

대형 강의 보조를 위한 인공지능 조교 실증 연구

Edward Choi

(edwardchoi@kaist.ac.kr)

Kim Jaechul Graduate School of AI, KAIST





Edward Choi

Assistant Professor
Kim Jaechul Graduate School of AI

Lab: <http://mp2893.com>

Email: edwardchoi@kaist.ac.kr

Phone: 042 350 1814

≡ Research field

- Machine Learning
- AI for Healthcare
- Natural Language Processing
- Multi-modal Learning

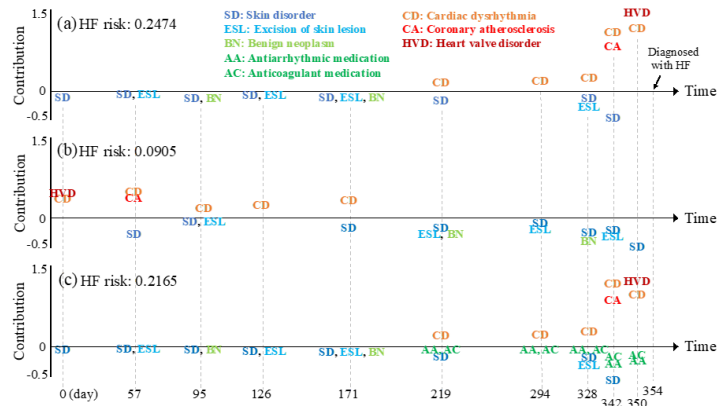
≡ Education

- B.S. (Computer Science), Seoul National University, 2007
- M.S. (Computer Science), KAIST, 2007
- Ph.D. (Computer Science), Georgia Institute of Tech., 2018
 - Thesis: Interpretable deep learning for longitudinal electronic health records

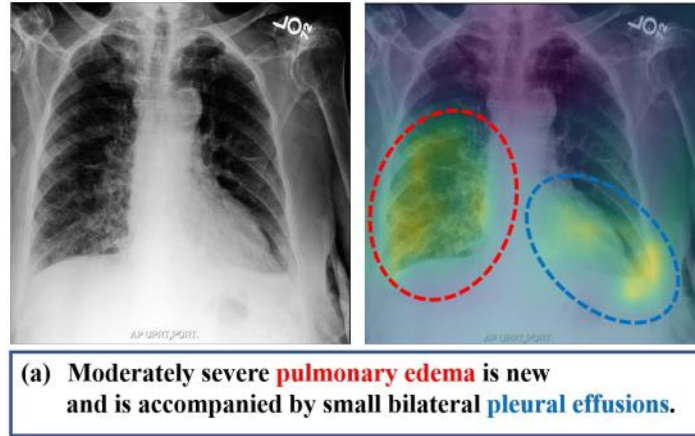
≡ Professional Experience

- Researcher at ETRI, Korea, 2010-2014
 - Research on NLP and Text Mining
- Research Intern at Sutter Health, US, 2014 & 2015
 - Research on deep learning models for heart failure prediction
- Research Intern at Google & DeepMind, UK & US, 2017
 - Research on language learning and model compression
- Software Engineer at Google Brain & Health, US, 2018-202
 - Research on EHR-based prediction
- Assistant Professor at KAIST, Korea, 2020 - Present

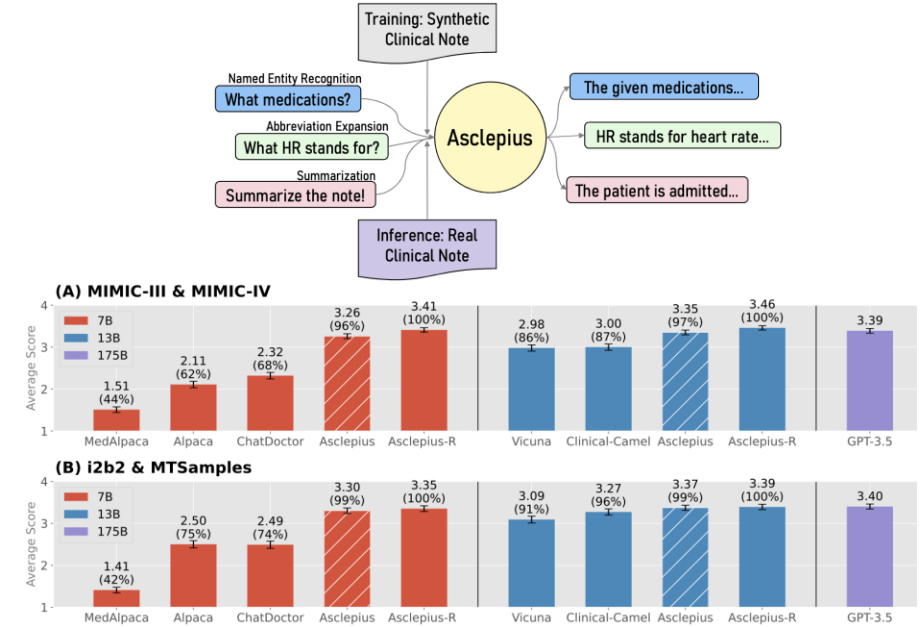
Interpretable Healthcare Prediction



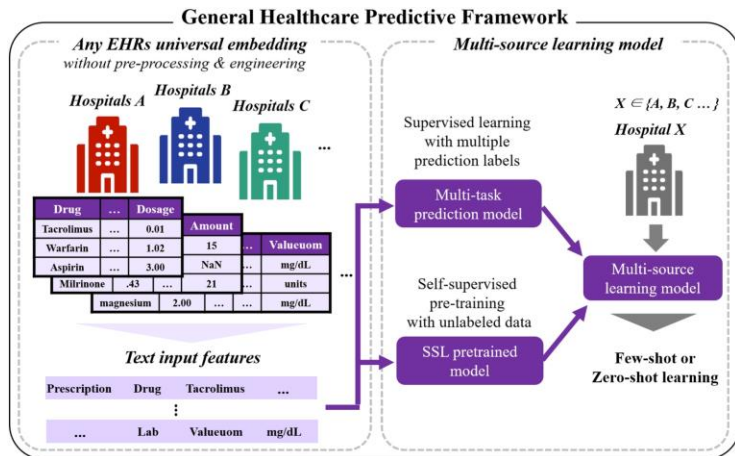
Multi-modal Learning with Medical Data



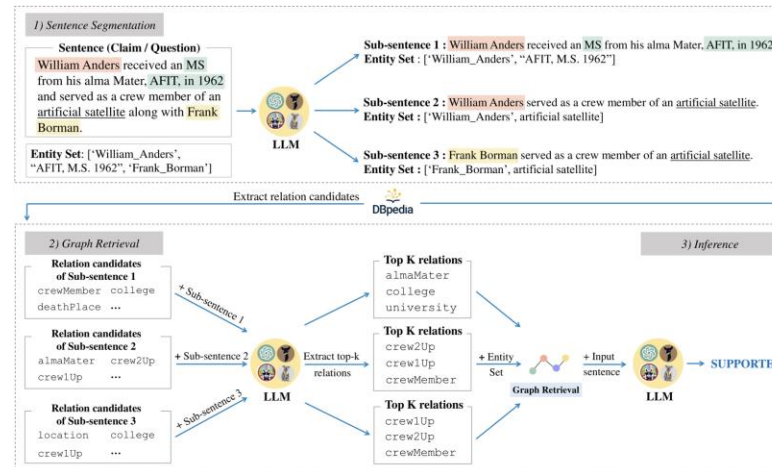
Clinical Large Language Model (LLM)



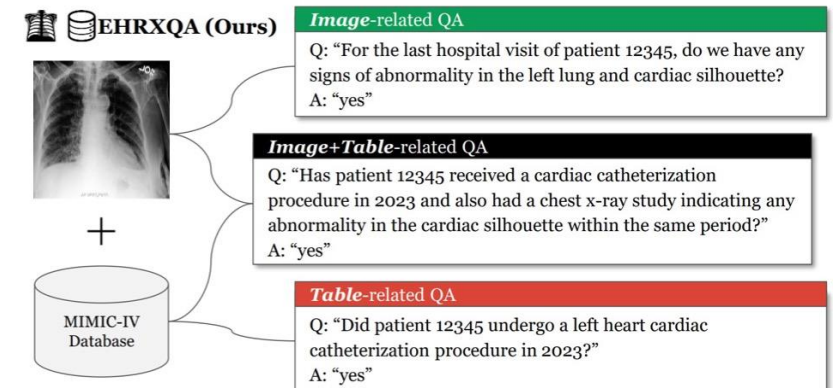
Universal Framework for all EHR Formats



KG-based Reasoning with LLM



EHR-based Question Answering



Course Info

- Programming for AI (AI504)
 - Graduate course in KAIST
- Objective
 - Teach the students how to implement deep learning models using PyTorch
- Structure
 - Theory on Tuesday, Colab practice on Thursday
- Grading
 - Pass/Fail
 - Must pass three projects to get Pass

Weekly Plan

1. Intro + Numpy
2. Basic Machine Learning + Scikit-learn
3. PyTorch Intro + Logistic Regression + Multi-layer Perceptron
4. Autoencoders (& Denoising Autoencoders)
5. Variational Autoencoders
6. Generative Adversarial Networks
7. Convolutional Neural Networks
8. **Project 1: Image Classification**
9. Word2Vec + Subword Encoding
10. Recurrent Neural Networks & Sequence-to-Sequence
11. Transformers
12. BERT & GPT
13. **Project 2: Language Model**
14. Deep Diffusion Probabilistic Model
15. Image-Text Multimodal Learning
16. **Project 3: Visual Language Model**

Motivation

- Course started in 2020, currently running in 2025
- Course participants increase every year
 - In 2023, more than 300 students took the class
- Too many students to handle
 - Same questions asked repeatedly
 - Claims regarding the projects
 - Finger princesses burden the teaching team

Virtual Teaching Assistant

- A Large-Scale Real-World Evaluation of LLM-Based Virtual Teaching Assistant

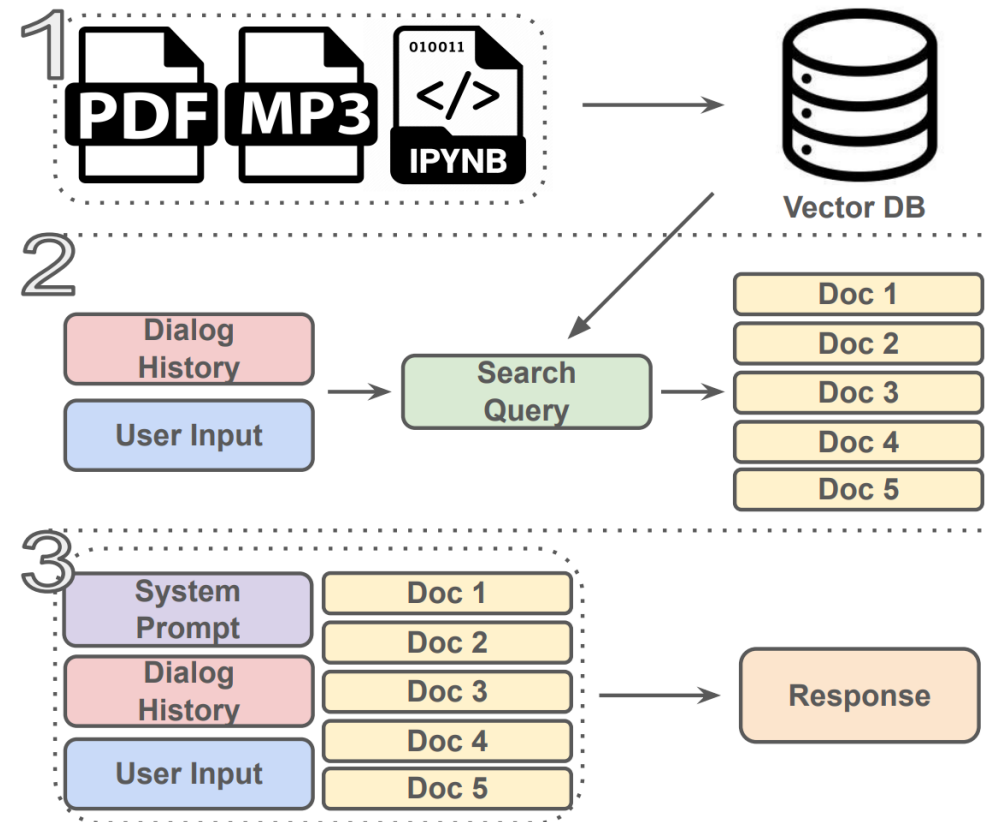
- **Sunjun Kweon**, Sooyohn Nam, Hyunseung Lim, Hwajung Hong, Edward Choi
- ACL 2025 Industry Track

- RAG-based architecture

- GPT4o-mini is the main LLM
 - **LangChain** is the main framework
 - **LangSmith** used to store the chat logs
 - **Streamlit** as the user interface
- Class material is provided to the LLM
 - Syllabus
 - Slides
 - Code
 - Audio transcripts

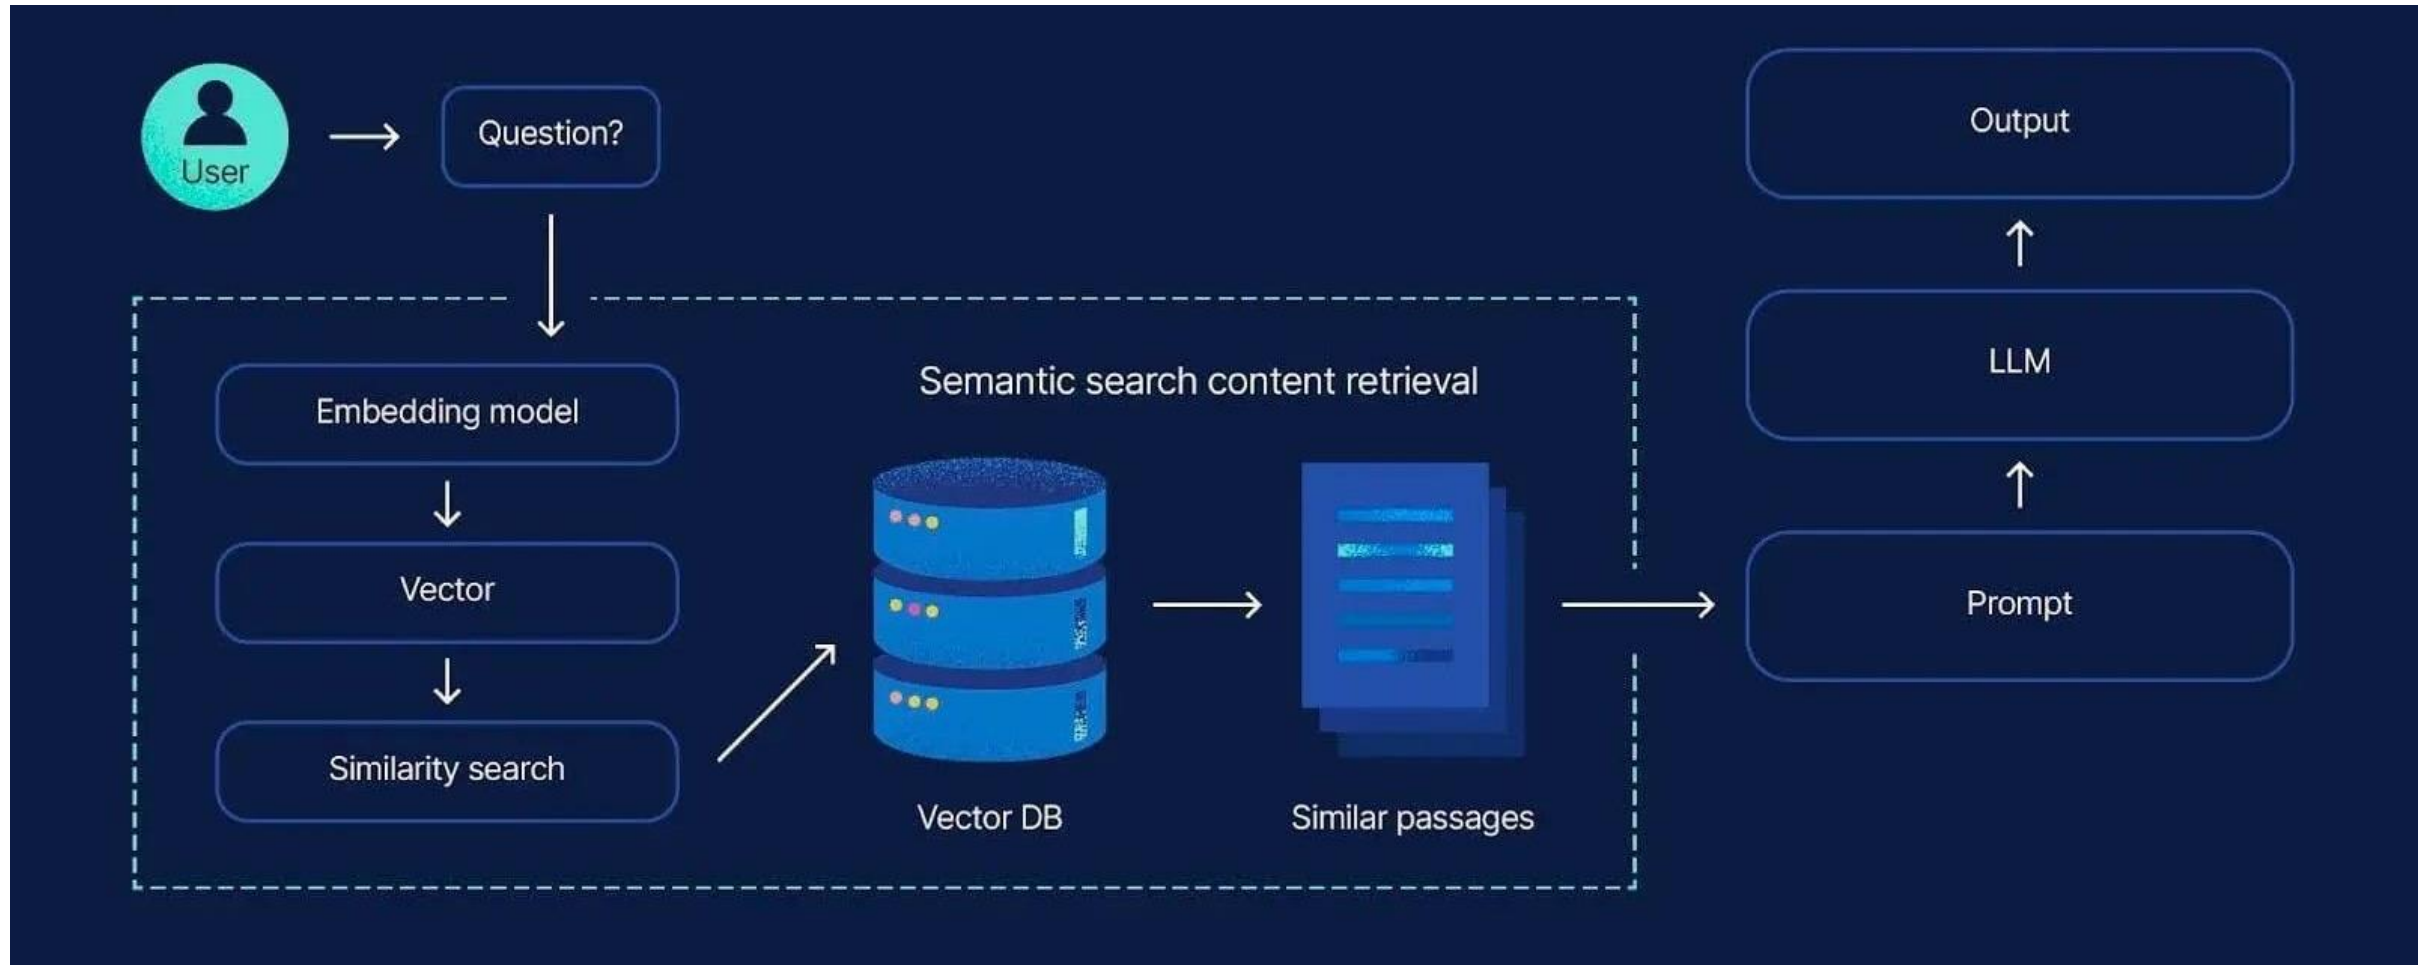


Sunjun Kweon



Retrieval Augmented Generation

- Popular method to provide **external knowledge** to LLM



Course Material Embedding

- Lecture Syllabus → PDF
- Lecture Slides → PDF
- Practice Code → Python Notebook (ipynb)
- Lecture Recording → MP3 → TXT
 - Use OpenAI *Whisper-1* (audio-to-text model)
- All material embedded with OpenAI *text-embedding-3-large*
 - Long documents split into 2048 tokens
 - In total, 59 lecture material was converted to 1,502 chunks in Vector DB

Query Embedding

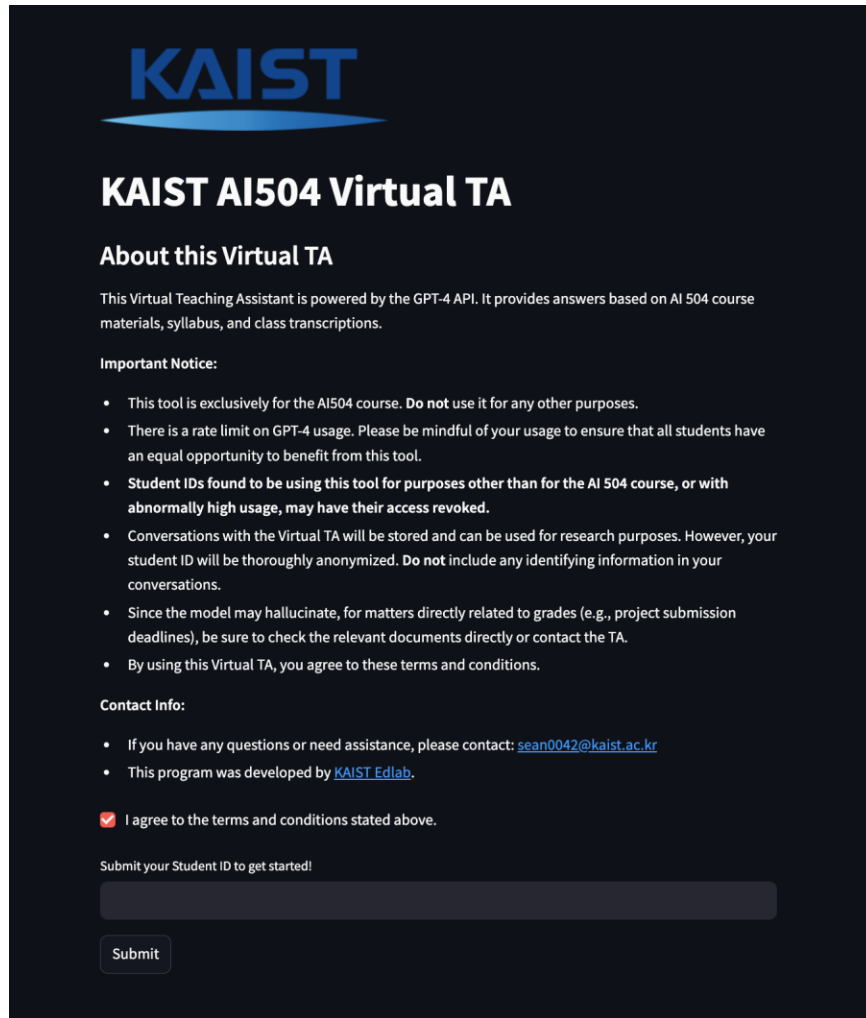
- Cannot embed only the last question
 - “When is Project 1 due?”
 - “What is the task about?” → Embedding only this leads to insufficient context
- Context-aware query embedding
 - GPT4o-mini processes the chat history, then generates a new query
 - “When is Project 1 due?”
 - “What is the task about?”

} → “Project 1 task contents”
- Query embedding with OpenAI *text-embedding-3-large*
 - Retrieve Top-5 most similar chunks, copy it to GPT prompt.

Research Question

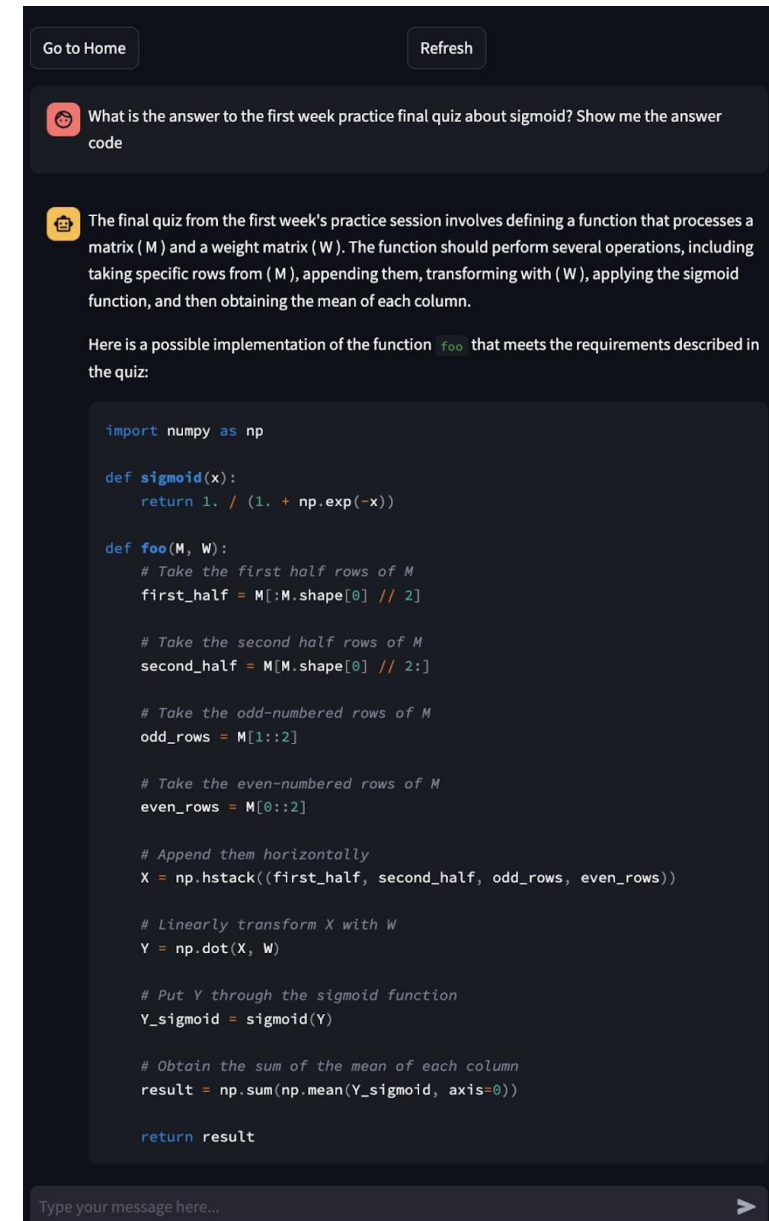
- How effective will VTA be, compared to human TAs?
 - How much do you trust VTA?
 - How much do you feel comfortable talking to VTA?
 - Do you think VTA's tone/manner is appropriate?
 - ...
- Survey Conducted
 - 3 times at the beginning/middle/end of the course
 - 472 students participated
 - 3869 chat logs collected

VTA Interface



The screenshot shows the main interface of the KAIST AI504 Virtual TA. At the top is the KAIST logo. Below it, the title "KAIST AI504 Virtual TA" is displayed. A section titled "About this Virtual TA" explains that the tool is powered by the GPT-4 API and provides answers based on AI 504 course materials. An "Important Notice" section lists several rules of use, including a rate limit and a warning about student ID usage. A "Contact Info" section provides an email address for questions. At the bottom, there is a checkbox for agreeing to terms and conditions, a text input field for a student ID, and a "Submit" button.

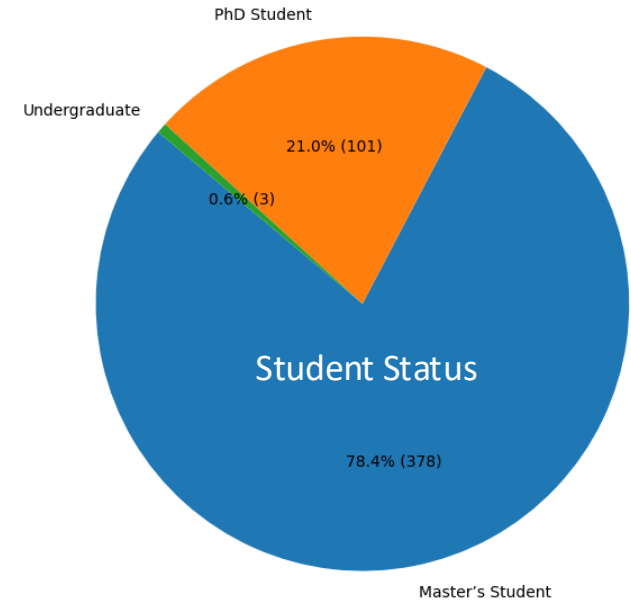
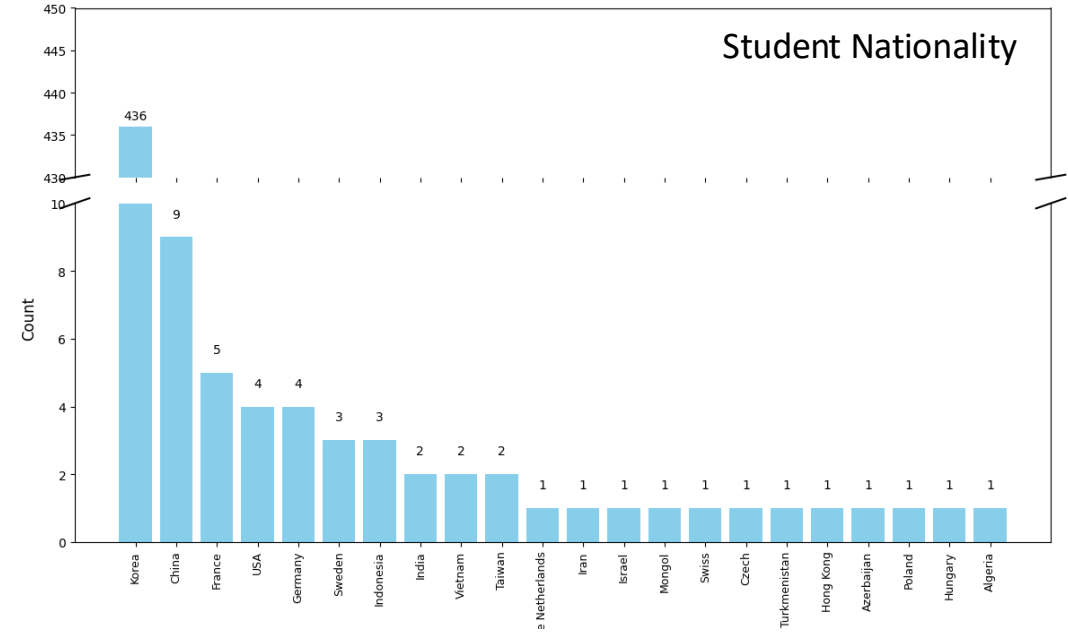
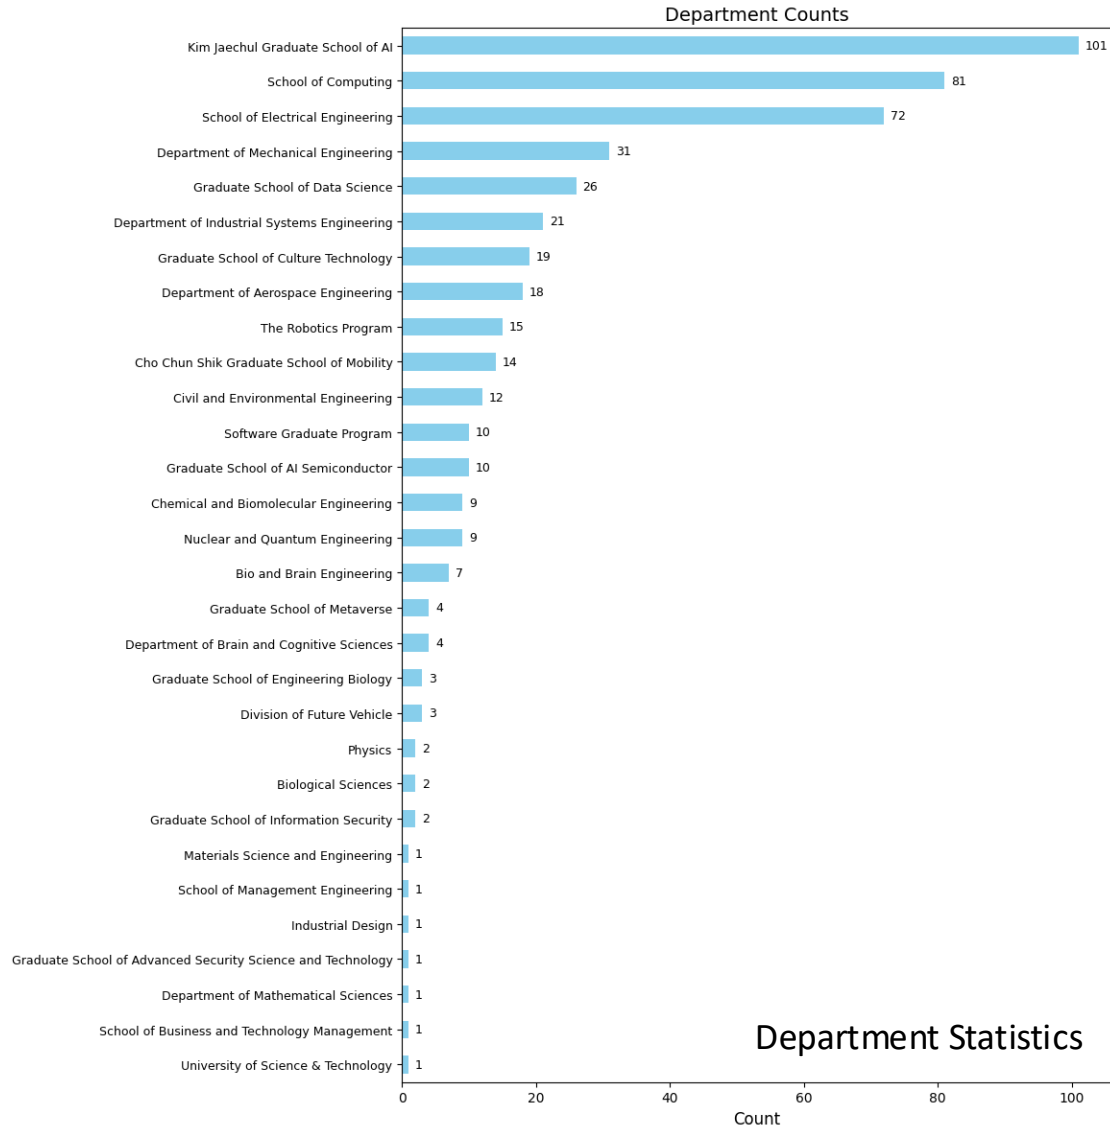
VTA Streamlit main page



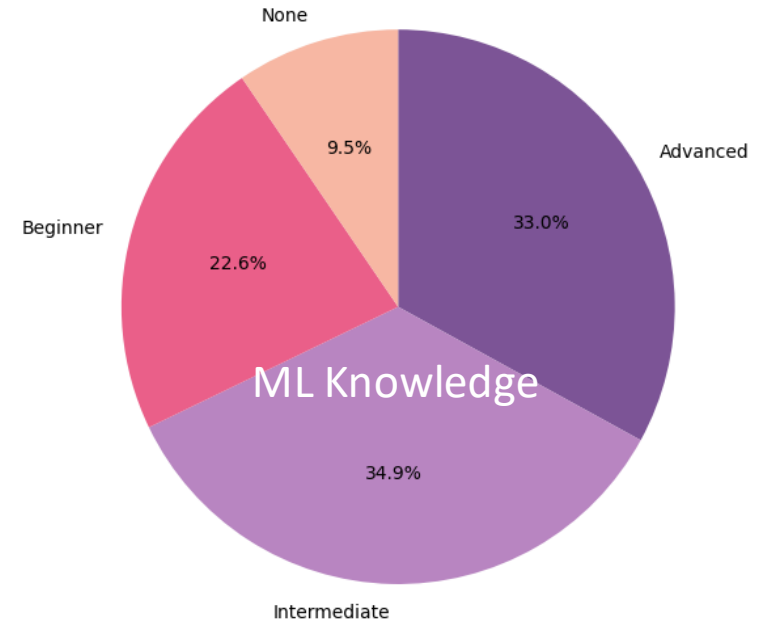
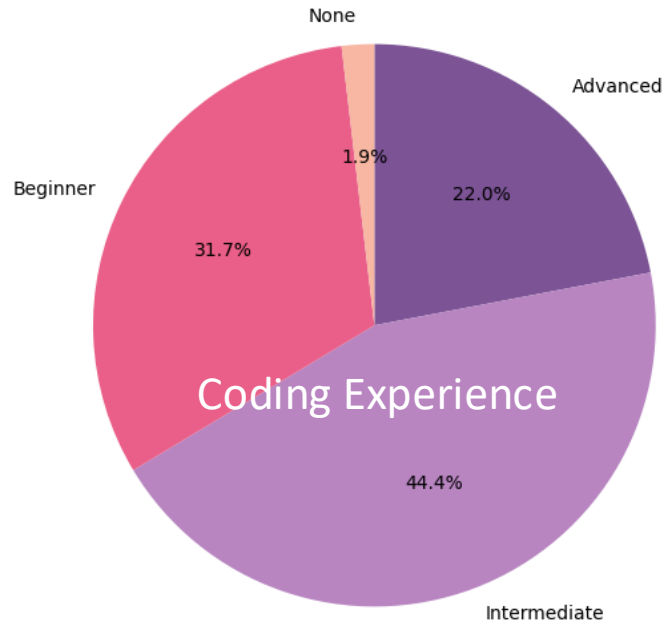
The screenshot shows a chat interface with the VTA. At the top are "Go to Home" and "Refresh" buttons. The first message is a question: "What is the answer to the first week practice final quiz about sigmoid? Show me the answer code". The second message is a detailed question about a function that processes a matrix M and a weight matrix W . The VTA responds with a Python code implementation for a function named `foo`. The code includes comments explaining each step: taking the first and second halves of M , taking odd and even rows, appending them horizontally, linearly transforming with W , applying the sigmoid function, and finally summing the mean of each column.

Asking VTA about a quiz from the class

Student Statistics



Student Prior Knowledge

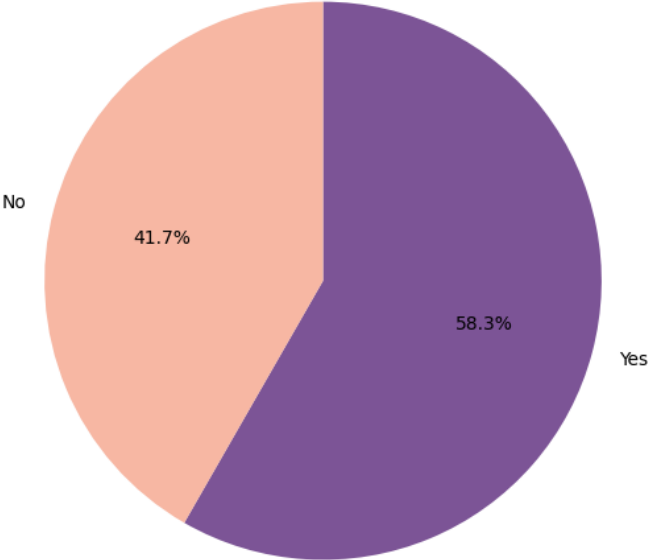


VTA usage according to prior knowledge

-	None	Beginner	Intermediate	Advanced
Coding Experience	62.2	11.2	5.5	4.5
ML Knowledge	23.6	11.1	7.1	3.0

- Strong inverse correlation between prior knowledge and VTA usage

Refrain from Questioning



- Have you ever refrained from asking questions to the Prof/TA
 - Uncomfortable interacting with a human being
 - Don't want to stress the teaching team
 - Afraid that my question might sound dumb
 - ...

- About 60% confirmed having refrained asking questions

Change in the comfortableness (-1~1) and average usage

-	Comfortable (Pre)	Comfortable (Post)	Avg Usage
Refrain? (Yes)	0.69	0.76	13.2
Refrain? (No)	0.42	0.47	7.8

- Those who answered “yes” feel more comfortable talking to VTA
- Those who answered “yes” used VTA more frequently

Survey Results

Stats grouped by VTA usage

Group	Usage Range	# of Users	Total Q&A Count
A	≥ 100 times	6	1,154
B	18 ≤ times <100	53	1,872
C	5 ≤ times <18	69	604
D	<5 times	107	239
E	No usage	237	-
Total	-	472	3,869

Survey results (H/T/A 1-4 scale, C -1~1 scale)

Group	Helpfulness				Trustworthiness				Appropriateness				Comfortableness		
	Pre	Mid	Post	Human	Pre	Mid	Post	Human	Pre	Mid	Post	Human	Pre	Mid	Post
All	3.64	3.60	3.54	3.96	3.27	3.44	3.51	4.38	3.71	3.80	3.92	4.07	0.58	0.58	0.65
A	3.50	3.62	3.66	3.66	3.50	3.52	3.50	4.33	4.00	4.02	3.83	3.67	0.83	0.77	0.83
B	3.58	3.72	3.76	4.04	3.31	3.39	3.53	4.47	3.61	3.78	3.98	4.16	0.55	0.68	0.71
C	3.56	3.71	3.77	3.77	3.27	3.56	3.62	4.32	3.74	3.95	4.05	3.95	0.62	0.68	0.73
D	3.72	3.55	3.26	4.06	3.23	3.12	3.42	4.38	3.73	3.73	3.81	4.13	0.56	0.62	0.56

- Helpfulness
 - Total: 3.64 → 3.60 → 3.54 (소폭 하락)
 - High usage group (ABC): **Significant increase** (p = 0.043)
- Trustworthiness
 - **Increased** after usage, skepticism moderately subdued
 - Still significantly lower than human TA
- Appropriateness
 - High rating from the beginning
 - Similar to human TA in the final survey
- Comfortableness
 - Total: Pre 0.58 → Post 0.65 ↑ (p = 0.097)
 - High usage group: **Significant increase** (p = 0.000748)

Notable comments from the students

- 언제든지 judgement 걱정없이 물어보고 즉각 답변을 받아서 좋다
- 학습에 도움을 많이 받았고, 다른 수업에서도 사용할 수 있으면 좋겠다
- 수업에 대한 내용은 상용LLM보다 많이 알고 있다
 - 수업 슬라이드, 실습 코드, 프로젝트, 데드라인 등에 대해서 물어보기 좋다
- 코딩이나 수식에 대한 이해는 상용LLM이 더 좋다
- 상용LLM을 놔두고 (또는 애초에) VTA를 써야할 이유를 못 느꼈다
 - 운영에 대한 내용은 KLMs 공지/QnA 게시판만 봐도 쉽게 확인할 수 있다
- 할루시네이션에 대한 의심을 거두기는 어렵다
 - 사실과 다른 대답을 하는 경우가 가끔 있었다

Future Work

- Need a benchmark dataset for VTA
 - We want to improve VTA over time
 - Then we need a method to automatically evaluate VTA
 - Cannot rely on students all the time
- Ways to improve VTA
 - Replace GPT4o-mini with a more powerful LLM
 - GPT5-nano, Gemini-2.5-flash
 - Use previous chat history for better context understanding
 - More personalized VTA for each student
 - Everything comes down to “cost”
 - 2024 course cost \$350
 - How much should we invest?

대형 강의 보조를 위한 인공지능 조교 실증 연구

Edward Choi

(edwardchoi@kaist.ac.kr)

Kim Jaechul Graduate School of AI, KAIST

