
UNIVERSITY
OF LONDON
Goldsmiths
UNIVERSITY OF LONDON
london.ac.uk

UNIVERSITY OF LONDON

Goldsmiths
UNIVERSITY OF LONDON

Applications of AI in the BSc CS degree

Professor Matthew Yee-King
Programme Director
BSc Computer Science

Summary

- Some terms
- Local LLMs for survey analysis
- LLM embeddings for collusion detection
- WIP: Generating exam papers
- WIP: course chatbo

Introduction

- PhD AI/CS (creative signal processing) and CS educator
- BSc CS fully online undergraduate programme:
- UoL, Goldsmiths Coursera
- Launched 2019, ~4000 active students

coursera

Position statement

- Enhance not replace
- Alignment: how to constrain the current and future generation of AI

coursera

Goldsmiths
UNIVERSITY OF LONDON

UNIVERSITY OF LONDON

Value alignment is hard

"The goal of this workshop is to bring together researchers and thought leaders from various disciplines such as computer science, social sciences, philosophy, law etc. to discuss the ethical considerations and implications of aligning AI to human values."

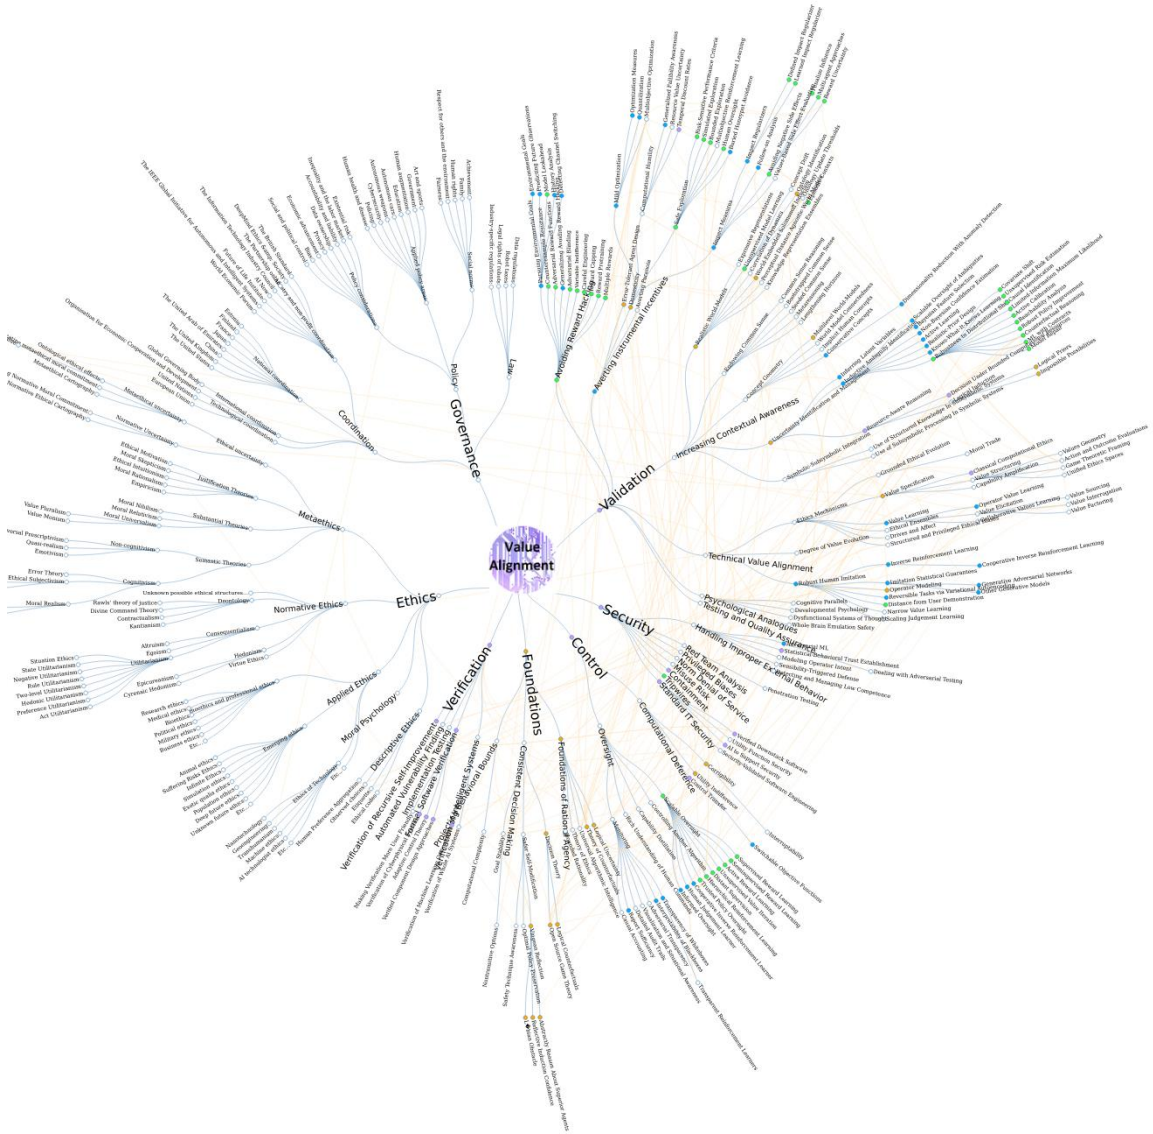

- https://futureoflife.org/valuealignmentmap/ (2017)

Seven Top level Categories

Governance

Ethics

Validation

Verification

Foundation

Control

Security

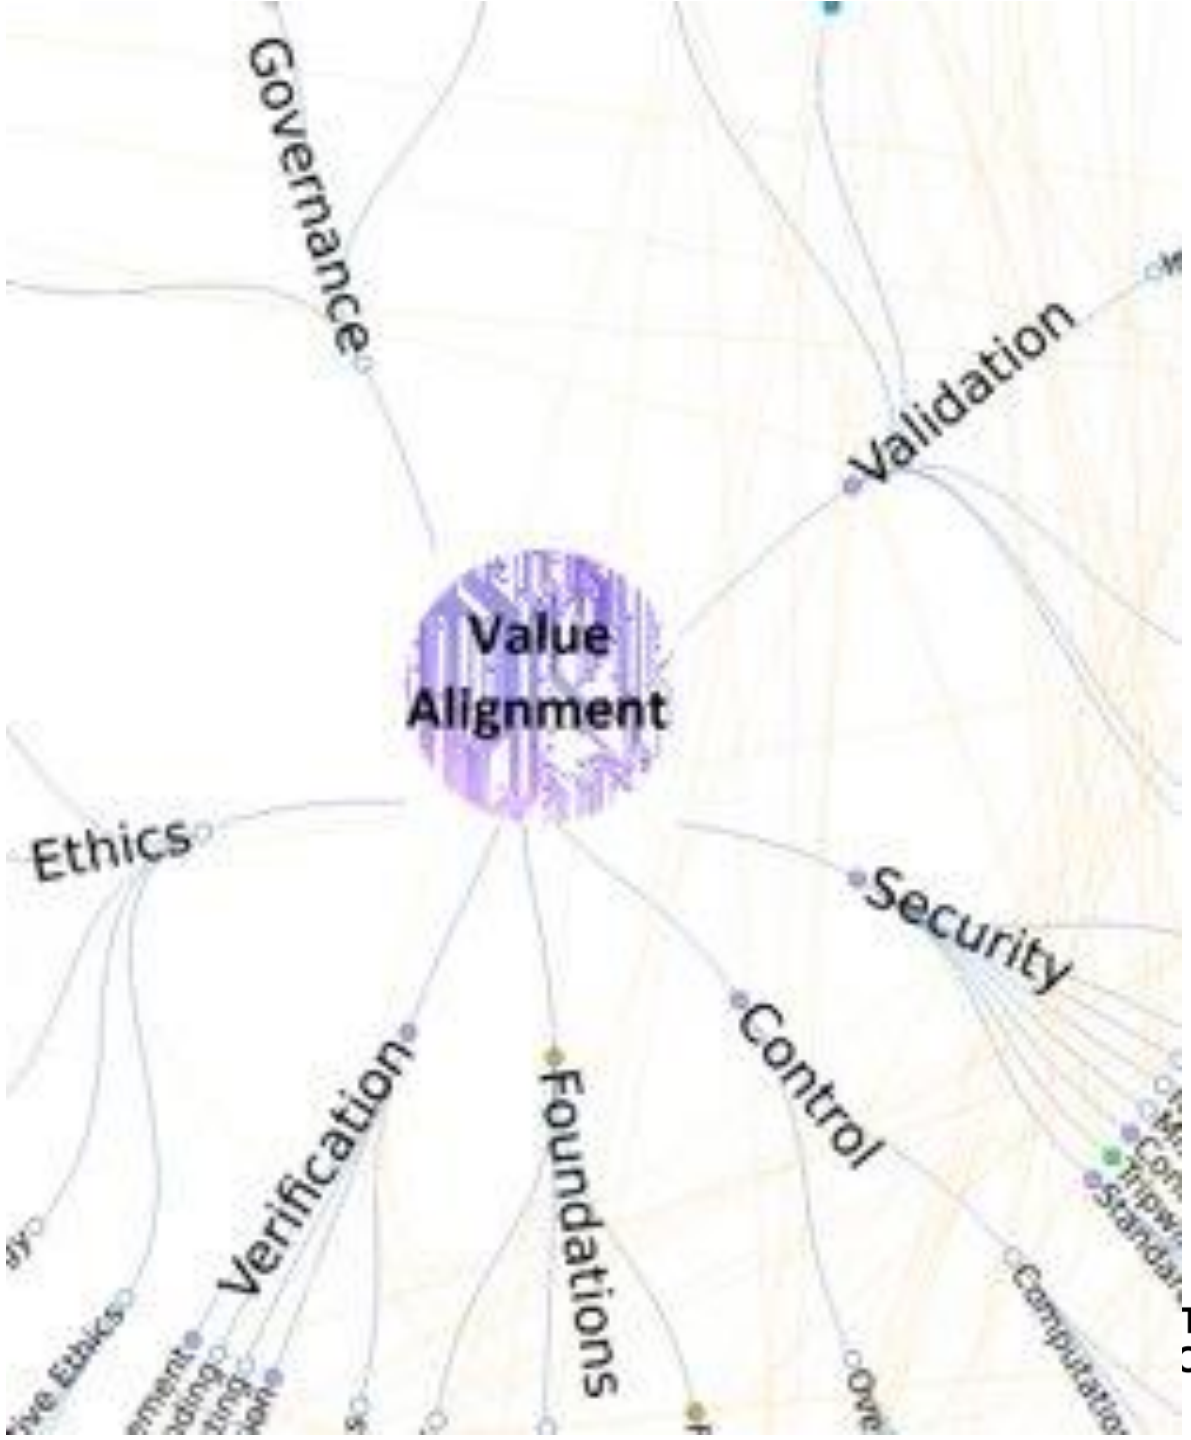

Exposure to AI in different jobs

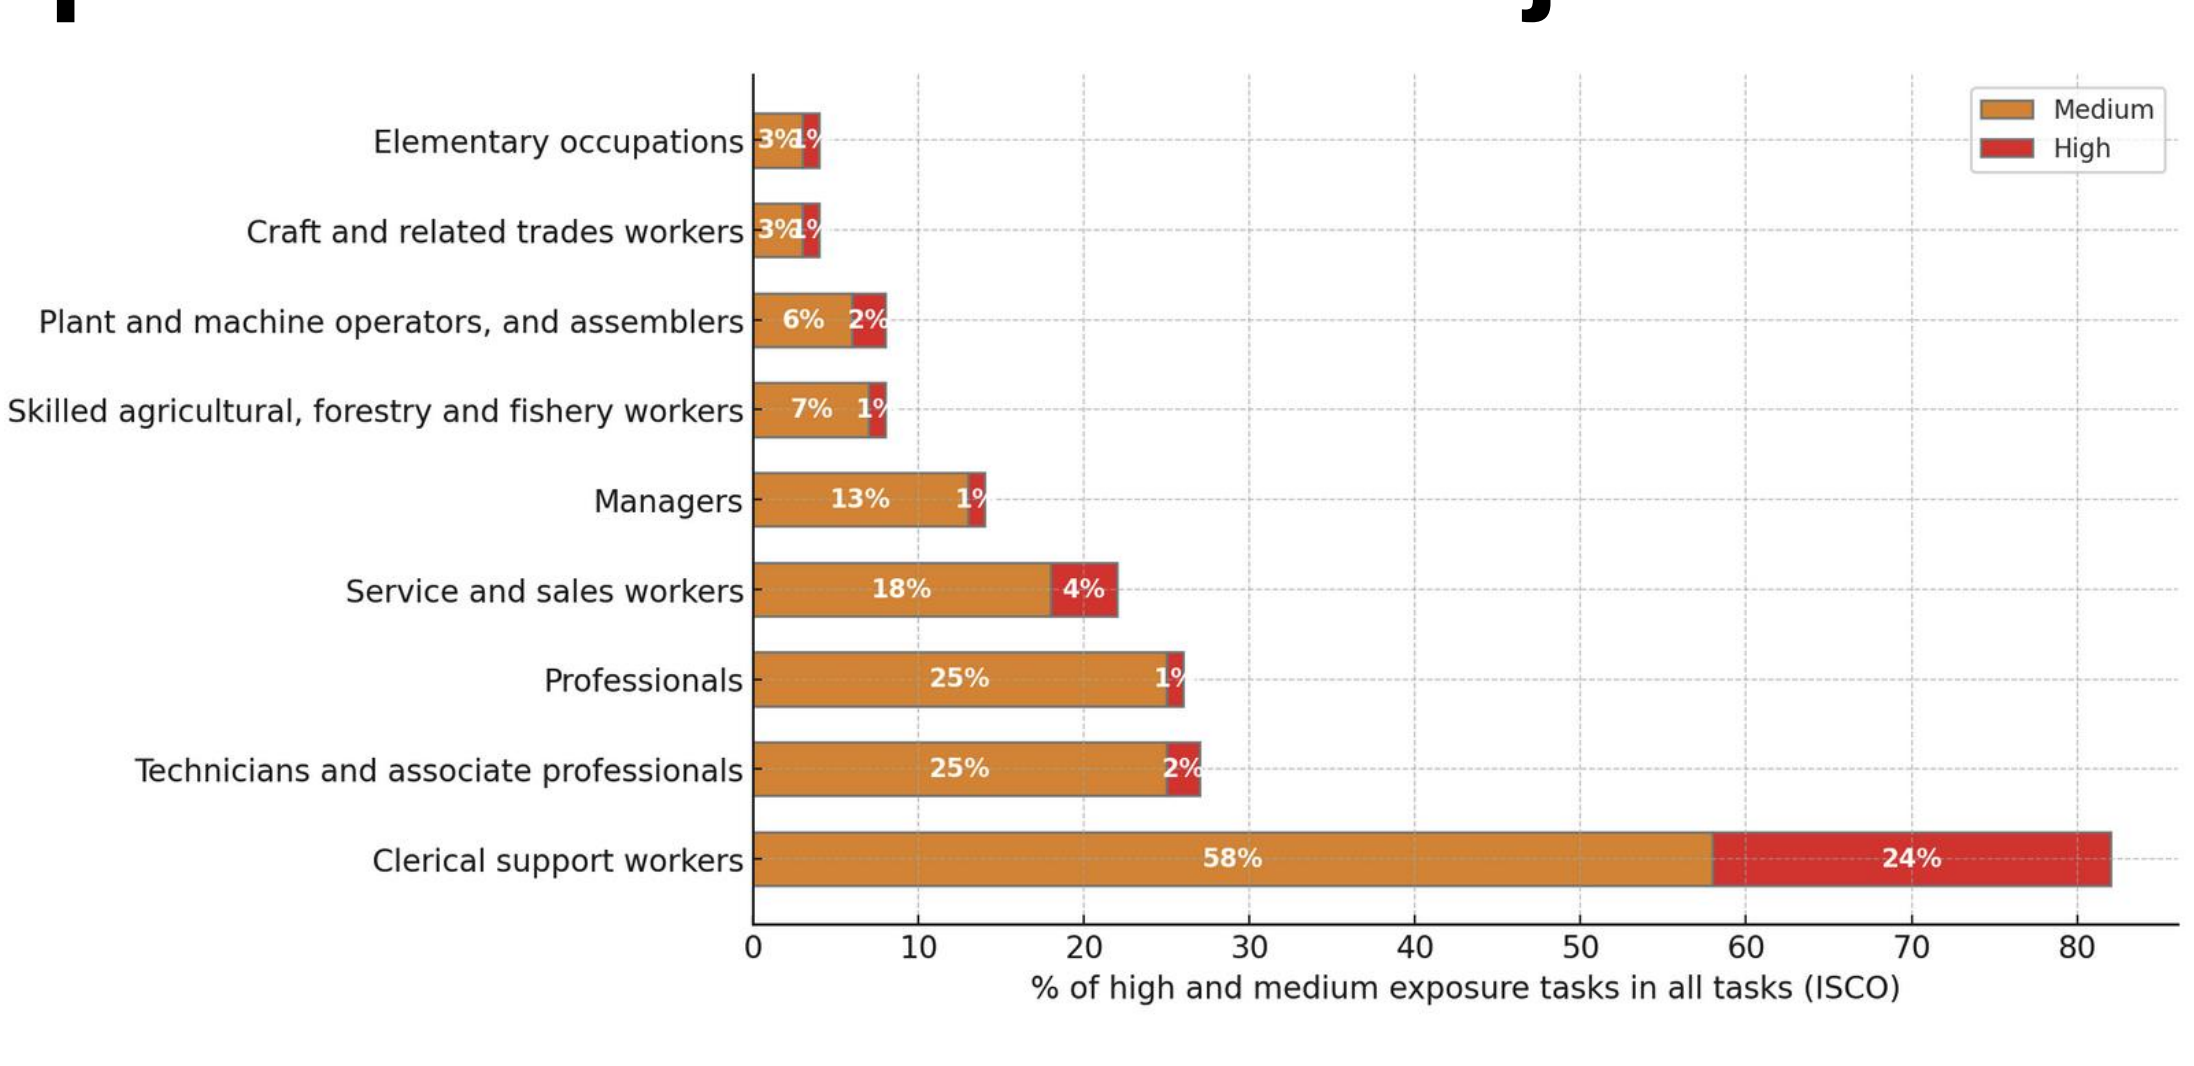

Tasks with medium and high GPT-exposure, by occupational category

Recreated after: https://www.econstor.eu/bitstream/10419/278614/1/1857683005.pdf (cc license)

Possible interactive thing

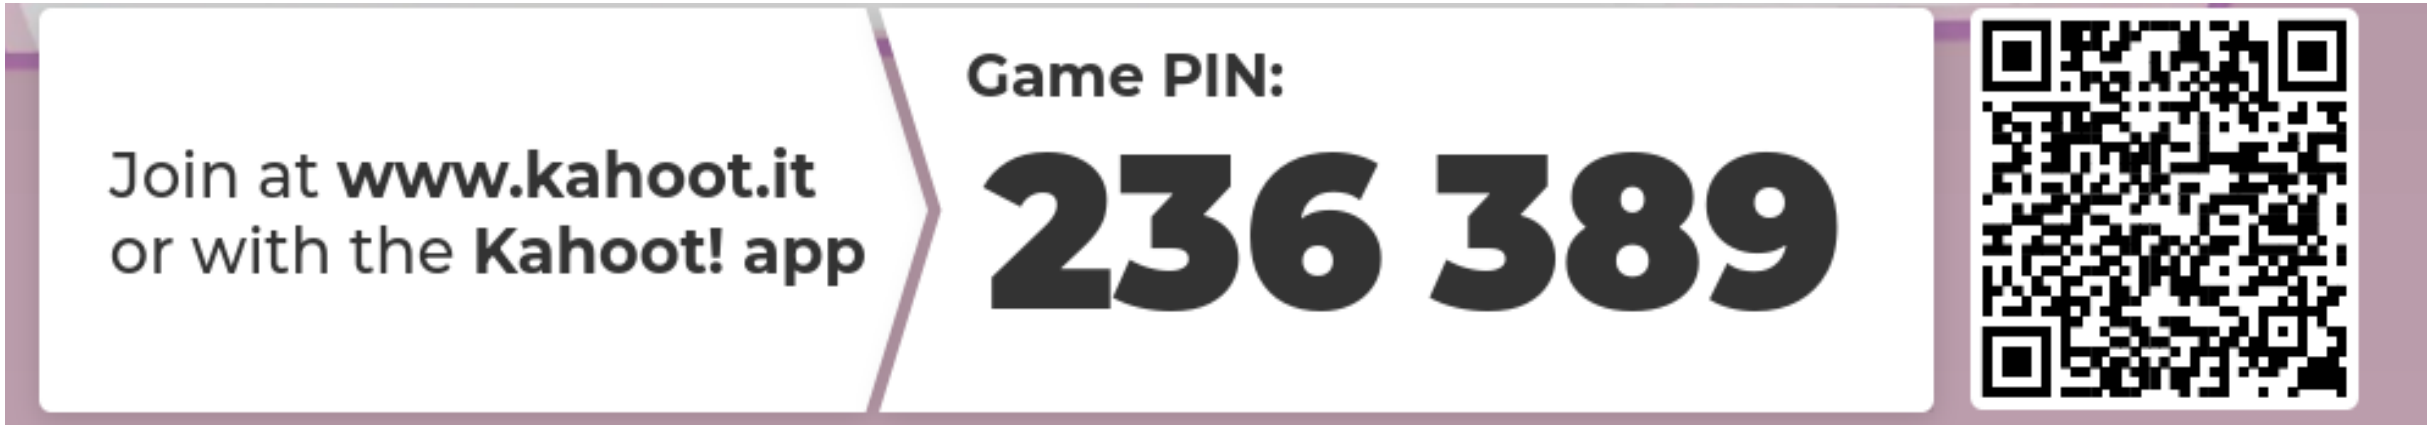

coursera

Goldsmiths
UNIVERSITY OF LONDON

UNIVERSITY OF LONDON

Ethical HE frameworks

- Lots of people are developing HE AI frameworks

Case study 1: analysing survey data

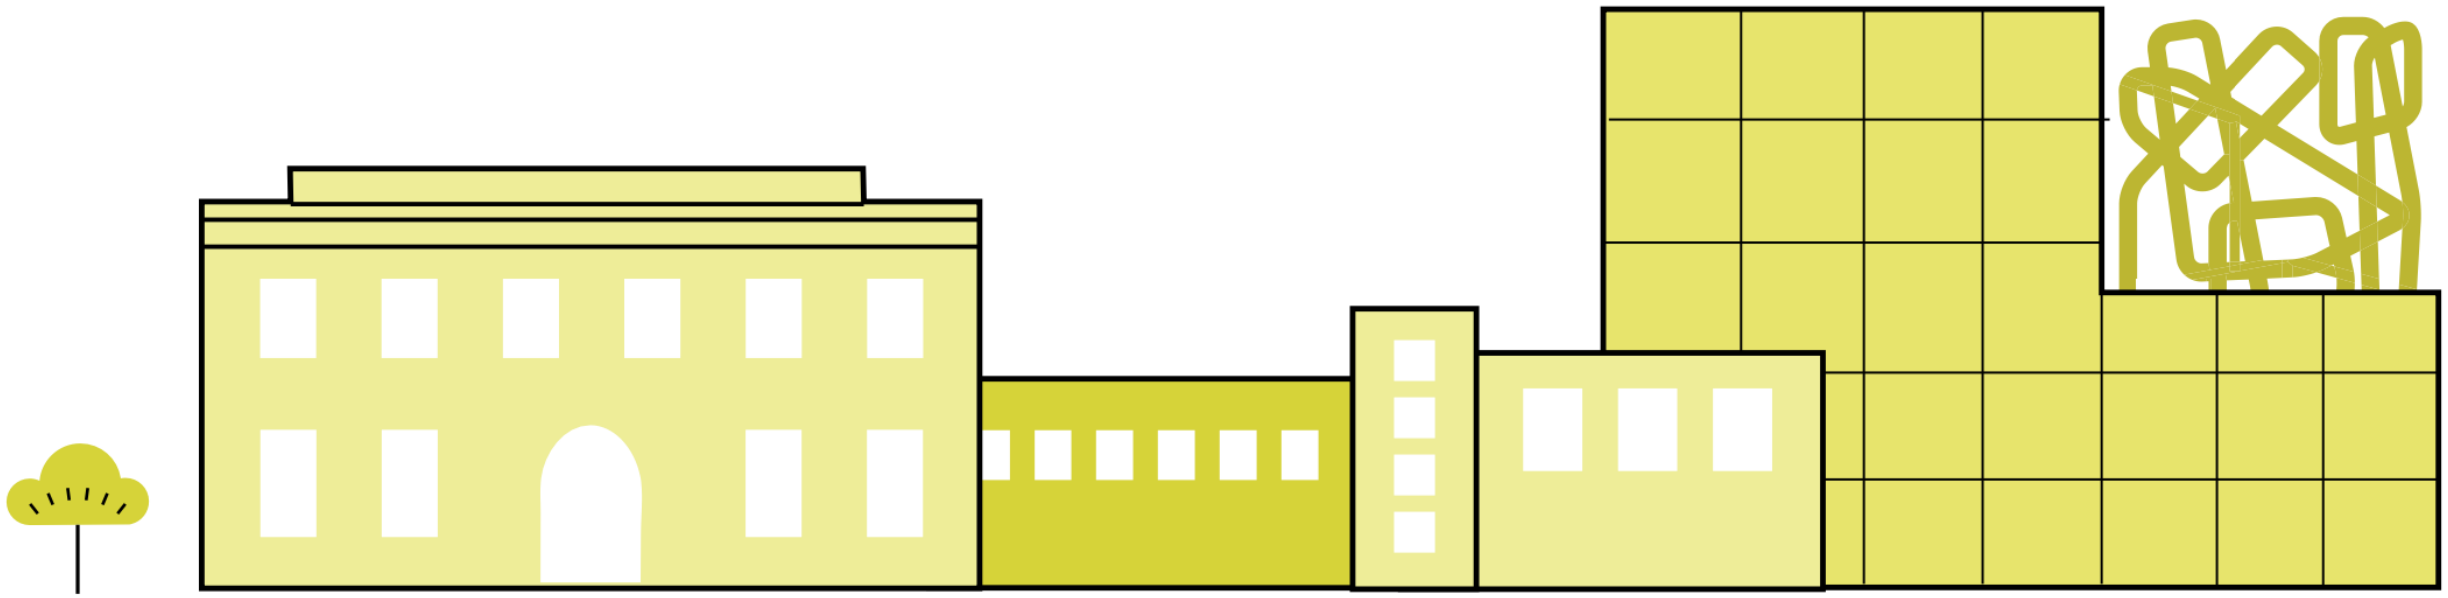

Some terms

- Natural language processing
- Sentiment analysis
- Embeddings
- Classification
- Zero-shot
- Large language models

Analysing survey data

What is it? / Technology used: Semantic analysis, word frequency analysis, zero-shot classification

Motivation: Large number of free text comments, need to extract main themes

Benefits: Can read and categorise all free text comments, effectively automating qualitative research process

Ethics: Data exposure: use local LLM, exposure of human qual. researcher role (we don't have one, they can design and orchestrate the workflow)

Survey analysis: the beginning... 2022 word clouds and sentiment analysis

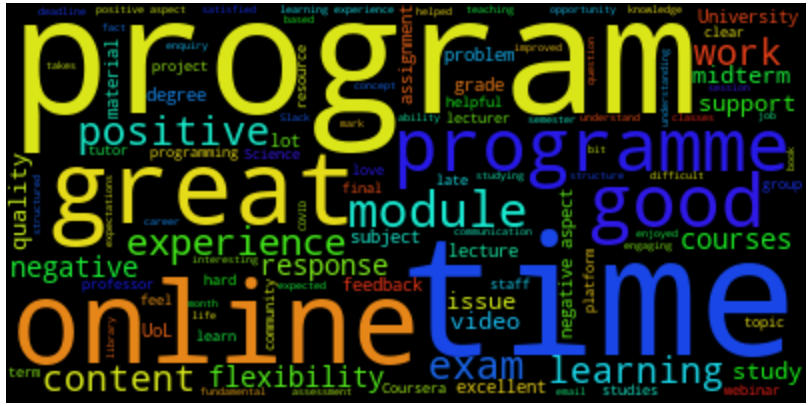

Comment	pos sentiment	neg sentiment
Really my biggest issue with the program is the lack of a direct communica	0.0001752172	0.9998248
The new tuition fee platform is bad. The exchange rate is too high, and I lo	0.0001793209	0.9998206
This program is a cash-grab with barely any follow-through. I left the progr	0.000180058	0.9998199
Midterm results are tending to take a bit long, rendering the feedback usele	0.0001866855	0.9998133

Survey analysis: assigning sentiment to themes

word	high_neg	med_meg	middle	med_pos	high_pos
organis	1.3	5.6	0	0	2.3
communicat	7.6	11.1	5.9	0	4.6
response	16.6	33.3	5.9	0	9.7
reply	3.5	0	5.9	0	1.7
efficien	0.2	0	0	0	0.6
delay	9	0	0	8.3	1.7
flexi	3.5	5.6	0	8.3	13.7
run	0.7	0	0	0	1.7
change	2.8	5.6	11.8	0	2.3

Local LLMs for survey analysis 2023

Local LLM: runs on your own machine but limited capability

Zero-shot classification with local LLM: assign comments to themes

Sentiment analysis

Summarisation with sets (works with a small model)

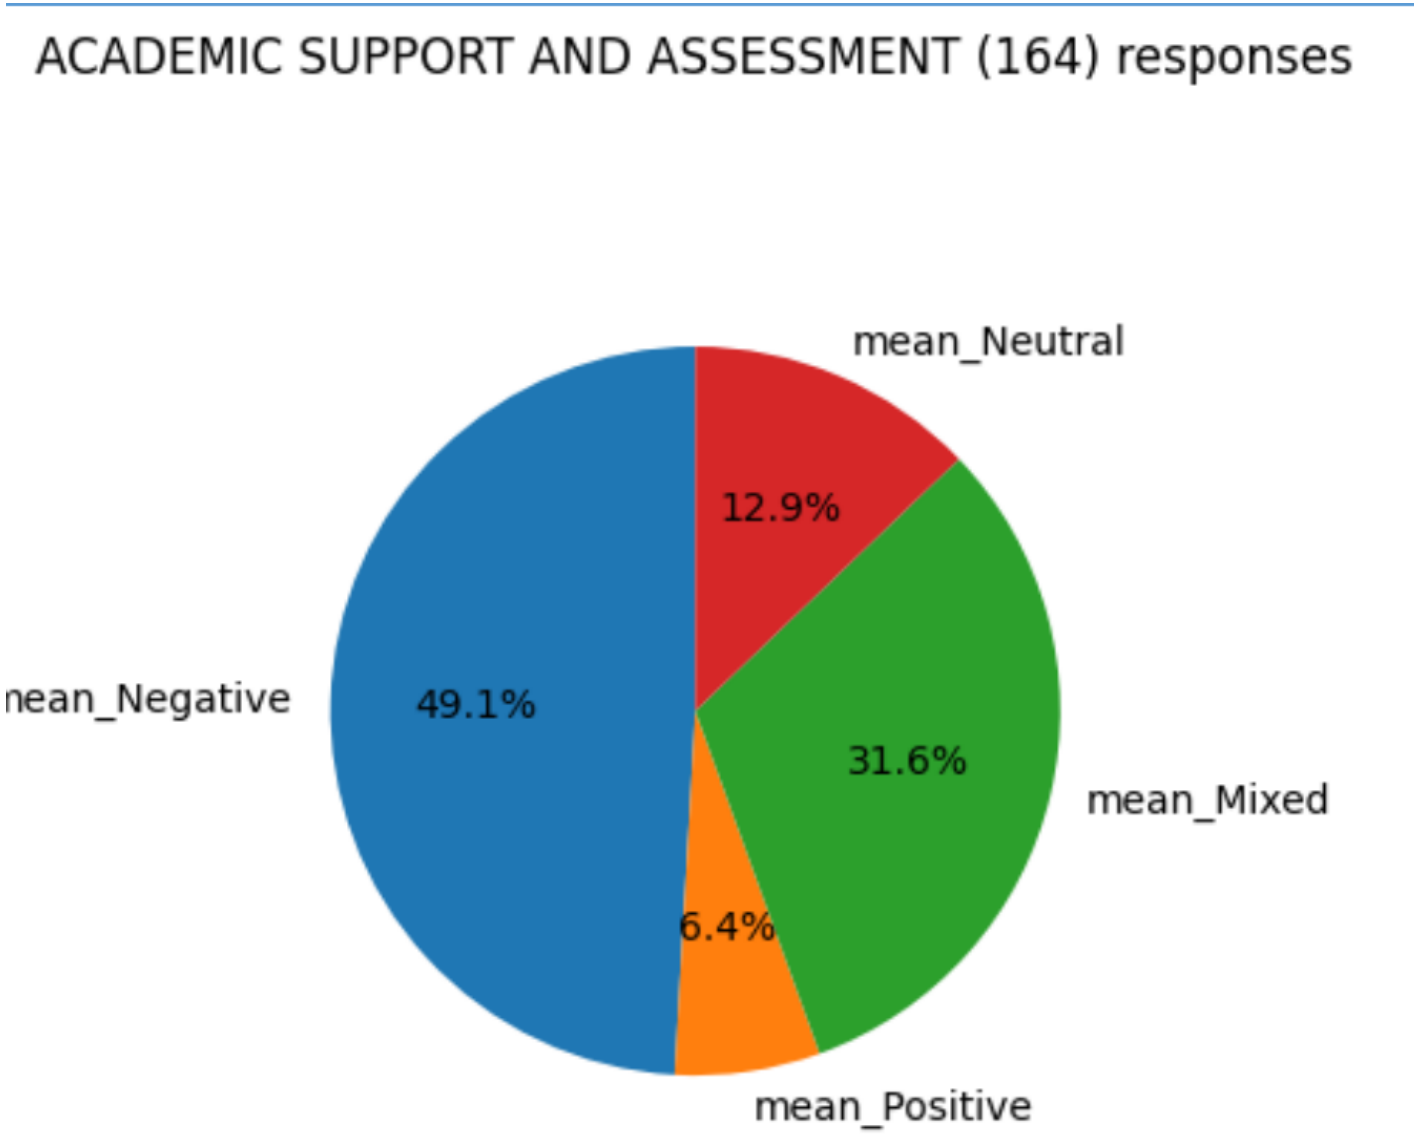

Case study 2: collusion detection

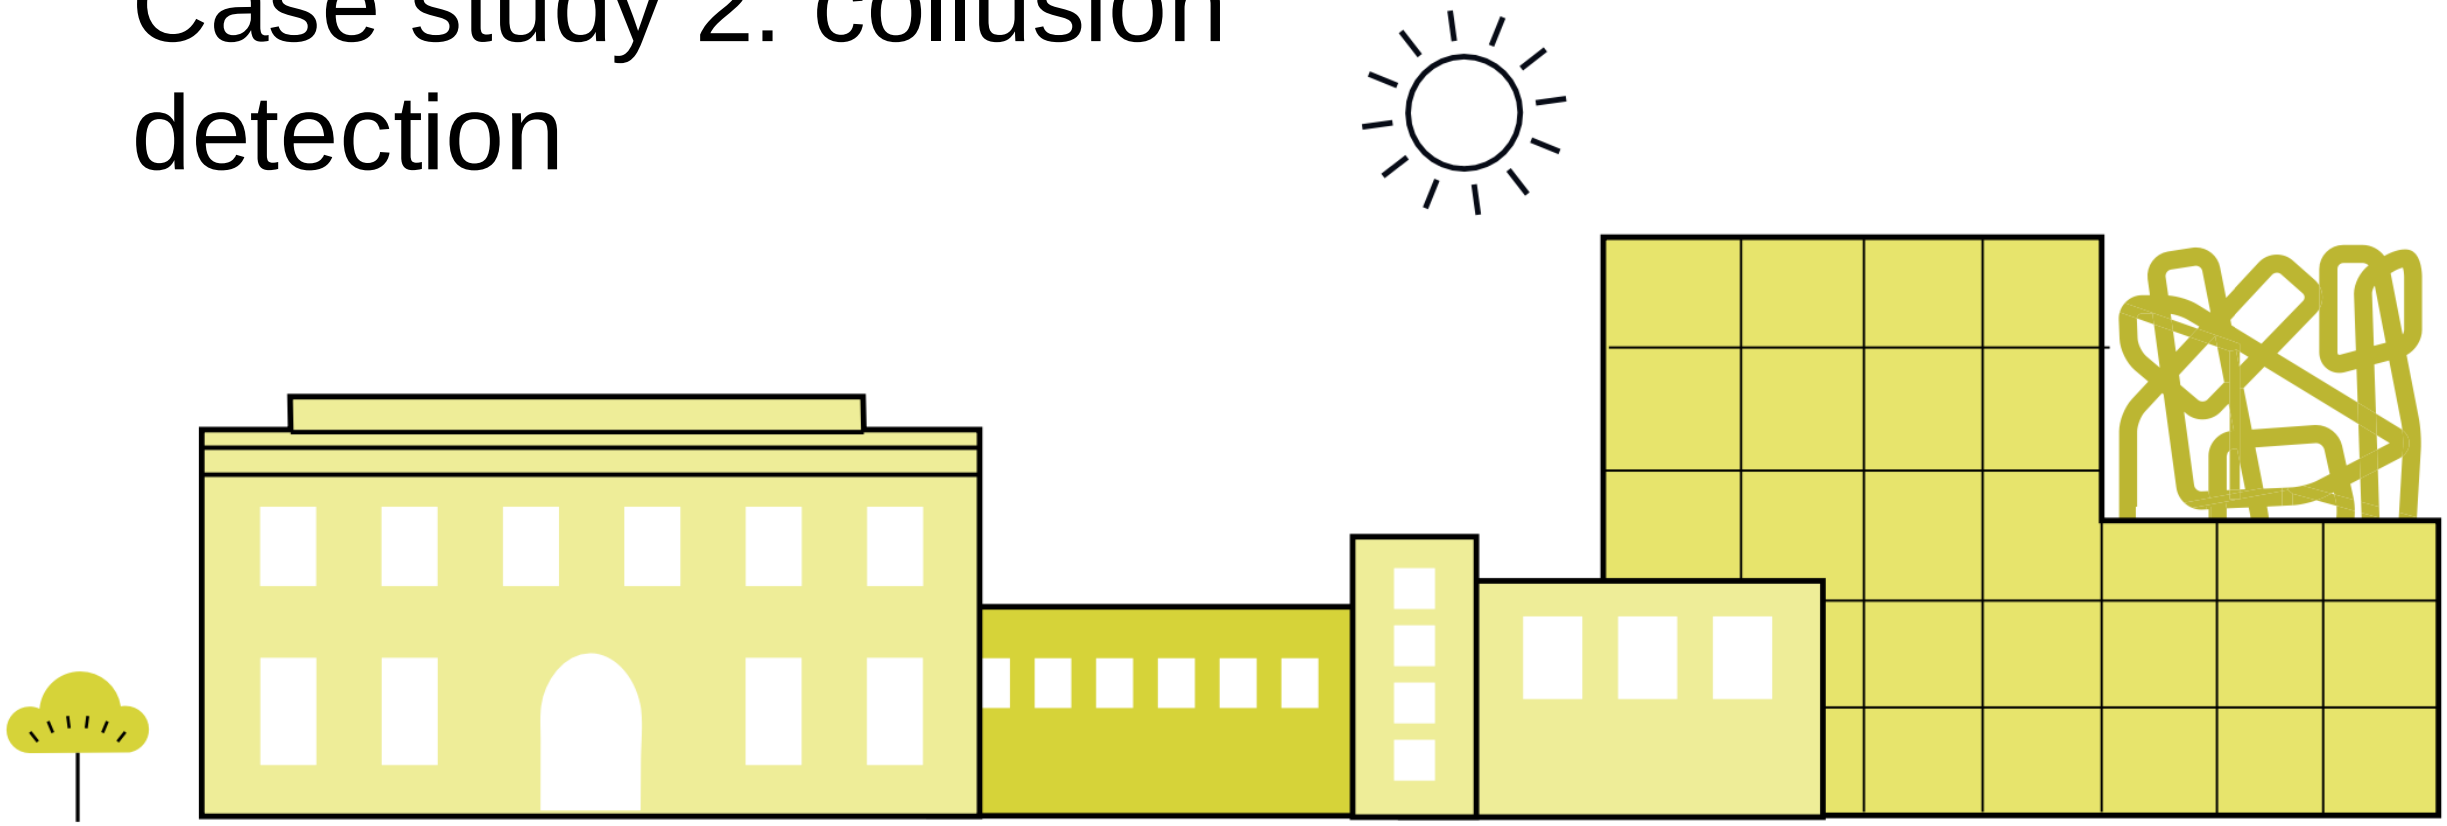

Collusion detection

What is it? / Technology used: Semantic embeddings

Motivation: Large number of exams, current system inadequate, impossible for human to porcess all comparisons

Benefits: Can process all answers to all exams, present information in compact digestible form

Ethics: Data exposure: use local model, exposure of examiners to automation - they didn't do this task before anyway, false positives: examiners have the final say

Collusion detection: research question

Is student A's answer to this question excessively similar to anyone else's?

Identify excessively similar pairs - dataset

Exam	Number of submissions	Unique questions	Number of answers
CM1030	494	10	2477
CM2010	562	18	6017
CM2015	518	17	4510
CM2025	326	14	2575
CM2040	311	17	3303
CM3005	215	14	1615
CM3010	306	11	1534
CM3020	109	14	904
CM3045	132	9	548
CM3055	80	13	585
CM3060	98	13	770
Total	3151	150	24838

TABLE I

Collusion detection: turnitin is not enough

B

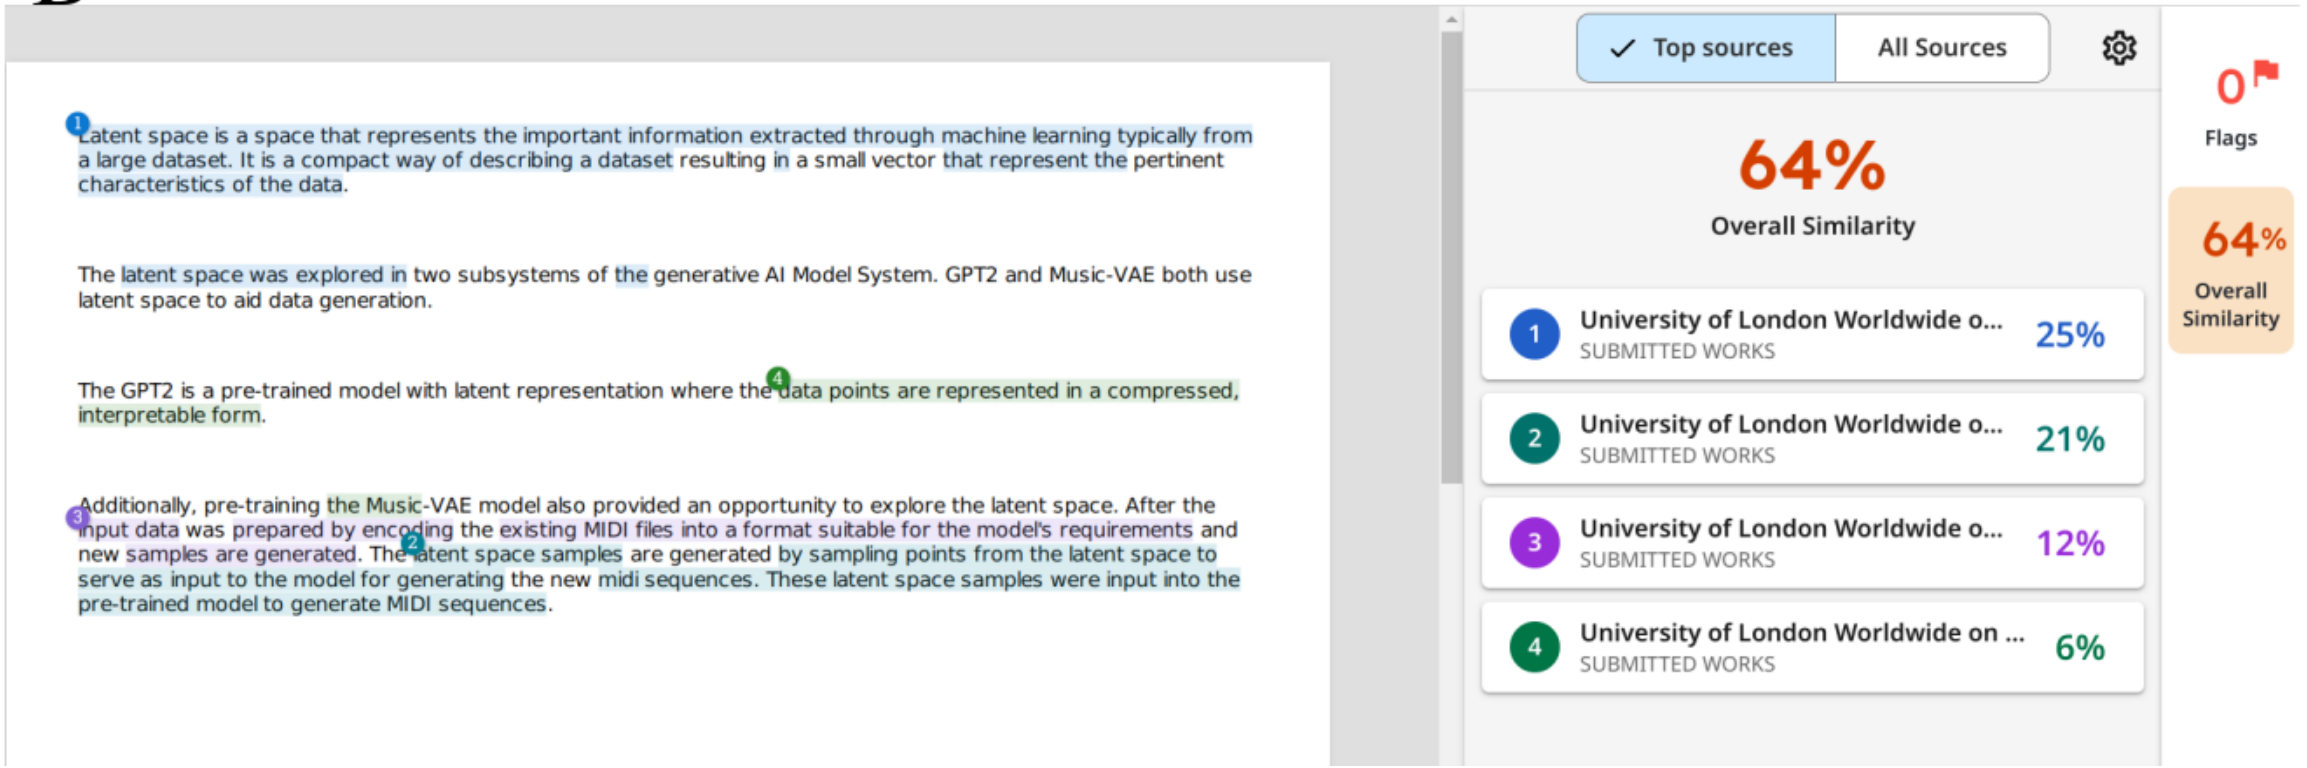

Collusion detection

Techniques used:

Texture Coordinate Manipulation: I
coordinate based on time.
This was done using the function u
* speed),cos(time * speed))* displa
where,
time - represents the shader's inter
speed- controls the rate of animati
displacement- determines how far

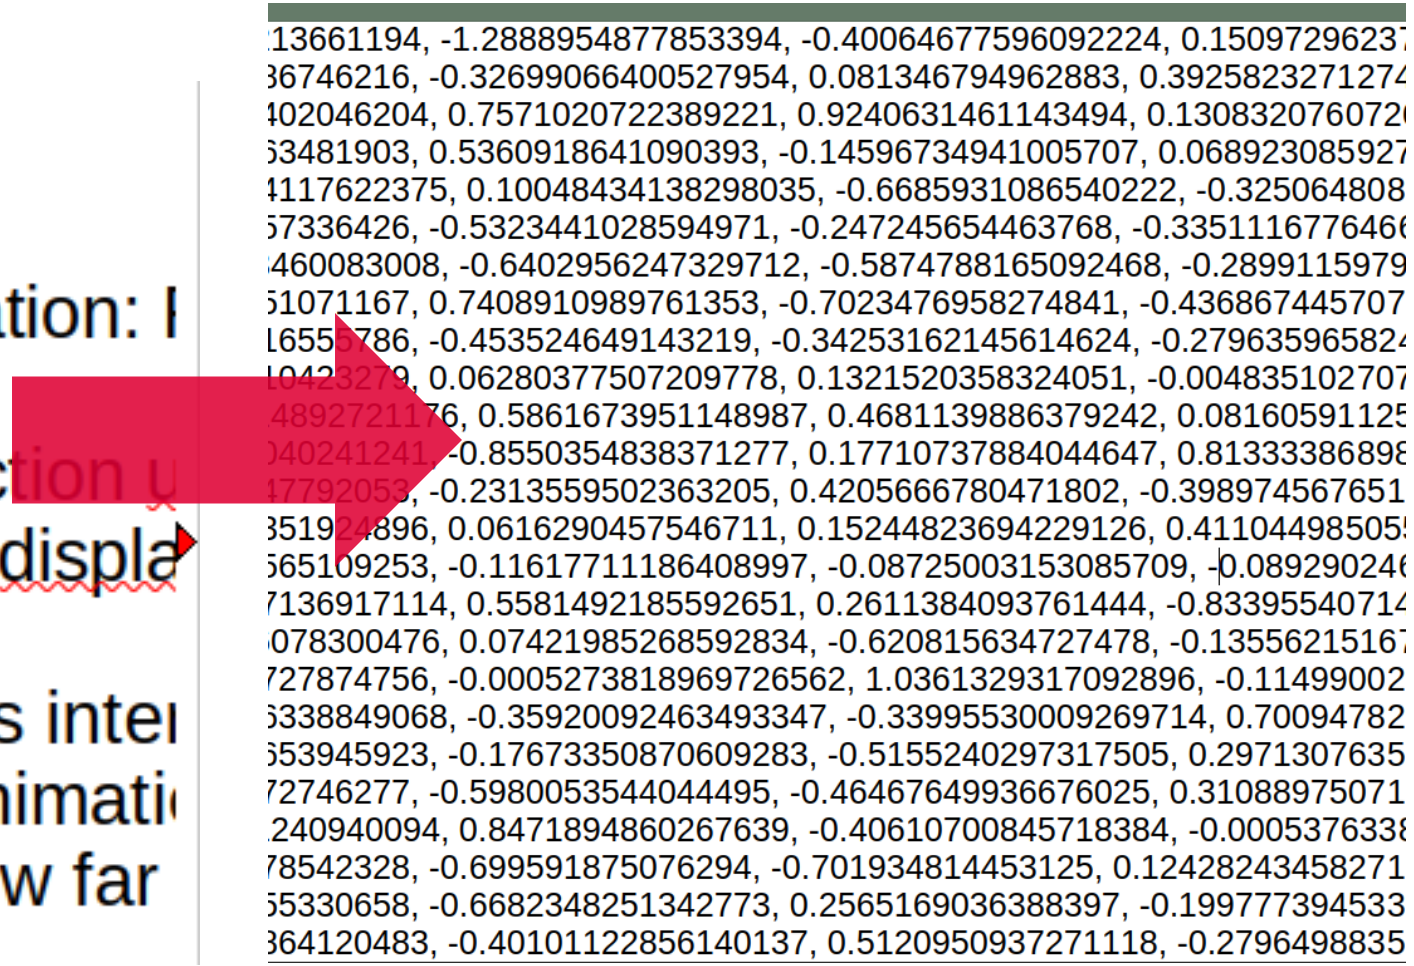

Embeddings: compute distance in semantic space

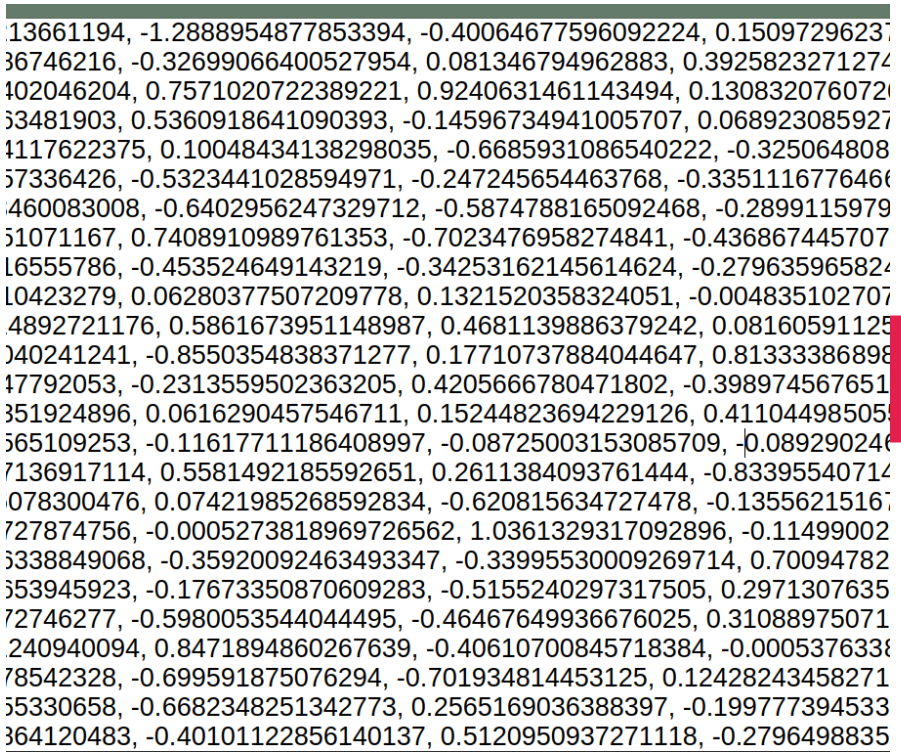

Student A's answer

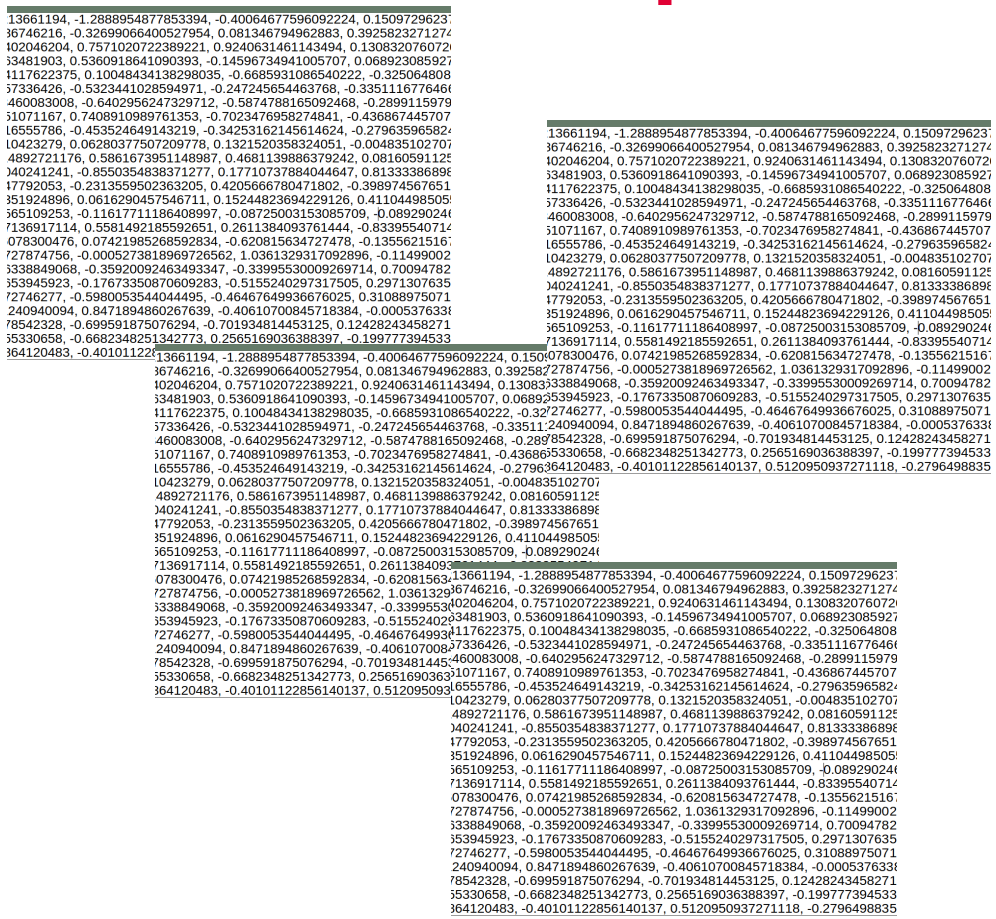

Other answers

Identify excessively similar pairs - report

Table summarising hits per filename

Filename	Number of Hits (where this exam had an answer very close to someone else)
SP3404_210905812.pdf	15
SP1675_210905812.pdf	15
SP2200_210905812.pdf	12
SP3502_210905812.pdf	4
SP0306_210905812.pdf	4

Identify excessively similar pairs - report

course_code	Filename	closest_file	Question Number	Answer	closest_answer	question_z_score
CM2010	SP3404_210905812.pdf	SP1675_210905812.pdf	4.4	The given function is not a robust program because of division by zero and input validation. For division by zero, the function does not account for the scenario where the elapsed time t is zero. If t is zero, division by zero will occur,leading to a runtime error. This is a common issue that can cause the program to crash. For input validation, the function does not perform any input validation. It assumes that v, v0 and t will always be provided as valid numerical values. If invalid inputs , such as strings or other non-numeric values, are passed to the function, it may produce unexpected results or errors. To make the program more robust, we can implement the following mechanism of input validation and exception handling. For input validation, we can check if t is zero before performing the division to avoid division by zero error and validate that v, v0 and t are valid numerical values before proceeding with the calculation. For exception handling, we can implement try-except blocks to catch potential errors, such as division by zero, and handle them gracefully , raise specific exceptions, or provide meaningful error messages to inform the user of the issue.	The given function is not a robust program because of division by zero and input validation. For division by zero, the function does not account for the scenario where the elapsed time t is zero. if t is zero, division by zero will occur, leading to a runtime error. This is a common issue that can cause the program to crash. For input validation, the function does not perform any input validation. It assumes that v, v0 and t will always be provided as valid numerical values. If invalid inputs, such as strings or other non-numeric values, are passed to the function, it may produce unexpected results or errors. To make the program more robust, we can implement the following mechanisms of input validation and exception handling. For input validation, we can check if t is zero before performing the division to avoid division by zero error and validate that v, v0 and t are valid numerical values before proceeding with the calculation. For exception handling, we can implement try-except blocks to catch potential errors, such as division by zero and handle them gracefully, raise specific exceptions or provide meaningful error messages to inform the user of the issue.	-3.551538
					The three git commands are git clone, git status and git	

Results

Minimal effort for module leaders:

Cases are auto-detected

Reports are auto-generated

Module leaders have sign-off

14 cases brought forward

Next step: detect AI content

Generate lots of answers to questions with LLMs of various sorts

Is student A's answer excessively semantically similar to the LLM answers?

Case study 3: generating exams

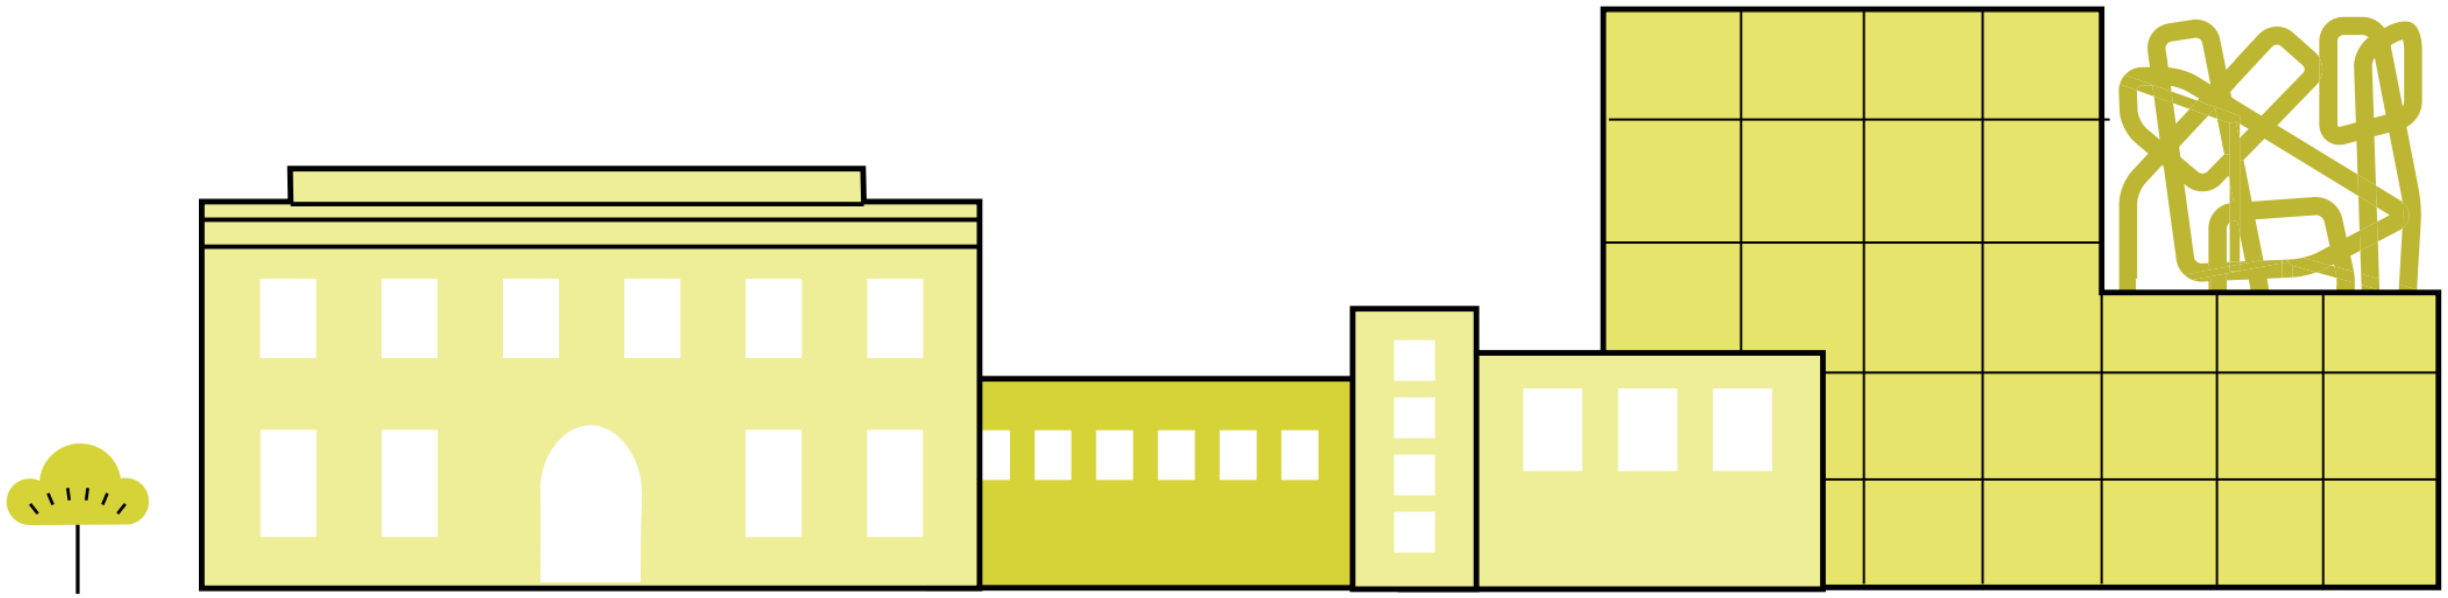

WIP: generating exams with LLMs

What is it? / Technology used: LocalLLM, pipelines, prompt engineering

Motivation: Continuous exam setting cycle, difficult to complete on time, inconsistency of exam design. Must have exams: gold standard

Benefits: Break perceived barrier of initial draft, exam setting task is simplified to an iteration and improvement task

Ethics: Data exposure: use local model, exposure of examiners to automation or increased workload

Side note: generating quizzes

Various projects to generate quizzes for scaffolding purposeS:

Context: video transcript + learning objectves:

Generate answers + feedback

Example exam structure: not just generate me an exam about x!

Request from the LLM, several questions

1) Fact recall
2) Given a scenario/ problem, provide a high level description of a solution
3) Describe the technical details of the solution
4) Critically evaluate your approach and other possible approaches

Implementation

Use local Llama 70b model via ‘open-ai’-like API

Data on courses all prepared in consistent form: topic lists, learning objectives, weekly descriptions of content

Scripts to iterate over courses and generate exams

Prompt

I am setting an exam for an undergraduate computer science course.

Here are the main topics in the course:

##topis##

Here ar descriptions of each week of content in the course:

##weeklies##

Here are the weekly learning objectives for the course.

##weeklylos##

Here are the most important descriptors of the course, the course-level learning objecti

ves. These are the items the exam should aim to evaluate the student against:

Prompt

Please write one exam question for me. The question should start with some factual questions which check the student's recall of basic facts taught in the course. The question should then ask the student to apply some skills learnt in the course to a particular case study or problem. The student should have to describe their approach or solution at a high level, then to give a more technical and specific description of their solution, using an appropriate format shown in the course. The final part of the question should ask the student to critically evaluate two approaches to solving a particular problem, highlighting advantages and disadvantages of the different approaches using knowledge gained in the course.

Case study 4: student assistant [very much WIP]

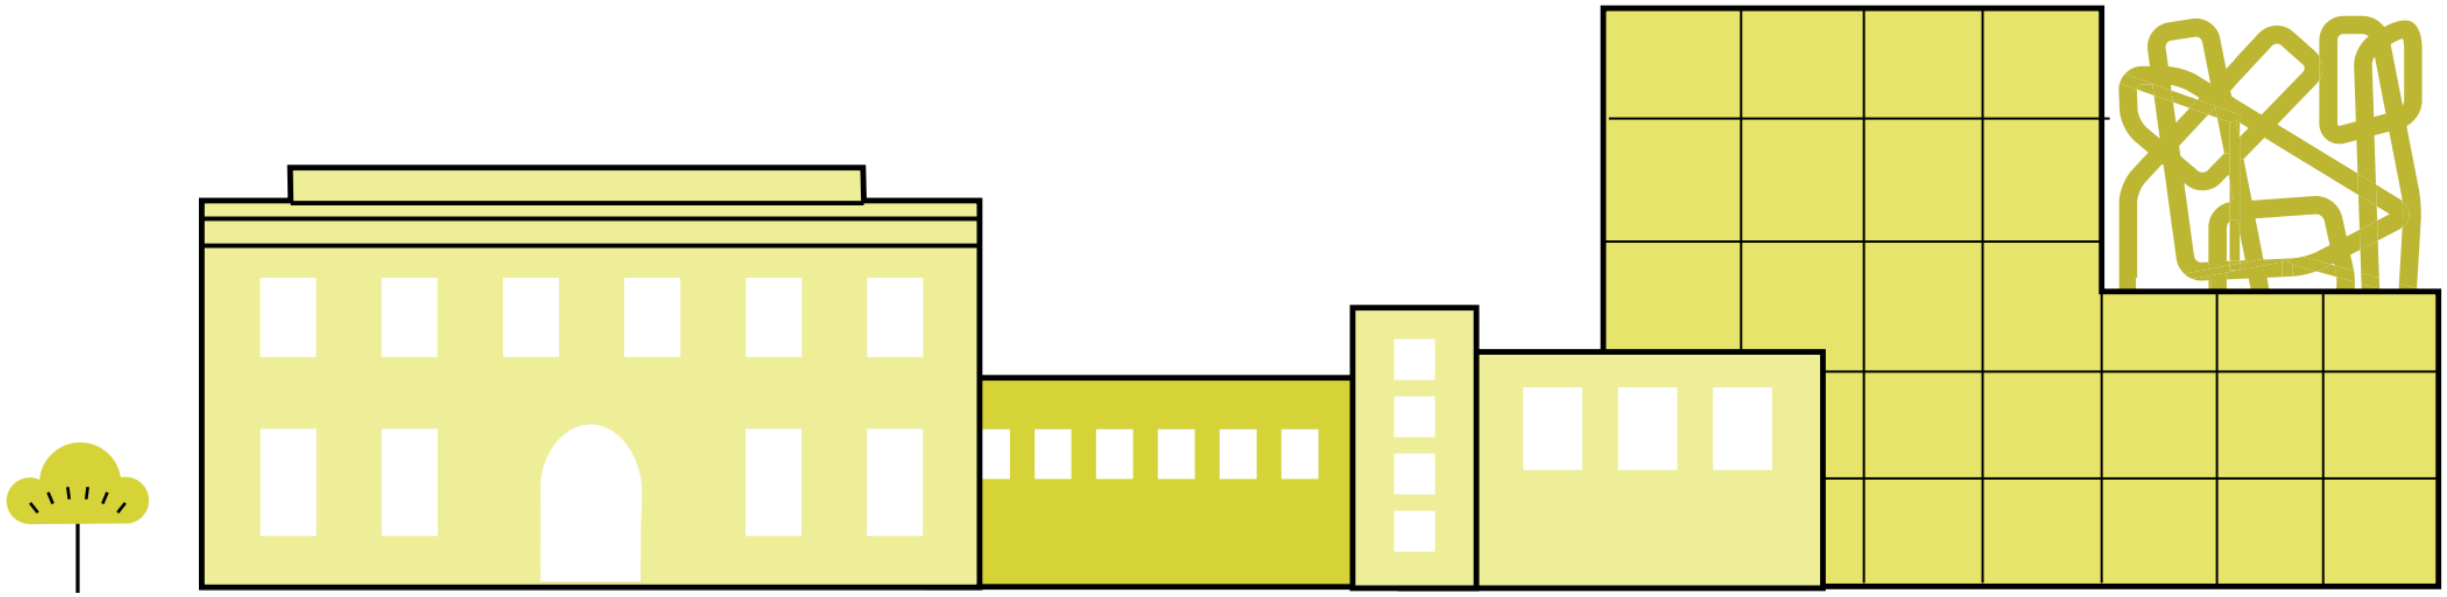

Student assistant

What is it? / Technology used: LocalLLM, RAG pipelines, prompt engineering,

Motivation: salespeople, hype, seems like a good idea, surely we can do it better?, they are already doing it with GPT

Benefits: Structured pedagogic workflows, reduce data exposure for students, proper access to local resources (actually fair use!)

Ethics: Data exposure: use local model, exposure of tutors to automation, cheating on courseworks

Student assistant agent: design

RAG pipeline:

Compose prompts from fragments of course materials through semantic search

Structured interaction: menu of options

Local LLM: Llama 3, Phi etc. Good enough? Fine-tuning? Local RAG?

Running locally: no inference cost for institution

Summary

- Some terms
- Local LLMs for survey analysis
- LLM embeddings for collusion detection
- WIP: Generating exam papers
- WIP: course chatbo

Goldsmiths
UNIVERSITY OF LONDON

UNIVERSITY
OF LONDON

london.ac.uk

Senate House, Malet Street,
London WC1E 7HU