Obed Junias

obed.junias@colorado.edu | obedjunias.github.io | obed-junias | 720-266-0046

Research Statement

I am a graduate researcher exploring **Commonsense Reasoning, Neuro-symbolic NLP, and Responsible ML**. My research focuses on developing reasoning frameworks that integrate structured logic with neural representations, with the goal of improving the interpretability, factual consistency, and fairness of LLMs. I aspire to contribute to foundational AI research at the intersection of cognition, reasoning, and ethical deployment.

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

Current Research Projects and Experience

Commonsense Reasoning with Logical Entailment Trees, *BLAST Lab* — June 2025 – Present
Advisor: Dr. Maria L. Pacheco, CU Boulder

- Investigating evaluation of multi-fact commonsense QA using structured compositional and logical reasoning.
- Designing a neuro-symbolic framework for evaluating structured reasoning combining informal logic, logical structure, entailment modeling, and social consensus reasoning.
- Creating a commonsense QA benchmark with logically composed options and multi-step reasoning trees.
- Benchmarked reasoning quality using LLM prompting strategies and step-level inference quality.
- Targeting submission to **ACL Findings or EMNLP 2026.**

Responsible AI Researcher, Institute of Cognitive Science, CU Boulder — Jan 2025 – Present
Advisor: Dr. Theodora Chaspari, HUBBS Lab

- First-author (**IEEE-EMBS 2025**; arXiv:2509.25795) on algorithmic bias in LLM-based depression detection.
- Led large-scale fairness evaluation of LLaMA, GPT-4, and other LLMs for mental health classification.
- Implemented in-context prompting strategies and fairness-aware losses to mitigate bias in low-resource settings.
- Developed diagnostic probes and metrics to quantify bias in model outputs, influencing model selection.

Adapter-Based Multi-Agent Systems for Scalable LLM Inference — Aug 2025 – Present
Course: Systems for Machine Learning, University of Colorado Boulder

- Building a modular serving system where each agent is represented by a dynamic, independently trainable LoRA adapter on top of a frozen base model.
- Designing infrastructure for dynamic adapter loading, chaining, caching, and orchestration.
- Measuring system metrics (latency, throughput) and quality trade-offs across adapter-based inference pipelines.
- Project focuses on scaling lightweight agents, serving efficiency, and modularity for real-world deployment.

Publications

- **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 (BHI), 2025. (Accepted). arXiv:2509.25795 (*Presenting Author)
- **Junias, Obed**, et al. "*Commonsense Reasoning via Logical Entailment Trees: Towards Neuro-symbolic Inference Alignment.*" (In Preparation) Target: *ACL Findings / EMNLP 2026.*

Skills

Languages	Python, C++, SQL
AI/ML Frameworks	PyTorch, TensorFlow, Transformers, LangChain, CUDA
Tools & Platforms	Git, Docker, Kubernetes, Redis, MongoDB, Selenium
Cloud Technologies	GCP, Oracle Cloud (OCI), AWS

Relevant Projects

- **Medical Ethics Assessment of LLMs**

  - **Objective:** Assessed the ethical reasoning capabilities of large language models (LLMs) in clinical contexts, an underexplored dimension of Responsible AI.
  - **Approach:** Designed a controlled multiple-choice evaluation pipeline using Retrieval-Augmented Generation (RAG) to measure the effect of external knowledge on ethical reasoning.
  - **Findings:** Through logit-based confidence and rationale analysis, revealed that while LLMs show partial moral understanding, they lack reliability and depth for ethically sensitive tasks.

- **Lost in Plot: Contrastive Learning for Movie Retrieval**
  - **Objective:** Developed a system for "tip-of-the-tongue" movie retrieval, addressing vague natural language queries.
  - **Approach:** Fine-tuned a BERT encoder with a contrastive multi-task loss aligning user descriptions with movie metadata in a shared semantic space.
  - **Result:** Outperformed a GPT-4 few-shot baseline on Recall@K and MRR, creating a new synthetic dataset of 100K+ movie descriptions.

- **Resource-Efficient LLM Fine-tuning for Mental Health Support**
  - **Objective:** Adapted a large language model for mental health conversation support under limited computational and data resources.
  - **Approach:** Applied parameter-efficient fine-tuning (PEFT) with Quantized Low-Rank Adaptation (QLoRA) to the Falcon 7B model, improving task specialization efficiency.

- **GradCompass: AI-Powered Graduate Application Assistant**
  - **Objective:** Built a multi-agent LLM system to deliver personalized graduate school guidance, document evaluation, and recommendations.
  - **Approach:** Implemented an **8-agent framework** with role-specific LLMs for university recommendation and advising, integrated with a RAG pipeline and an AI-based interview simulator.

- **E2E Text-to-Comic Generation Pipeline**
  - **Objective:** Created an end-to-end web system that transforms textual narratives into comic strips.
  - **Approach:** Integrated GPT-3.5 and Stable Diffusion via a modular pipeline orchestrated on Google Kubernetes Engine (GKE) with Redis caching for high-throughput, concurrent generation.

Industry Experience

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

- Architected and implemented **TestNG and BATS** automation frameworks, focusing on scalable and robust end-to-end testing for Oracle Rest Data Services and SQLcL , reducing manual effort by 95%.
- Engineered a Python-based migration framework to automate the transfer of Oracle APEX applications, improving deployment efficiency and supporting **over 200 employees**.
- **Led a 15-member team** to validate a new DB Tools feature on Oracle Cloud in 2 months.

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

- Built an intelligent automation solution with Groovy and Workfusion OCR, optimizing data workflows and integrating text extraction into software robots, **increasing speed by 40% and reducing errors by 80%.**
- Delivered a Proof of Concept (PoC) on Intelligent Business Process Management Software tools, showcasing automation potential to stakeholders.

Research Interests

Commonsense Reasoning • Neuro-symbolic NLP • Responsible Machine Learning • Logical Inference • Fairness and Interpretability in LLMs • LLM Systems and Scalability • Multi-Agent AI Architectures