

The Hawk AI Research Project: Developing Trusted LLMs for Real World Solutions

Research proposal

Principal Investigators: Christopher Shane Elliott, Douglas Engelman, Sang Teck Oh

Abstract

The rapid adoption of large language models (LLMs) for therapeutic and wellness purposes has raised significant ethical and practical concerns, particularly among vulnerable populations such as college students. What is the University's role in creating and curating knowledge that LLM applications use? Could LLM technology in universities become powerful tools enhancing student wellbeing and success? The Hawk AI Research Project develops both a prototype application and an institutional process for further debating these questions. This project creates an LLM's knowledge corpus, tests its utility, and debates the experience of using the application on campus and community forums. A research community will be invited to "finetune" this LLM, while meeting in focus groups to discuss its feasibility and effects. Our prototype will either prove successful or suggest the pitfalls such technologies will face.

Purpose and Primary Objective

In 2022, large language models (LLMs) were released for public consumption with many unintended consequences. Among these have been the adoption of LLMs for therapeutic or wellness purposes. While public universities and its scholars have for over a decade sought to responsibly, carefully develop LLMs for mental health applications, specific to the needs of either clinicians or patients. Meanwhile, college students are lucrative populations for commercial development. Since they are comfortable interacting with social media applications, and they are suffering from increasing wellness challenges, this combination could make students attractive for market development—where ethical oversight in developing these technologies lag.

One problem is that training data is protected by proprietary privacy laws, making critical scrutiny difficult. Therefore, biases shaping users' understanding—so called "garbage in"—become challenging to rectify. Secondly, the LLM's context for meaning when it interprets queries may not be clear, which leads to unclear or hallucinatory information that further strains vulnerable college student. Careful, conscientious scrutiny over the LLM's knowledge corpus is needed, and this process creates an opportunity. The institutional process in constructing this LLM could create campus wide connections in the production of wellness and success information databases. The "Hawk AI" Research Project develops a training protocol for LLMs that leverages its strengths, while addressing these ethical concerns. The objective of this project is to train and publish a Microsoft Copilot Studio Application that could assist students in their collegiate challenges.

Scope of work

This project proceeds in three phases. In the first, Drs. Shane Elliott and Douglas Engelman present the Hawk AI Research Project in student focus groups throughout the Fall semester. One has been conducted, with four more planned. These sessions present the LLM building strategy to students to discuss how such a device might be used. We also present research with experts in health care, social work and human computer interaction with the knowledge corpus, and research describing how trusted LLM interactions—as tutors who do not replace expertise—can be designed. A publicly accessible website is made available where corpus materials, LLM parameters, and LLM interactions are published. These presentations should form the cornerstone of practice for such LLM knowledge corpus construction projects: public engagement and discussion of the LLM construction process.

Secondly, in the spring 2025 semester, Drs. Elliott, Engelman, Sang Teck Oh and four UNCW student research assistants finetune the knowledge corpus with Hawk AI training. Additional faculty contributors will be recruited from previous presentations to assist with finetuning the LLM. In 2 sessions per month, persons who had interacted with the LLM over the previous period will bring transcripts for review and discussion. Eight sessions of content analysis during the 2025 Spring semester will be made open to public attendance, and advertised among campus experts who are stakeholders in student wellness. Collectively, we determine the meaning of LLM responses. These sessions comprise collective discussion on the feasibility for developing and deploying LLMs for student success in collegiate environments. The original research team will complete finetuning over the summer of 2025.

In our third step, we seek additional funding for developing the Hawk AI Research Project team for the 2025-2026 year. Here, we expect to publish the Hawk AI application for semester long student usage. Research funding to pay for student interactions with the neural network, support staff, funds for recruiting users into the study and monitoring the LLM's deployment will be needed. We expect this research effort to be an ongoing collaboration, in conjunction with partners developed through the first two phases. We have identified NC innovations as potential funding sources for this innovative technology. Furthermore, the Hawk AI Research Project could be developed into an institutional process for maintaining a trusted university LLM committed to student success and wellness. Such institutional processes could serve as a model for campuses across the country, as every other campus is currently dealing with these challenges. We will all be learning from the attempts made.

Timeline and Milestones

Phase One: Conceptual Development

1. Five Student Focus Group Observations
 - a. Emergent content analysis of trends in LLM technology expectations and possibilities.
 - b. Data gathered and analyzed by December
2. Five Presentations for Campus Experts
 - a. Discussions of the Hawk AI concept with campus stakeholders.
 - b. UNCW faculty discussions gathering feedback on the concept.
 - c. Build Hawk AI relationships with campus stakeholders.

- d. Once per month for five months.
- 3. Website Postings for Public Debate
 - a. Publish project concept to website for ongoing public and campus engagement.
 - b. Knowledge corpus references and materials.
 - c. Recruiting faculty and students to use Hawk AI during Spring trials.
 - d. Published by January 2025.

Phase Two: Build Knowledge Corpus and Finetune Hawk AI Interactions

- 1. Complete Knowledge Corpus for Student Wellness
 - a. Publicly Available January 2025
- 2. Interact with LLM to assess finetuning needs.
 - a. Researchers produce transcripts of LLM interactions.
 - b. Users reflect on experiences of LLM interactions.
 - c. Meet bi-monthly for eight sessions during the Spring 2025 semester to share reflections.
 - d. Discuss how to code the interactions.
 - e. Discuss how to finetune the LLM.
 - f. Sessions open to campus attendance.
- 3. Present findings on users' experiences and finetuning needs
 - a. Campus research event
 - e. Southern Sociology Conference
 - i. April 9-12, 2025 for three presenters
 - f. Recruit undergraduate students to use Hawk AI during the Fall 2025 semester.
- 4. Complete Hawk AI training
 - a. May – July 2025

Phase Three: Publish Hawk AI for Fall 2025 Student Wellness usage

- 1. Secure funding for supporting faculty and student for Hawk AI's 2025-2026 launch.
- 2. Recruit undergraduate students to use Hawk AI during the Fall 2025 semester.

End Products

- 1. Hawk AI Large Language Model Prototype
- 2. Hawk AI Student Wellness Open-Sourced Knowledge Corpus
- 3. Poster for student presentation (campus research events)
- 4. Paper describing the collaborative, institutional process of LLM construction