UsherGPT: Supporting students with a generative AI model specialising in public health and medical texts

Katya Lait[1], Kenneth Li[2], Marshall Dozier[3], Andrew Horne[4], Neneh Rowa-Dewar[1],
Sian Bayne[5], Mattia Opper[6], Mike Sharples[7], Kevin Tsang[1]

[1]Usher Institute, University of Edinburgh; [2]DialPad, Canada; [3]Library Academic Support, University of Edinburgh;
[4]EDINA, University of Edinburgh; [5]Moray House School of Education and Sport, University of Edinburgh;
[6]School of Informatics, University of Edinburgh; [7]Institute of Educational Technology, Open University

We are developing a specialised generative AI model that would be fine-tuned on public health and medical data. It will be made available to students at the Usher Institute via a chat-style interface.

"I'm a non-native English speaker, explain to me what relative risk is in epidemiology."

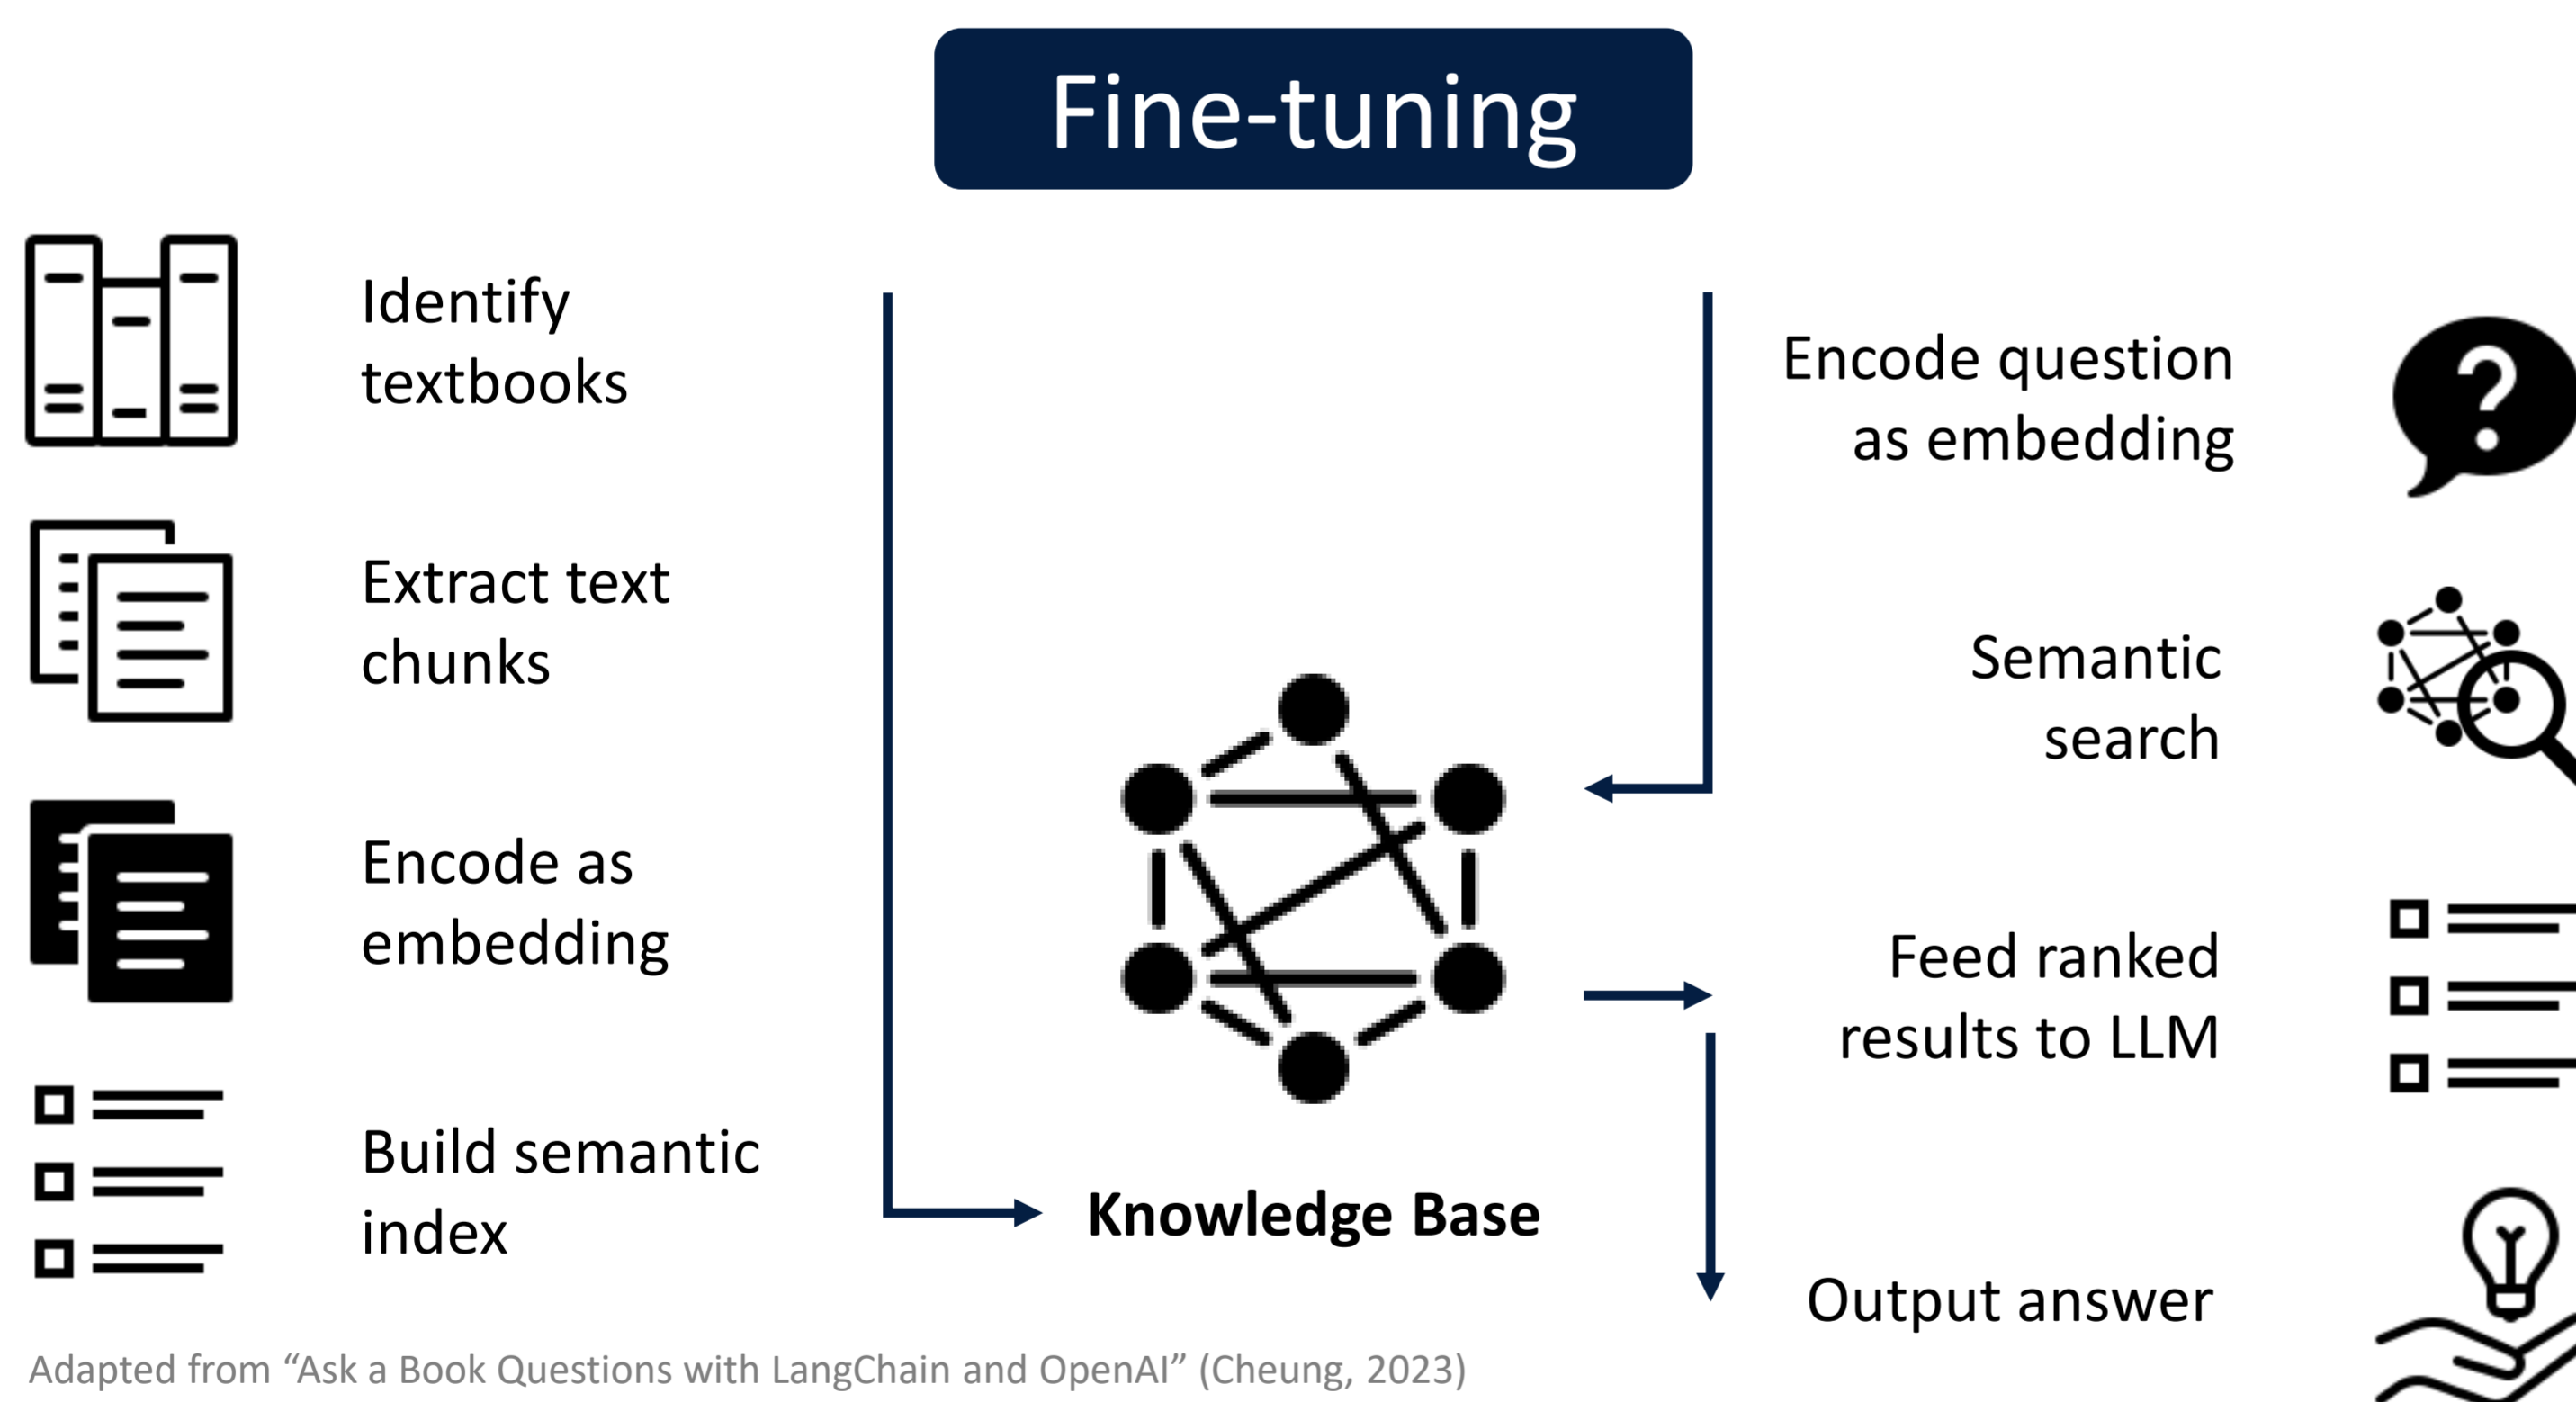*"What are the major determinants of health in a population?"*

Today, most implementations of generative AI are generalist in nature and not useful in specialist contexts. When fine-tuning has been applied, models can provide outputs that are more accurate for particular use cases. Generative AI can also become a tool to enable paraphrasing for different audiences, like non-native English speaking and neurodivergent students or students struggling with foundational concepts. What if these groups could specify exactly how they wanted information about a certain subject to be explained to them – as if they had a personal tutor?

Objective 1
Develop guidance on safe use of generative AI in higher education and framework for preparing fine-tuning data

Objective 2
Build an interactive generative AI platform fine-tuned on public health and medical data for staff and students

Fine-tuning

- Identify textbooks
- Extract text chunks
- Encode as embedding
- Build semantic index

Knowledge Base

- Encode question as embedding
- Semantic search
- Feed ranked results to LLM
- Output answer

Adapted from "Ask a Book Questions with LangChain and OpenAI" (Cheung, 2023)

Benchmarking
Multiple large language models (LLMs) will be benchmarked by evaluating the generated answers for a set list of questions about public health and medical topics.

Next steps
- Beta testing and focus groups with users [~ winter 2023]
- Launch of model and release of guidance about safe use of tool and framework to prepare fine-tuning data [~ end of 2023]

THE UNIVERSITY of EDINBURGH | usher institute

kcht

IAD INSTITUTE FOR ACADEMIC DEVELOPMENT · dialpad · The Open University