. Song. "Fast Optimal Circular Clustering and Applications on Round Genomes." *IEEE/ACM Transactions on Computational Biology and Bioinformatics (TCBB)*, 18(6):2061–2071, 2021. DOI: 10.1109/TCBB.2021.3077 **[CRAN: OptCirClust]**

[4] A.D. Lopez, **T. Debnath**, M. Pinch, I.A. Hansen. "Phosphoproteomics Analyses of *Aedes aegypti* fat body reveals blood meal-induced Signaling and Metabolic Pathways." 2024.

[5] S. Bardhan, S. Nath, **T. Debnath**, D. Bhattacharjee, M.K. Bhowmik. "Designing of an inflammatory knee joint thermogram dataset for arthritis classification using Deep Convolution Neural Network." *Quantitative InfraRed Thermography Journal*, 19(3):145–171, 2020.

[6] M.K. Bhowmik, **T. Debnath**, D. Bhattacharjee, P. Dutta. "EF-Index: Determining number of clusters (K) to estimate number of segments (S) in an image." *Image and Vision Computing*, 88:29–40, 2019.

[7] S. Bardhan, M.K. Bhowmik, **T. Debnath**, D. Bhattacharjee. "RASIT: Region shrinking based accurate segmentation of inflammatory areas from thermograms." *Biocybernetics and Biomedical Engineering*, 38(4):903–917, 2018.

[8] Additional publications in computational biology and computer vision — full list at Google Scholar (81+

citations total).

SOFTWARE

- **OptCirClust** — CRAN R package. Fast, optimal circular clustering in O(KN $\log^2$N). Applied to bacterial, mitochondrial, and chloroplast genomes. CRAN · GitHub
- **CircularSilhouette** — CRAN R package. Linear-time silhouette for circular data, 3000× faster than standard. CRAN · GitHub
- **Pohori** (in development) — Model-free statistical method for alternative splicing detection in RNA-seq (short-read and long-read). GitHub

RESEARCH EXPERIENCE

Graduate Research Assistant *2020–Present*
Department of Computer Science, New Mexico State University ·

- Designed the FOCC algorithm for optimal circular clustering (O(KN $\log^2$N)), published in IEEE/ACM TCBB and released as CRAN R package *OptCirClust*.
- Co-developed circular silhouette method for unsupervised cluster number selection, published in IEEE TPAMI and released as CRAN R package *CircularSilhouette*.
- Developing Pohori, a model-free statistical method for detecting alternative splicing in short-read and long-read RNA-seq data with applications to other omics datasets.

Research Collaborator, Hansen Lab (Molecular Vector Physiology) *2023–2025*
Department of Biology, New Mexico State University

- Conducted computational analysis of radiation-induced alternative splicing in *Aedes aegypti* using rMATS on RNA-seq data, identifying 197 genes with splicing changes across Hippo and Notch signaling pathways (*Scientific Reports* 2025).
- Contributed bioinformatics analysis to phosphoproteomics study of blood meal-induced signaling in mosquito fat body (2024).

Research Collaborator, Castillo Lab *2021–2023*
Department of Biology, New Mexico State University

- Implemented bioinformatics pipelines for alignment, de novo assembly (Trinity), and differential immune response analysis in snail (*Physella acuta*) and squid RNA-seq datasets.

Volunteer Research Associate *2017–2018*
Department of Computer Science, Tripura University

- Assisted in setting up computer vision pipelines for real-time object detection in CCTV footage.
- Applied deep-learning based CNNs on arthritis knee joint thermograms.

TEACHING EXPERIENCE

Graduate Teaching Assistant *2018–2021*
Department of Computer Science, New Mexico State University

- TA for Data Structures & Algorithms (CS 372), Analysis of Algorithms (CS 570), Artificial Intelligence I (CS 475), Machine Programming & Organization (CS 273), C++ (CS 151), Java (CS 172).
- Graded assignments, held office hours, and delivered lectures as substitute instructor.

EDUCATION

- **Ph.D. in Computer Science** Expected 2026
  New Mexico State University · Advisor: Satyajayant Misra · GPA: 4.00
  Dissertation: Algorithms for Post Quantum secure Cryptographic techniques
- **M.S. in Computer Science** Spring 2026
  New Mexico State University · *Focus: AI/ML*
- **M.Tech in Computer Science and Engineering** 2017
  Tripura University · GPA: 9.83/10 · **Gold Medalist**
  Focus: Computer Vision — thermal image analysis and image segmentation
- **B.Tech in Computer Science and Engineering** 2015
  National Institute of Technology, Agartala · GPA: 7.62/10

AWARDS & HONORS

- Gold Medal, M.Tech — Branch topper, Tripura University (2017)
- Biopattern Award, New Mexico State University

- Ph.D. Tuition Scholarship, New Mexico State University (2018–Present)

ACADEMIC SERVICE

- Paper reviewer for *Springer Nature* and *IEEE BIBM*.

RELEVANT COURSEWORK

Generative Artificial Intelligence, Artificial Intelligence II, Data Mining, Applied ML I, Cloud Computing, Algorithms, Analysis of Algorithms, Computational Geometry, Statistical Methods, Programming Language Structure I.

CERTIFICATIONS

- Deep Learning Specialization (Coursera — Andrew Ng): Neural Networks and Deep Learning, Improving Deep Neural Networks, Structuring ML Projects, CNNs, Sequence Models.
- Finding Hidden Messages in DNA (Bioinformatics I), Coursera — with Honors.

LEADERSHIP

- President, Computer Science Graduate Student Association, NMSU (2020–2021)
- President, Indian Student Association, NMSU (2021–2022)
- Vice President, CS Graduate Student Association, NMSU (2019–2020)
- Vice President, Indian Student Association, NMSU (2020–2021)