

Obed Junias

✉ obed.junias@colorado.edu | 🌐 obedjunias.com | 🔄 obedjunias19 | 🌐 obed-junias | 📞 720-266-0046

RESEARCH SUMMARY

I study **model collapse in language models**, focusing on how recursive training on model-generated data degrades distributional diversity, factual reliability, and long-term model quality. My current work combines **measurement, detection, and controlled experimentation** to identify early signals of collapse under synthetic-data feedback loops.

More broadly, I am interested in reliable reasoning, evaluation under distribution shift, and scalable systems for studying foundation models.

Long-term goal: develop principled methods to measure, detect, and mitigate model collapse in foundation models

PRIMARY RESEARCH FOCUS

Model Collapse in Generative AI, Synthetic Data, Recursive Training, Reliability of Foundation Models

SECONDARY INTERESTS

Commonsense Reasoning, LLM Evaluation and Benchmarking, Responsible AI, Scalable LLM Systems

EDUCATION

M.S. in Computer Science, University of Colorado Boulder Aug 2024 – May 2026

Courses: NLP, Deep Natural Language Understanding, Neurosymbolic NLP Methods (GPA: 4.0/4.0)

B.E. in Computer Science, BMS College of Engineering Aug 2017 – Aug 2021

Relevant Courses: Algorithms, Databases, Operating Systems, Machine Learning

PUBLICATIONS

- **Junias, Obed**, and Pacheco, Maria Leonor. *LOGICAL-COMMONSENSEQA: A Benchmark for Logical Commonsense Reasoning*. **ACL 2026 (Main Conference)** [paper]
- **Junias, Obed***, et al. *Assessing Algorithmic Bias in Language-Based Depression Detection: A Comparison of DNN and LLM Approaches*. **IEEE-EMBS International Conference on Biomedical and Health Informatics (IEEE-BHI), 2025** [paper]

RESEARCH EXPERIENCE

Measuring Model Collapse under Recursive Summarization Training Feb 2026 - Present

Advisor: [Maria Leonor Pacheco](#), *BLAST Lab*

- Designing a controlled recursive training pipeline where human-written corpora are progressively replaced with **model-generated (synthetic) data** across generations.
- Training language models under synthetic-data feedback loops to study degradation in **distributional diversity, factual reliability, and knowledge fidelity**.
- Developing measurement methods for collapse using **entropy-based metrics, lexical diversity, and distributional divergence**.
- Investigating early warning signals such as **entropy decay, repetition, and detectability of synthetic text** as indicators of collapse before major downstream failures.

Structured Commonsense Reasoning and Benchmarking June 2025 – Present

Advisor: [Maria Leonor Pacheco](#), *BLAST Lab*, *Associated publication listed above*

- Introduced **LOGICAL-COMMONSENSEQA**, a benchmark designed to isolate compositional and multi-fact reasoning failures in LLMs.
- Identified systematic mismatches between **chain-of-thought fluency and logical correctness**, demonstrating cases where coherent reasoning traces yield incorrect conclusions.
- Developed structured evaluation protocols to analyze reasoning depth and failure modes.
- Designed entailment-like reasoning representations to study how performance varies with task structure and inference depth.

Responsible LLM Evaluation in Mental Health Applications Jan 2025 – Nov 2025

Advisor: [Theodora Chaspari](#), *HUBBS Lab*, *Associated publication listed above*

- Built **large-scale LLM evaluation pipelines** to assess demographic bias and subgroup performance of LLMs (LLaMA, GPT-4) for mental health classification tasks.
- Studied how prompting strategies and fairness-aware losses affect reasoning consistency and robustness in low-resource and imbalanced data settings.
- Developed diagnostic metrics to characterize bias and reasoning instability beyond aggregate accuracy.
- Led research design for a **peer-reviewed publication**, demonstrating limitations of LLM reliability in sensitive domains.

Efficient Multi-Agent LLM Inference with Parameter-Efficient Adaptation Aug 2025 – Dec 2025

- Designed a modular inference framework to study how parameter-efficient adapters enable scalable multi-agent LLM systems.
- Implemented dynamic **LoRA adapter loading and scheduling** under strict GPU memory constraints, streaming adapters from CPU to GPU at runtime.
- Conducted controlled experiments comparing adapter-based specialization against full-model switching, analyzing latency, throughput, and memory efficiency.
- Demonstrated that adapter-based agents reduce cumulative switching overhead by **20–90×**, enabling complex agent workflows on constrained hardware.

RELEVANT RESEARCH PROJECTS

Lost in Plot: Contrastive Learning for Logical Semantic Retrieval

- Formulated a semantic retrieval task for vague queries, targeting underspecified and ambiguous natural-language descriptions requiring semantic rather than lexical matching.
- Designed a dual-encoder model (BERT) with multi-task contrastive learning objectives to align free-form user descriptions with structured movie metadata.
- Constructed a 100K+ synthetic dataset to enable scalable training and controlled evaluation.
- Evaluated retrieval quality using Recall@K and MRR, outperforming a GPT-4 few-shot baseline.

Medical Ethics Assessment of Large Language Models

- Designed evaluation protocols to assess ethical reasoning of LLMs in medical decision-making scenarios.
- Built a multiple-choice benchmarking pipeline to analyze how external knowledge affects reasoning consistency.
- Analyzed rationales and confidence signals to identify systematic failure modes in ethically sensitive settings.

INDUSTRY EXPERIENCE

Senior Member of Technical Staff at Oracle Corporation Aug 2021 - Aug 2024

- Designed and implemented scalable automation and validation frameworks (**TestNG, BATS**) for end-to-end testing of **Oracle REST Data Services (ORDS)** and **SQLcl**, improving validation efficiency by **95%**.
- Led technical development, validation and delivery of a new Oracle Cloud **DB Tools** feature, coordinating a 15-engineer effort and achieving production readiness within 2 months.

Machine Learning Engineer Intern at Hewlett Packard Enterprise Feb 2021 - Jul 2021

- Built OCR-based automation pipelines for processing unstructured enterprise documents, integrating learned document extraction into production workflows.
- Evaluated model performance across noisy and heterogeneous inputs, contributing to improvements in robustness and downstream automation reliability.

SKILLS

Languages & Frameworks	Python, PyTorch, Triton, LangGraph, SQL, Transformers, C++
LLM Systems	LLM Serving, LLM Evaluation, RAG Pipelines, Multi-Agent Systems and Tools
LLM Training	Pre-Training, Supervised Fine-Tuning, Transformers, PyTorch DDP
ML Infrastructure	CUDA, vLLM, FlashAttention, Quantization, GPU Memory Optimization

AWARDS AND RECOGNITION

- **NSF–EMBS–Google Young Professional NextGen Scholar**, *IEEE EMBS BHI Conference*, 2025