

Retrieval Augmented Generation (RAG)



Dr. Antonio Mastropaolo

Instructor

Mr. Alvi Haque



Teaching Assistant



WILLIAM & MARY

CHARTERED 1693

Spring 2026



antoniomastropaolo.com



[aura-se-lab.github.io](https://github.com/aura-se-lab)



Retrieval Augmented Generation



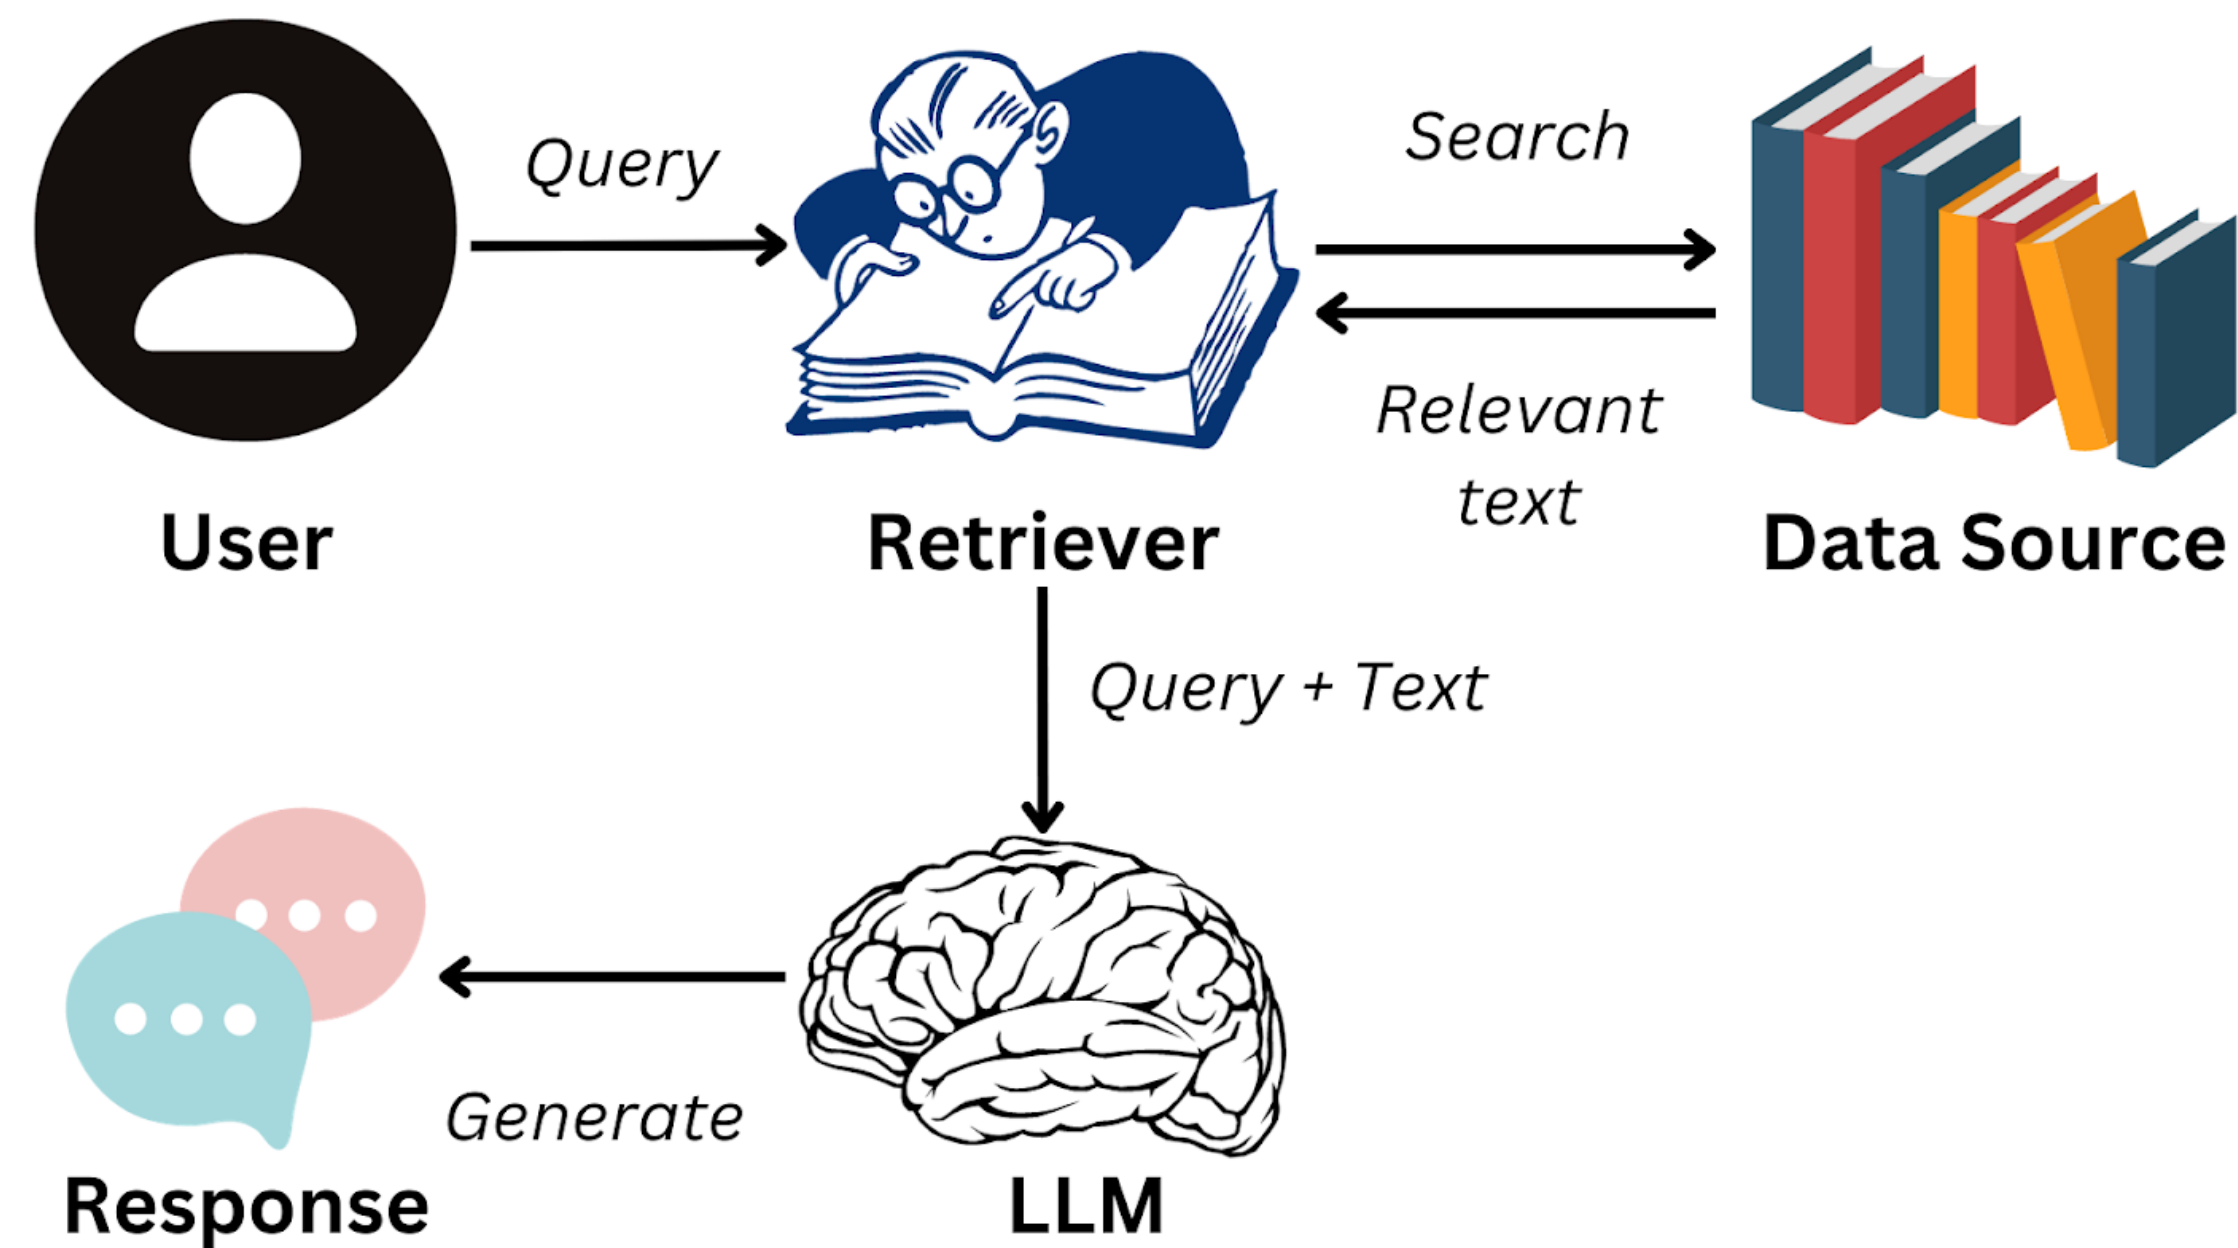
antoniomastropaolo.com



[aura-se-lab.github.io](https://github.com/aura-se-lab)



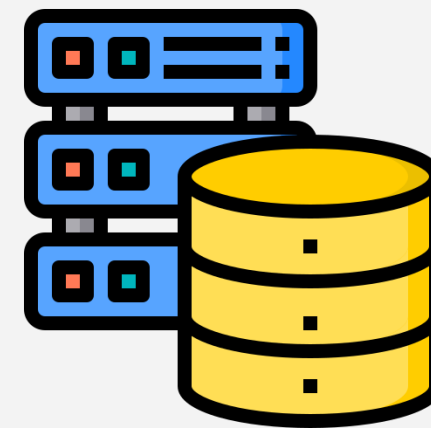
Retrieval Augmented Generation



Retrieval Augmented Generation

KB (Knowledge Base)

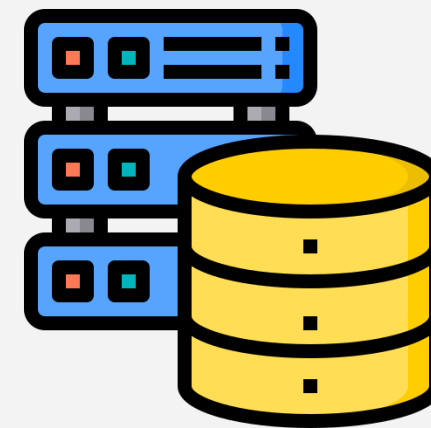
Homework can be up to three days late with a 10% penalty per day.
Assignments submitted beyond three days receive no credit.
The midterm is closed-book; one page of notes is allowed.
The final exam is closed-book; a single formula sheet may be used.
Quizzes are closed-book and closed-notes.
Quizzes count for ten percent of the final grade.
Attendance is tracked; more than two unexcused absences lowers your participation score.
Active discussion in class contributes to participation.
Office hours are Tuesdays and Thursdays from two to four in room 201.
The TA holds office hours on Wednesdays three to five on Zoom.
You may book a one-on-one appointment for extra help.
Lecture slides are posted weekly on the course page.
Readings are listed on the schedule and updated each week.
Use the discussion forum for questions and peer support.
Slides and readings are available online.
Homework up to three days late loses ten percent per day.
Work one day late loses ten percent; two days late loses twenty.
Submitting beyond three days late earns zero points.
One formula sheet is permitted on the midterm.
Participation is worth ten percent overall.
Office hours are available twice a week.



Retrieval Augmented Generation

KB (Knowledge Base)

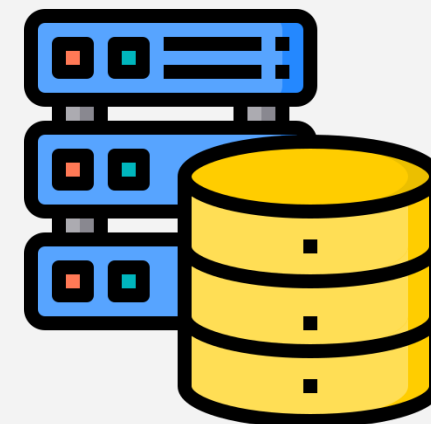
Homework can be up to three days late with a 10% penalty per day.
Assignments submitted beyond three days receive no credit.
The midterm is closed-book; one page of notes is allowed.
The final exam is closed-book; a single formula sheet may be used.
Quizzes are closed-book and closed-notes.
Quizzes count for ten percent of the final grade.
Attendance is tracked; more than two unexcused absences lowers your participation score.
Active discussion in class contributes to participation.
Office hours are Tuesdays and Thursdays from two to four in room 201.
The TA holds office hours on Wednesdays three to five on Zoom.
You may book a one-on-one appointment for extra help.
Lecture slides are posted weekly on the course page.
Readings are listed on the schedule and updated each week.
Use the discussion forum for questions and peer support.
Slides and readings are available online.
Homework up to three days late loses ten percent per day.
Work one day late loses ten percent; two days late loses twenty.
Submitting beyond three days late earns zero points.
One formula sheet is permitted on the midterm.
Participation is worth ten percent overall.
Office hours are available twice a week.



Retrieval Augmented Generation

KB (Knowledge Base)

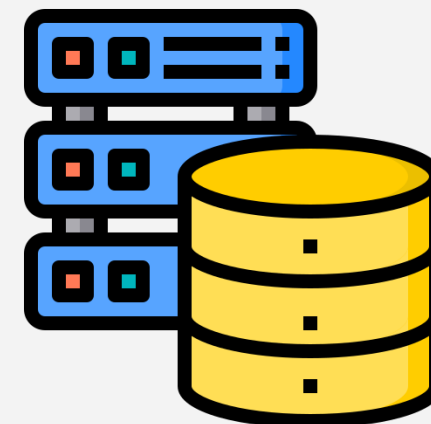
Homework can be up to three days late with a 10% penalty per day.
Assignments submitted beyond three days receive no credit.
The midterm is closed-book; one page of notes is allowed.
The final exam is closed-book; a single formula sheet may be used.
Quizzes are closed-book and closed-notes.
Quizzes count for ten percent of the final grade.
Attendance is tracked; more than two unexcused absences lowers your participation score.
Active discussion in class contributes to participation.
Office hours are Tuesdays and Thursdays from two to four in room 201.
The TA holds office hours on Wednesdays three to five on Zoom.
You may book a one-on-one appointment for extra help.
Lecture slides are posted weekly on the course page.
Readings are listed on the schedule and updated each week.
Use the discussion forum for questions and peer support.
Slides and readings are available online.
Homework up to three days late loses ten percent per day.
Work one day late loses ten percent; two days late loses twenty.
Submitting beyond three days late earns zero points.
One formula sheet is permitted on the midterm.
Participation is worth ten percent overall.
Office hours are available twice a week.



Retrieval Augmented Generation

KB (Knowledge Base)

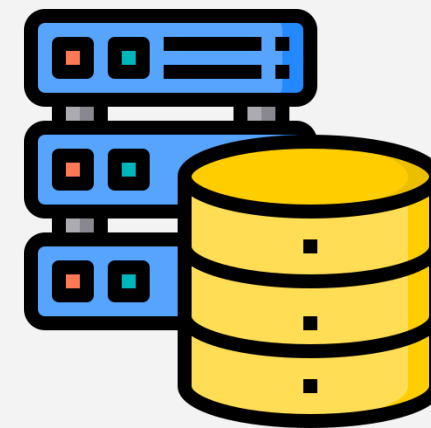
Homework can be up to three days late with a 10% penalty per day.
Assignments submitted beyond three days receive no credit.
The midterm is closed-book; one page of notes is allowed.
The final exam is closed-book; a single formula sheet may be used.
Quizzes are closed-book and closed-notes.
Quizzes count for ten percent of the final grade.
Attendance is tracked; more than two unexcused absences lowers your participation score.
Active discussion in class contributes to participation.
Office hours are Tuesdays and Thursdays from two to four in room 201.
The TA holds office hours on Wednesdays three to five on Zoom.
You may book a one-on-one appointment for extra help.
Lecture slides are posted weekly on the course page.
Readings are listed on the schedule and updated each week.
Use the discussion forum for questions and peer support.
Slides and readings are available online.
Homework up to three days late loses ten percent per day.
Work one day late loses ten percent; two days late loses twenty.
Submitting beyond three days late earns zero points.
One formula sheet is permitted on the midterm.
Participation is worth ten percent overall.
Office hours are available twice a week.



Retrieval Augmented Generation

KB (Knowledge Base)

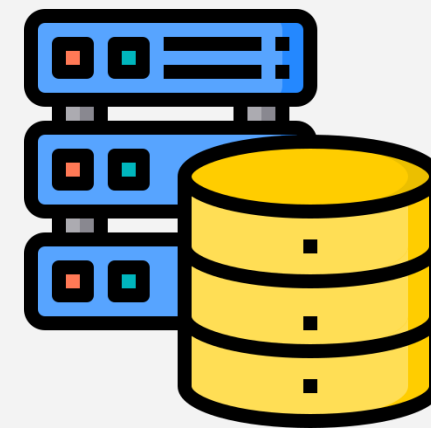
Homework can be up to three days late with a 10% penalty per day.
Assignments submitted beyond three days receive no credit.
The midterm is closed-book; one page of notes is allowed.
The final exam is closed-book; a single formula sheet may be used.
Quizzes are closed-book and closed-notes.
Quizzes count for ten percent of the final grade.
Attendance is tracked; more than two unexcused absences lowers your participation score.
Active discussion in class contributes to participation.
Office hours are Tuesdays and Thursdays from two to four in room 201.
The TA holds office hours on Wednesdays three to five on Zoom.
You may book a one-on-one appointment for extra help.
Lecture slides are posted weekly on the course page.
Readings are listed on the schedule and updated each week.
Use the discussion forum for questions and peer support.
Slides and readings are available online.
Homework up to three days late loses ten percent per day.
Work one day late loses ten percent; two days late loses twenty.
Submitting beyond three days late earns zero points.
One formula sheet is permitted on the midterm.
Participation is worth ten percent overall.
Office hours are available twice a week.



Retrieval Augmented Generation

KB (Knowledge Base)

Homework can be up to three days late with a 10% penalty per day.
Assignments submitted beyond three days receive no credit.
The midterm is closed-book; one page of notes is allowed.
The final exam is closed-book; a single formula sheet may be used.
Quizzes are closed-book and closed-notes.
Quizzes count for ten percent of the final grade.
Attendance is tracked; more than two unexcused absences lowers your participation score.
Active discussion in class contributes to participation.
Office hours are Tuesdays and Thursdays from two to four in room 201.
The TA holds office hours on Wednesdays three to five on Zoom.
You may book a one-on-one appointment for extra help.
Lecture slides are posted weekly on the course page.
Readings are listed on the schedule and updated each week.
Use the discussion forum for questions and peer support.
Slides and readings are available online.
Homework up to three days late loses ten percent per day.
Work one day late loses ten percent; two days late loses twenty.
Submitting beyond three days late earns zero points.
One formula sheet is permitted on the midterm.
Participation is worth ten percent overall.
Office hours are available twice a week.



INPUT

You are a syllabus policy assistant.
Use only the retrieved facts from
`coll100_policy_kb.txt`.
If evidence is insufficient, reply
exactly: "Not found in the KB."

Can I skip class in
your opinion?



antoniomastropaolo.com



[aura-se-lab.github.io](https://github.com/aura-se-lab)



Retrieval Augmented Generation

KB (Knowledge Base)

Homework can be up to three days late with a 10% penalty per day. ←

Assignments submitted beyond three days receive no credit.

The midterm is closed-book; one page of notes is allowed.

The final exam is closed-book; a single formula sheet may be used.

Quizzes are closed-book and closed-notes.

Quizzes count for ten percent of the final grade.

Attendance is tracked; more than two unexcused absences lowers your participation score.

Active discussion in class contributes to participation.

Office hours are Tuesdays and Thursdays from two to four in room 201.

The TA holds office hours on Wednesdays three to five on Zoom.

You may book a one-on-one appointment for extra help.

Lecture slides are posted weekly on the course page.

Readings are listed on the schedule and updated each week.

Use the discussion forum for questions and peer support.

Slides and readings are available online.

Homework up to three days late loses ten percent per day.

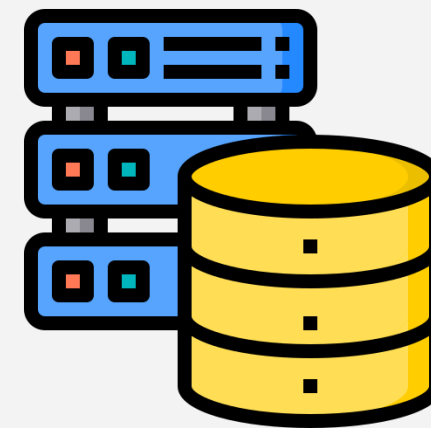
Work one day late loses ten percent; two days late loses twenty.

Submitting beyond three days late earns zero points.

One formula sheet is permitted on the midterm.

Participation is worth ten percent overall.

Office hours are available twice a week.



INPUT

You are a syllabus policy assistant.
Use only the retrieved facts from
`coll100_policy_kb.txt`.
If evidence is insufficient, reply
exactly: "Not found in the KB."

Can I skip class in
your opinion?

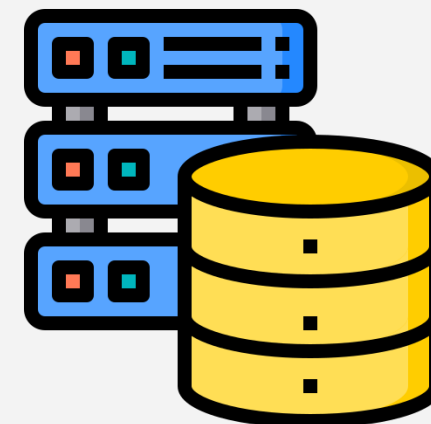


Retrieval Augmented Generation

KB (Knowledge Base)

Homework can be up to three days late with a 10% penalty per day.
Assignments submitted beyond three days receive no credit.
The midterm is closed-book; one page of notes is allowed.
The final exam is closed-book; a single formula sheet may be used.
Quizzes are closed-book and closed-notes.
Quizzes count for ten percent of the final grade.
Attendance is tracked; more than two unexcused absences lowers your participation score.
Active discussion in class contributes to participation.
Office hours are Tuesdays and Thursdays from two to four in room 201.
The TA holds office hours on Wednesdays three to five on Zoom.
You may book a one-on-one appointment for extra help.
Lecture slides are posted weekly on the course page.
Readings are listed on the schedule and updated each week.
Use the discussion forum for questions and peer support.
Slides and readings are available online.
Homework up to three days late loses ten percent per day.
Work one day late loses ten percent; two days late loses twenty.
Submitting beyond three days late earns zero points.
One formula sheet is permitted on the midterm.
Participation is worth ten percent overall.
Office hours are available twice a week.

0.4



INPUT

You are a syllabus policy assistant.
Use only the retrieved facts from
`coll100_policy_kb.txt`.
If evidence is insufficient, reply
exactly: "Not found in the KB."

Can I skip class in
your opinion?



antoniomastropaolo.com



[aura-se-lab.github.io](https://github.com/aura-se-lab)

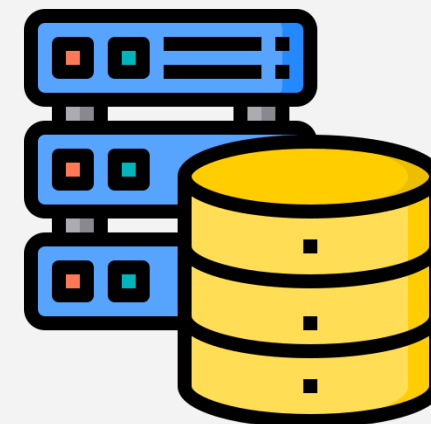
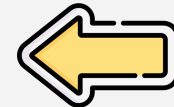


Retrieval Augmented Generation

KB (Knowledge Base)

Homework can be up to three days late with a 10% penalty per day.
Assignments submitted beyond three days receive no credit.
The midterm is closed-book; one page of notes is allowed.
The final exam is closed-book; a single formula sheet may be used.
Quizzes are closed-book and closed-notes.
Quizzes count for ten percent of the final grade.
Attendance is tracked; more than two unexcused absences lowers your participation score.
Active discussion in class contributes to participation.
Office hours are Tuesdays and Thursdays from two to four in room 201.
The TA holds office hours on Wednesdays three to five on Zoom.
You may book a one-on-one appointment for extra help.
Lecture slides are posted weekly on the course page.
Readings are listed on the schedule and updated each week.
Use the discussion forum for questions and peer support.
Slides and readings are available online.
Homework up to three days late loses ten percent per day.
Work one day late loses ten percent; two days late loses twenty.
Submitting beyond three days late earns zero points.
One formula sheet is permitted on the midterm.
Participation is worth ten percent overall.
Office hours are available twice a week.

0.4
0.7



INPUT

You are a syllabus policy assistant.
Use only the retrieved facts from
`coll100_policy_kb.txt`.
If evidence is insufficient, reply
exactly: "Not found in the KB."

Can I skip class in
your opinion?



antoniomastropaolo.com



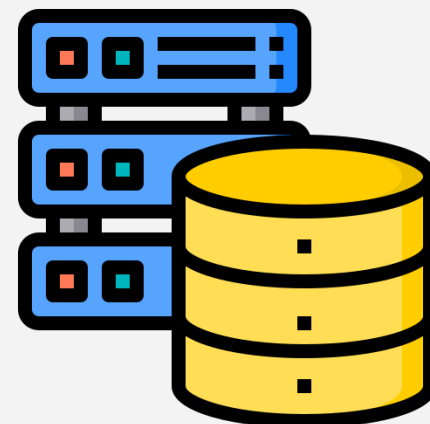
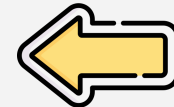
[aura-se-lab.github.io](https://github.com/aura-se-lab)



Retrieval Augmented Generation

KB (Knowledge Base)

Homework can be up to three days late with a 10% penalty per day.
Assignments submitted beyond three days receive no credit.
The midterm is closed-book; one page of notes is allowed.
The final exam is closed-book; a single formula sheet may be used.
Quizzes are closed-book and closed-notes.
Quizzes count for ten percent of the final grade.
Attendance is tracked; more than two unexcused absences lowers your participation score.
Active discussion in class contributes to participation.
Office hours are Tuesdays and Thursdays from two to four in room 201.
The TA holds office hours on Wednesdays three to five on Zoom.
You may book a one-on-one appointment for extra help.
Lecture slides are posted weekly on the course page.
Readings are listed on the schedule and updated each week.
Use the discussion forum for questions and peer support.
Slides and readings are available online.
Homework up to three days late loses ten percent per day.
Work one day late loses ten percent; two days late loses twenty.
Submitting beyond three days late earns zero points.
One formula sheet is permitted on the midterm.
Participation is worth ten percent overall.
Office hours are available twice a week.



0.4
0.7
0.1



INPUT

You are a syllabus policy assistant.
Use only the retrieved facts from
`coll100_policy_kb.txt`.
If evidence is insufficient, reply
exactly: "Not found in the KB."

Can I skip class in
your opinion?



antoniomastropaolo.com



[aura-se-lab.github.io](https://github.com/aura-se-lab)

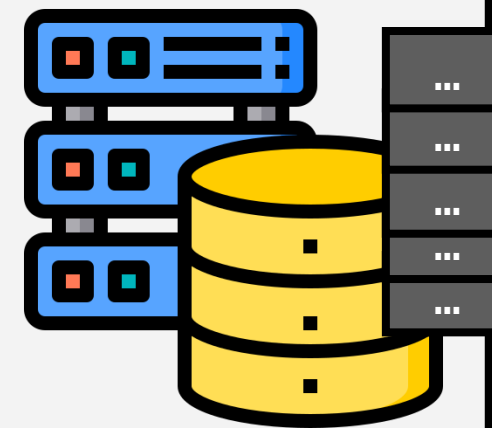


Retrieval Augmented Generation

KB (Knowledge Base)

Homework can be up to three days late with a 10% penalty per day.
Assignments submitted beyond three days receive no credit.
The midterm is closed-book; one page of notes is allowed.
The final exam is closed-book; a single formula sheet may be used.
Quizzes are closed-book and closed-notes.
Quizzes count for ten percent of the final grade.
Attendance is tracked; more than two unexcused absences lowers your participation score.
Active discussion in class contributes to participation.
Office hours are Tuesdays and Thursdays from two to four in room 201.
The TA holds office hours on Wednesdays three to five on Zoom.
You may book a one-on-one appointment for extra help.
Lecture slides are posted weekly on the course page.
Readings are listed on the schedule and updated each week.
Use the discussion forum for questions and peer support.
Slides and readings are available online.
Homework up to three days late loses ten percent per day.
Work one day late loses ten percent; two days late loses twenty.
Submitting beyond three days late earns zero points.
One formula sheet is permitted on the midterm.
Participation is worth ten percent overall.
Office hours are available twice a week.

0.4
0.7
0.1
...



INPUT

You are a syllabus policy assistant.
Use only the retrieved facts from
`coll100_policy_kb.txt`.
If evidence is insufficient, reply
exactly: "Not found in the KB."

Can I skip class in
your opinion?



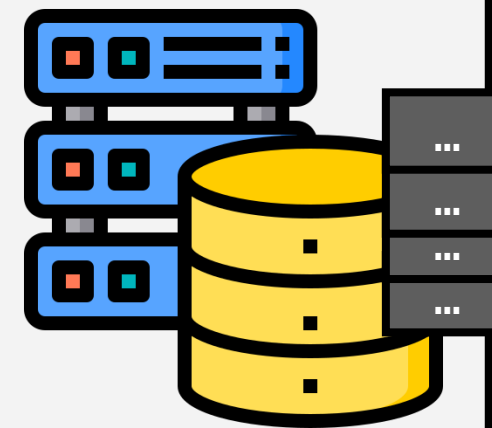
Retrieval Augmented Generation

KB (Knowledge Base)

Homework can be up to three days late with a 10% penalty per day.
Assignments submitted beyond three days receive no credit.
The midterm is closed-book; one page of notes is allowed.
The final exam is closed-book; a single formula sheet may be used.
Quizzes are closed-book and closed-notes.
Quizzes count for ten percent of the final grade.
Attendance is tracked; more than two unexcused absences lowers your participation score.
Active discussion in class contributes to participation.
Office hours are Tuesdays and Thursdays from two to four in room 201.
The TA holds office hours on Wednesdays three to five on Zoom.
You may book a one-on-one appointment for extra help.
Lecture slides are posted weekly on the course page.
Readings are listed on the schedule and updated each week.
Use the discussion forum for questions and peer support.
Slides and readings are available online.
Homework up to three days late loses ten percent per day.
Work one day late loses ten percent; two days late loses twenty.
Submitting beyond three days late earns zero points.
One formula sheet is permitted on the midterm.
Participation is worth ten percent overall.
Office hours are available twice a week.

0.4
0.7
0.1
...

0.85



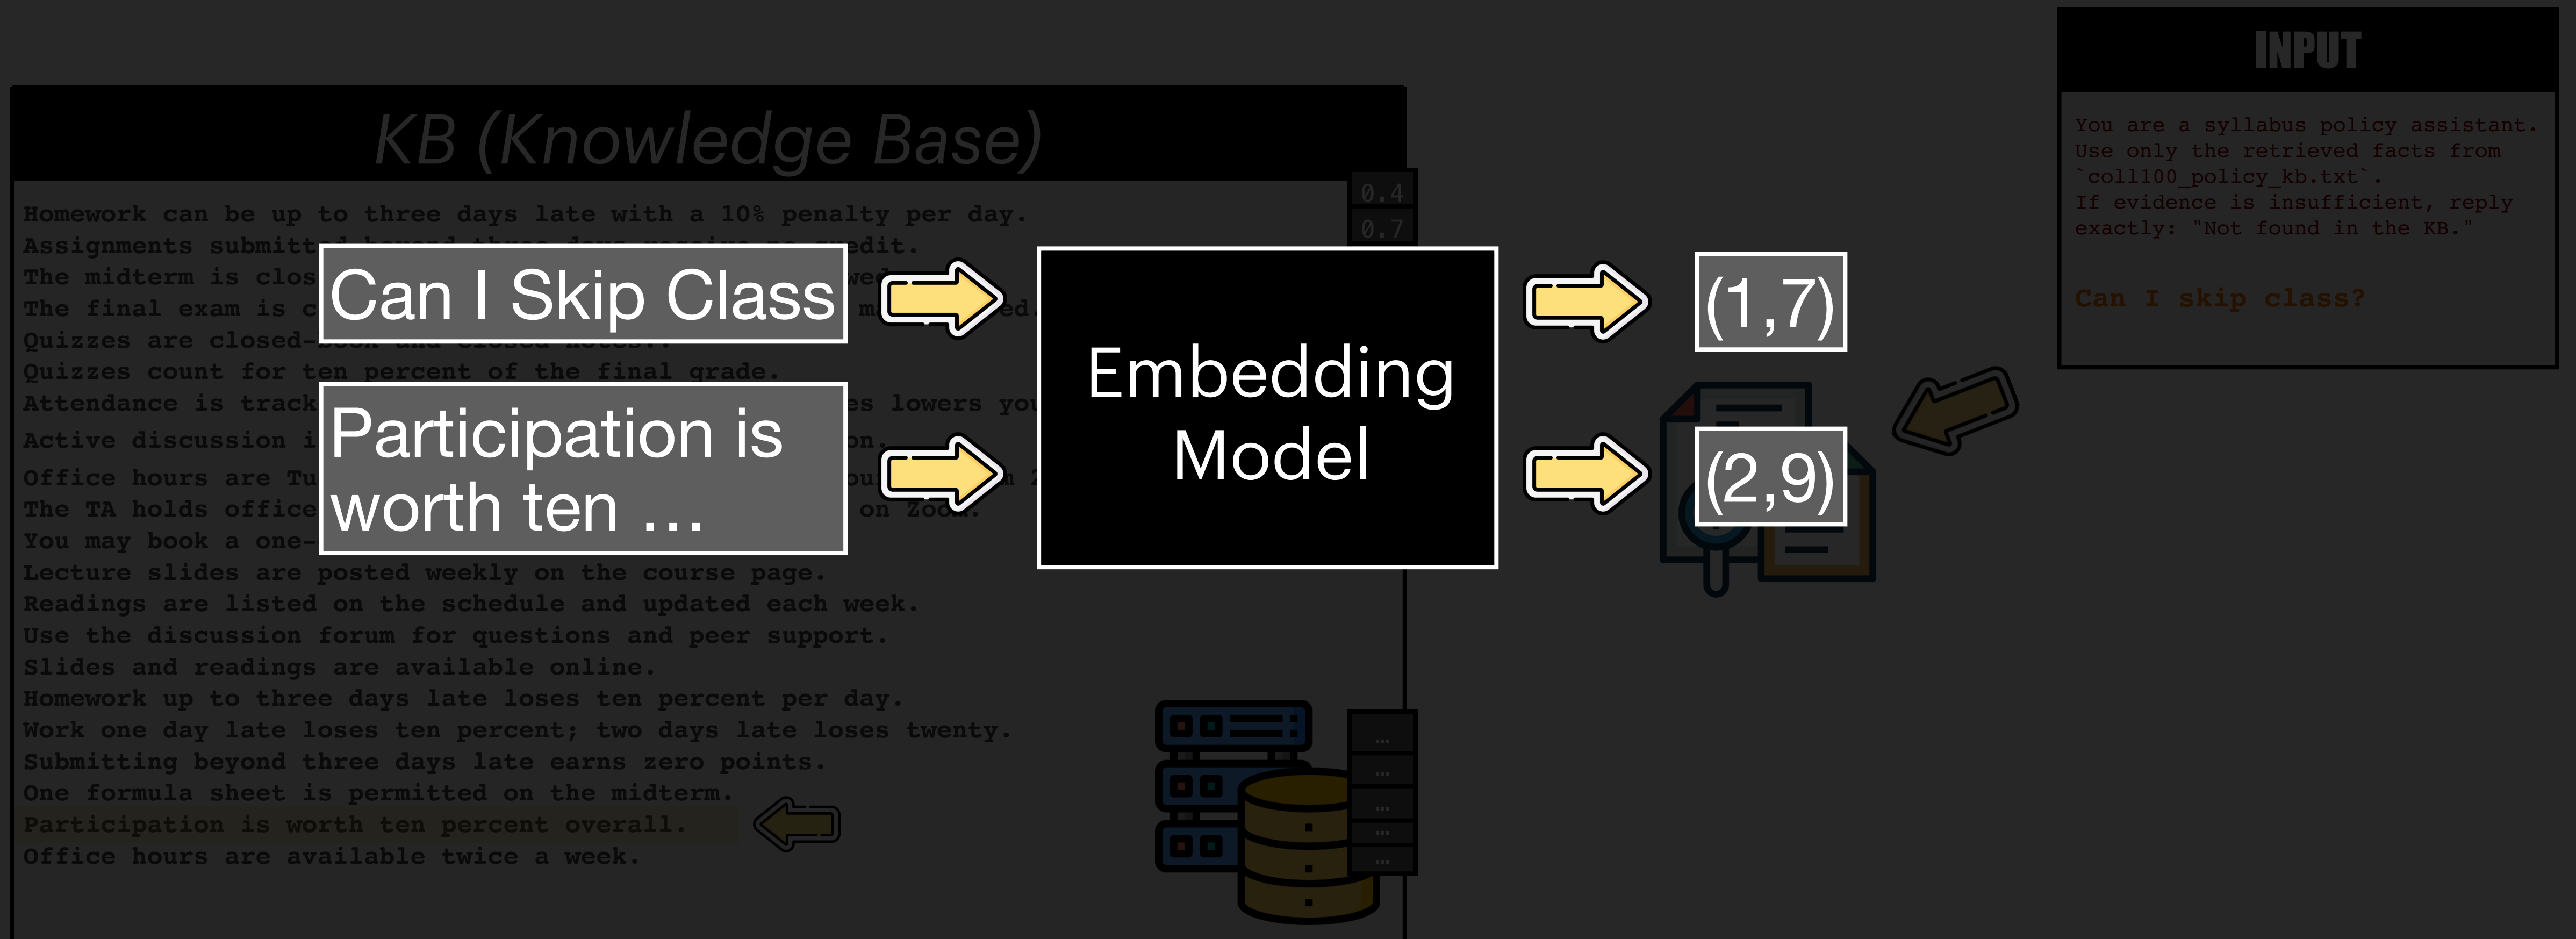
INPUT

You are a syllabus policy assistant.
Use only the retrieved facts from
`coll100_policy_kb.txt`.
If evidence is insufficient, reply
exactly: "Not found in the KB."

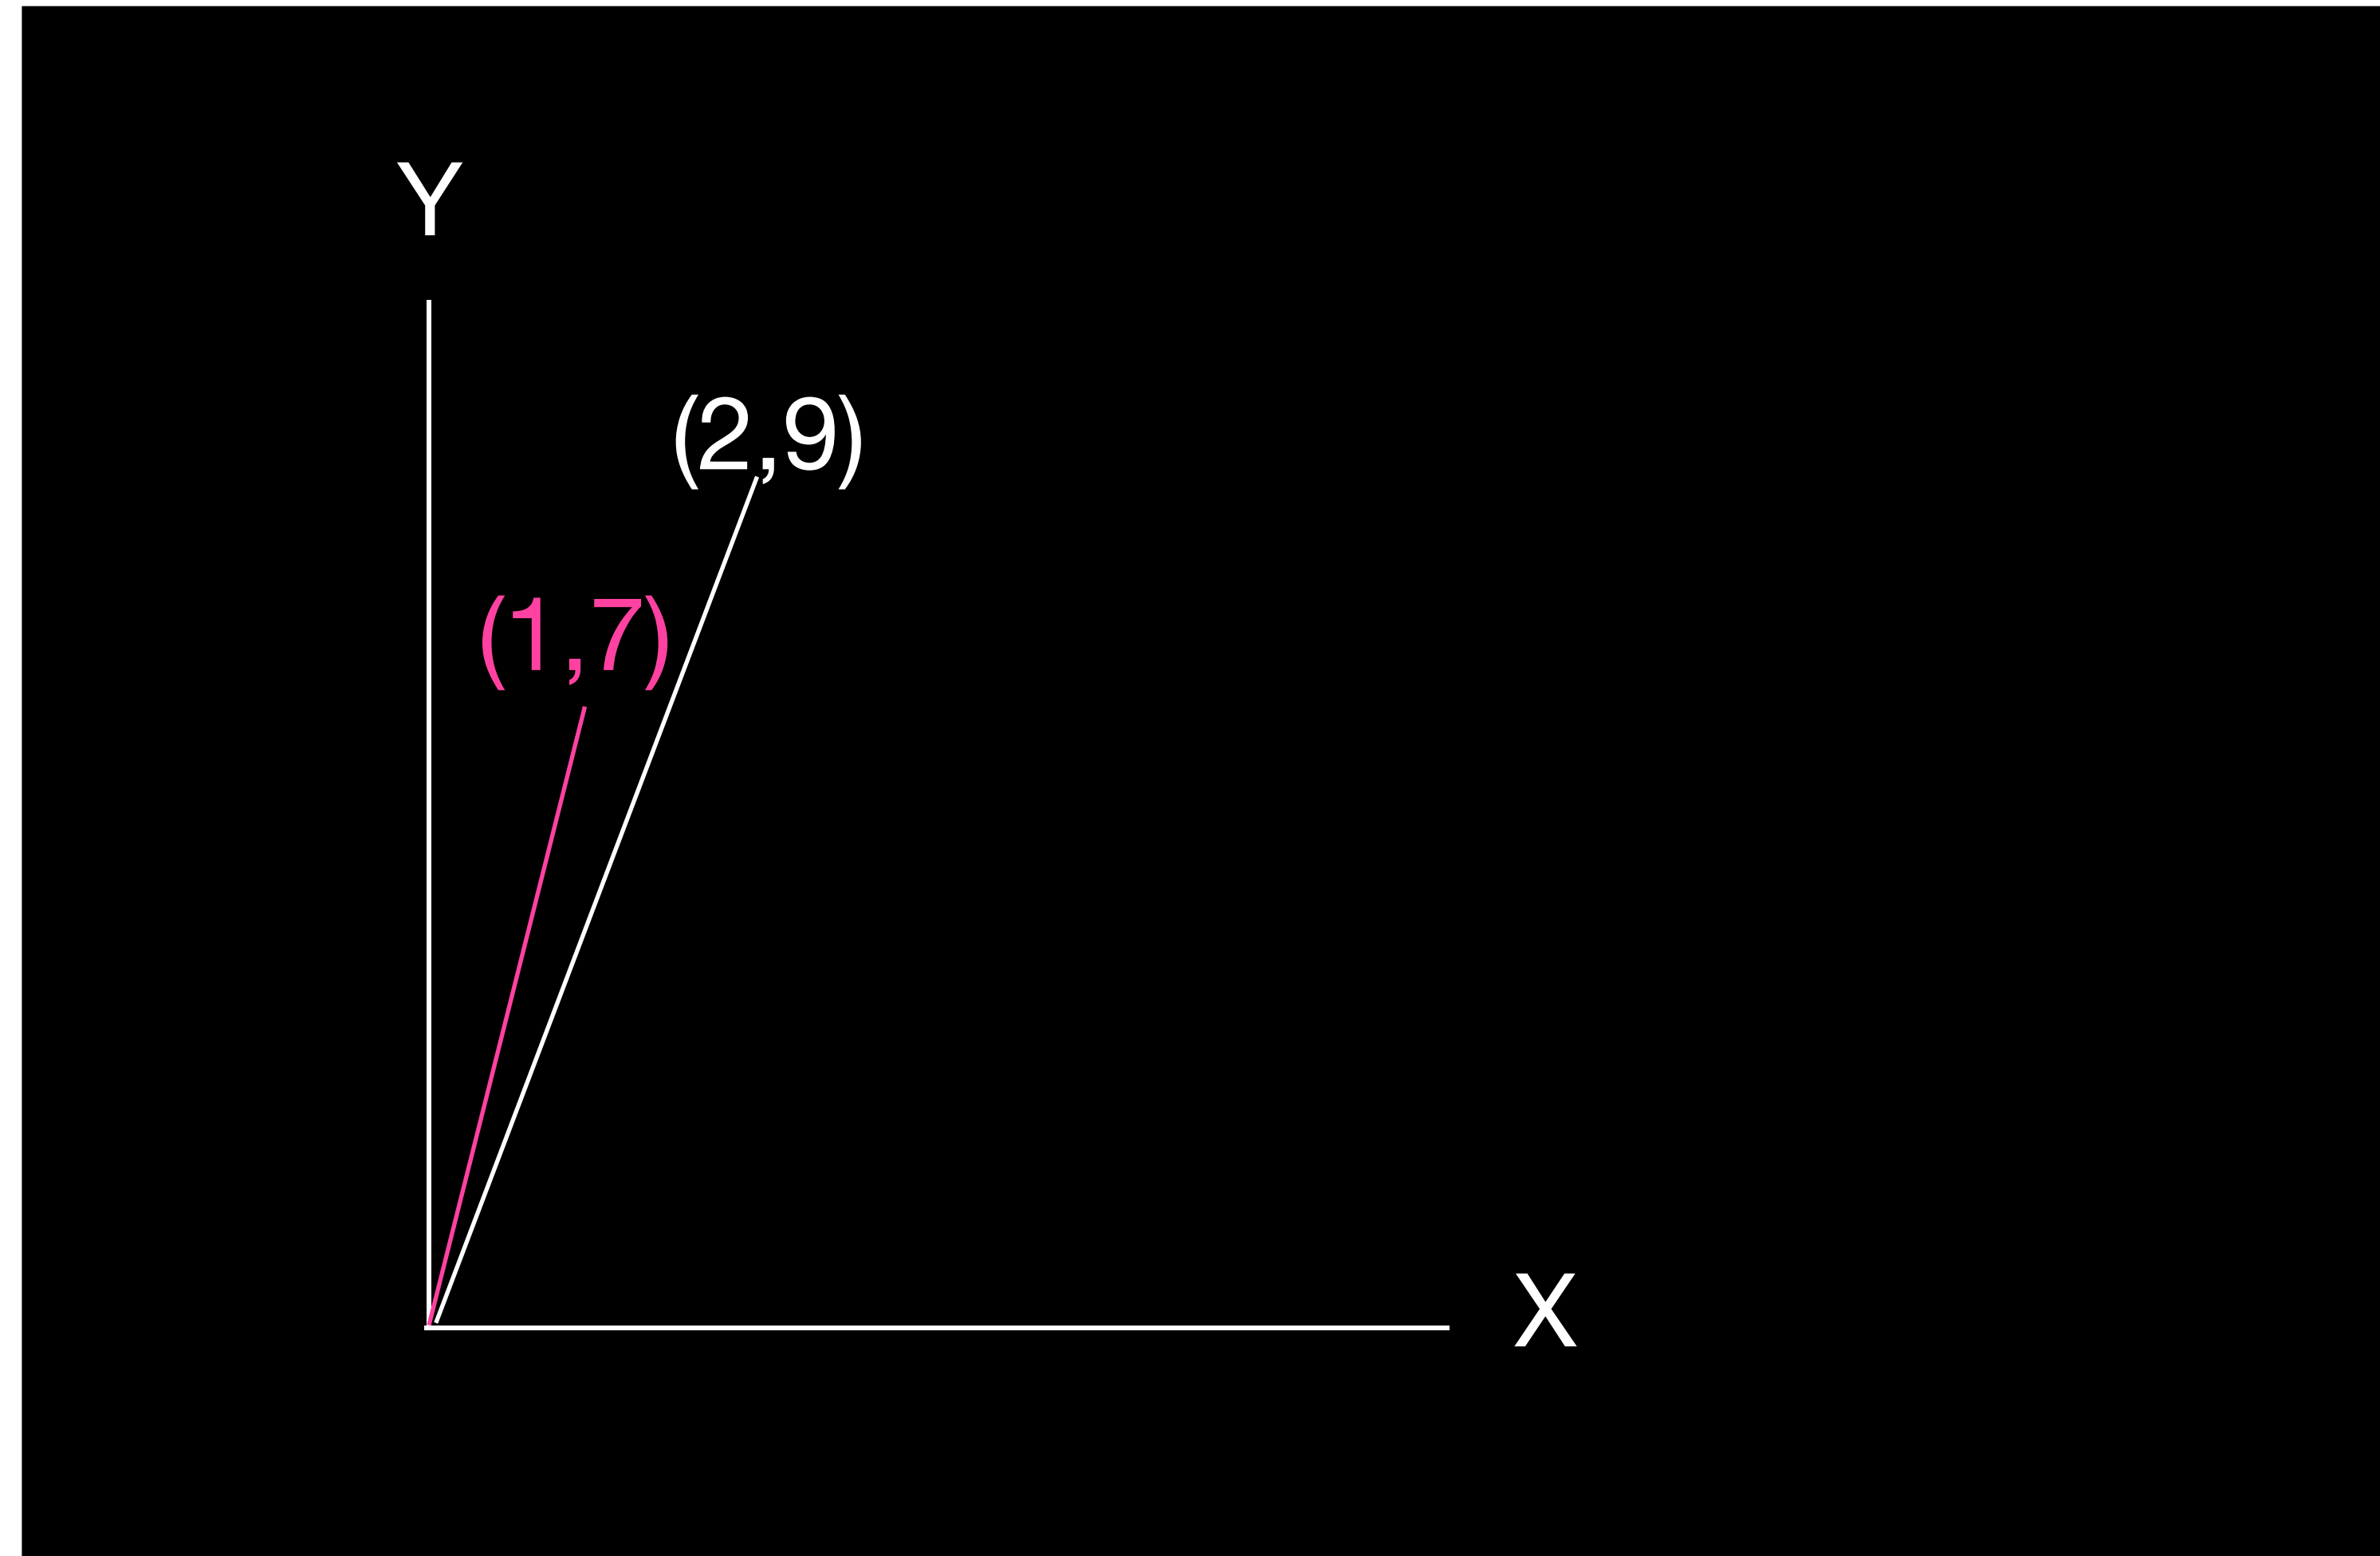
Can I skip class in
your opinion?



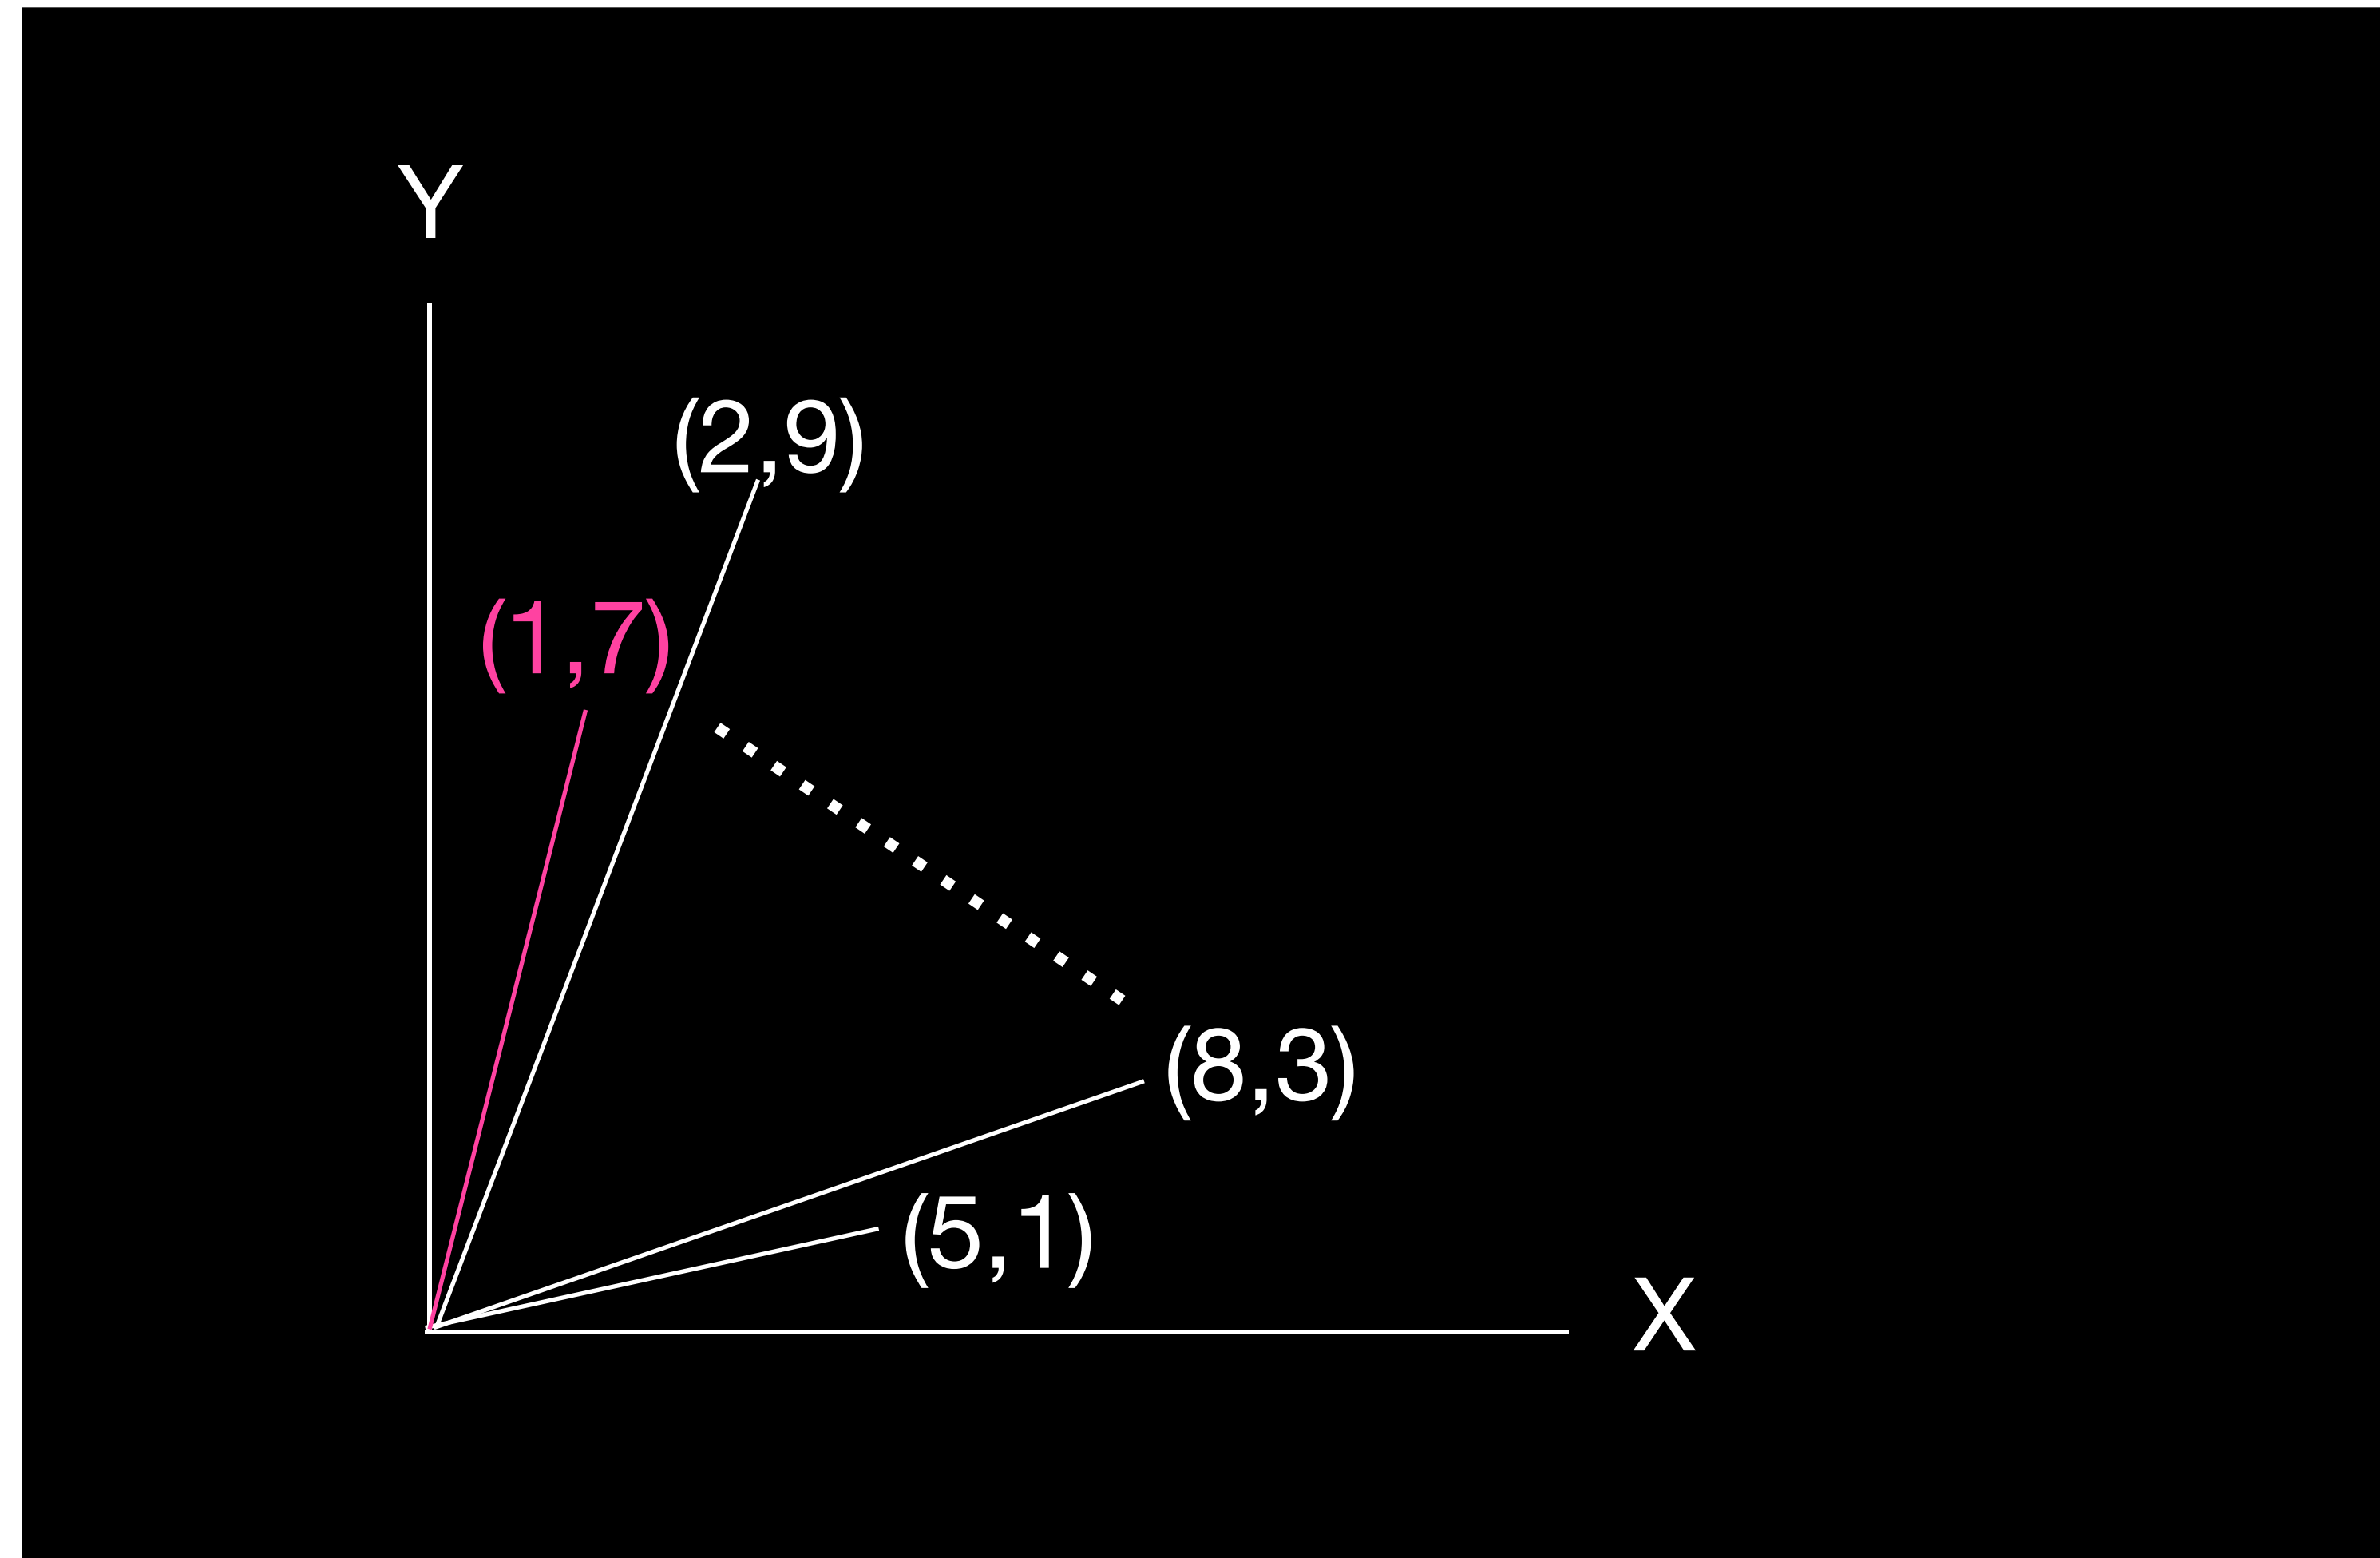
Retrieval Augmented Generation



Retrieval Augmented Generation



Retrieval Augmented Generation



Retrieval Augmented Generation

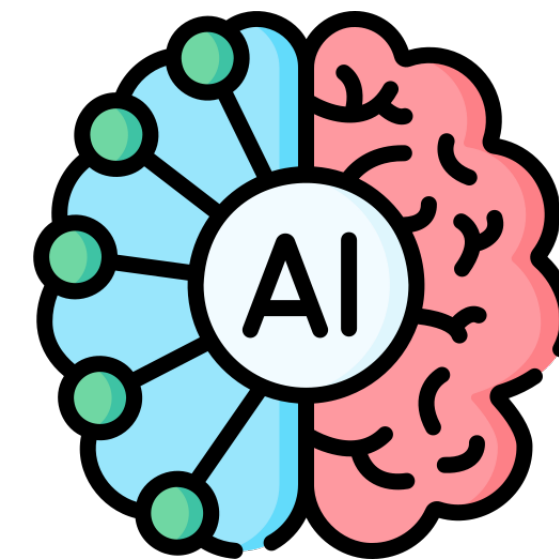
INPUT

You are a syllabus policy assistant. Use only the retrieved facts from `coll100_policy_kb.txt`.
If evidence is insufficient, reply exactly: "Not found in the KB."
Can I skip class in your opinion?

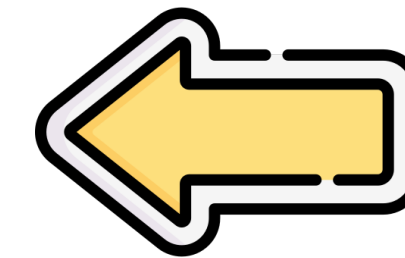
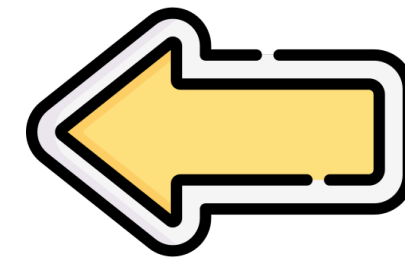
Concat

Participation is worth ten percent overall.

Closest Ex.



GENERATOR



Retrieval Augmented Generation

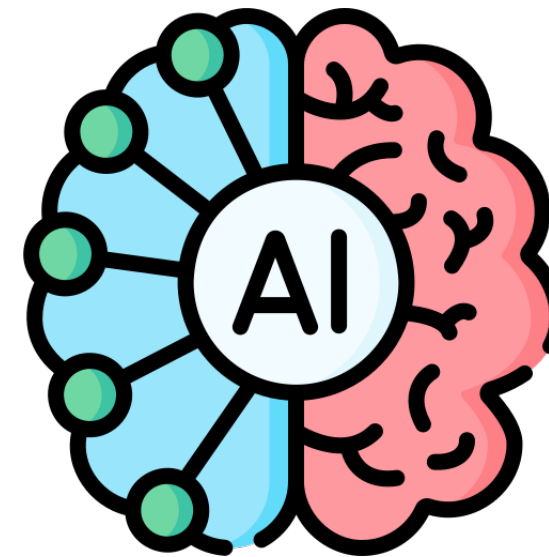
INPUT

You are a syllabus policy assistant. Use only the retrieved facts from `coll100_policy_kb.txt`.
If evidence is insufficient, reply exactly: "Not found in the KB."
Can I skip class in your opinion?

Concat

Participation is worth ten percent overall.

Closest Ex.



GENERATOR

You better not!
COLL:100 Prompt Engineering
cannot be skipped!



antoniomastropaolo.com



[aura-se-lab.github.io](https://github.com/aura-se-lab)



How to avoid hallucinated citations with RAG systems

How to avoid hallucinated citations with RAG systems

Find scientific papers researching system prompts for code-related applications



How to avoid hallucinated citations with RAG systems

Find scientific papers researching system prompts for code-related applications

Okay, this is a great question, and it's a growing area of research! System prompts, used strategically, are showing a lot of promise in software engineering. Here's a breakdown of scientific papers that have systematically explored system prompts for various software engineering tasks, categorized by task and with increasing levels of complexity. I'm including links where available and brief explanations of their approach. **Please read the "Caveats" section at the bottom – the field is evolving very rapidly.**

I. Code Generation & Completion

- **"Large Language Models Encode Human Preferences for Code" (Chen et al., 2023):** This is a foundational paper. They used system prompts to guide code generation based on different "personas" (e.g., "experienced developer," "junior developer") and observed how these personas affected code quality and style. They systematically varied the system prompt to explore these influences. <https://arxiv.org/abs/2309.04008>
- **"Program Synthesis with Large Language Models" (Austin et al., 2023):** Explores using system prompts to guide the generation of code snippets from natural language descriptions, with a focus on program synthesis. <https://arxiv.org/abs/2305.16217>
- **"Refining Code Generation with Large Language Models through System Prompt Engineering" (Wang et al., 2023):** Focuses directly on system prompt engineering to improve the quality and relevance of generated code. They experiment with different prompt templates and content. (This one is a bit older and may use earlier LLMs, but the principles still apply.) https://www.researchgate.net/publication/369712093/Refining_Code_Generation_with_Large_Language_Models_through_System_Prompt_Engineering



How to avoid hallucinated citations with RAG systems

Find scientific papers researching system prompts for code-related applications

Okay, this is a great question, and it's a growing area of research! System prompts, used strategically, are showing a lot of promise in software engineering. Here's a breakdown of scientific papers that have systematically explored system prompts for various software engineering tasks, categorized by task and with increasing levels of complexity. I'm including links where available and brief explanations of their approach. **Please read the "Caveats" section at the bottom – the field is evolving very rapidly.**

I. Code Generation & Completion

- **"Large Language Models Encode Human Preferences for Code" (Chen et al., 2023):** This is a foundational paper. They used system prompts to guide code generation based on different "personas" (e.g., "experienced developer," "junior developer") and observed how these personas affected code quality and style. They systematically varied the system prompt to explore these influences. <https://arxiv.org/abs/2309.04008>
- **"Program Synthesis with Large Language Models" (Austin et al., 2023):** Explores using system prompts to guide the generation of code snippets from natural language descriptions, with a focus on program synthesis. <https://arxiv.org/abs/2305.16217>
- **"Refining Code Generation with Large Language Models through System Prompt Engineering" (Wang et al., 2023):** Focuses directly on system prompt engineering to improve the quality and relevance of generated code. They experiment with different prompt templates and content. (This one is a bit older and may use earlier LLMs, but the principles still apply.) https://www.researchgate.net/publication/369712093/Refining_Code_Generation_with_Large_Language_Models_through_System_Prompt_Engineering



Beyond Reward: Offline Preference-guided Policy Optimization

Yachen Kang^{1,2} Diyuan Shi² Jinxin Liu² Li He² Donglin Wang²

Abstract

This study focuses on the topic of offline preference-based reinforcement learning (PbRL), a variant of conventional reinforcement learning that dispenses with the need for online interaction or specification of reward functions. Instead, the agent is provided with fixed offline trajectories and human preferences between pairs of trajectories to extract the dynamics and task information, respectively. Since the dynamics and task information are orthogonal, a naive approach would involve using preference-based reward learning followed by an off-the-shelf offline RL algorithm. However, this requires the separate learning of a scalar reward function, which is assumed to be an information bottleneck of the learning process. To address this issue, we propose the offline preference-guided policy optimization (OPPO) paradigm, which models offline trajectories and preferences in a one-step process, eliminating the need for separately learning a reward function. OPPO achieves this by introducing an offline hindsight information matching objective for optimizing a contextual policy and a preference modeling objective for finding the optimal context. OPPO further integrates a well-performing decision policy by optimizing the two objectives iteratively. Our empirical results demonstrate that OPPO effectively models offline preferences and outperforms prior competing baselines, including offline RL algorithms performed over either true or pseudo reward function specifications. Our code is available on the project website: <https://sites.google.com/view/oppo-icml-2023>.

¹College of Computer Science and Technology, Zhejiang University, Hangzhou, Zhejiang, China ²Machine Intelligence Lab (MiLAB) of the School of Engineering, Westlake University, Hangzhou, Zhejiang, China. Correspondence to: Donglin Wang <wangdonglin@westlake.edu.cn>, Yachen Kang <kangyachen@westlake.edu.cn>.

Proceedings of the 40th International Conference on Machine Learning, Honolulu, Hawaii, USA. PMLR 202, 2023. Copyright 2023 by the author(s).

1. Introduction

Deep reinforcement learning (RL) offers a versatile framework for acquiring task-oriented behaviors, as evidenced by a growing body of literature (Kohl & Stone, 2004; Kober & Peters, 2008; Kober et al., 2013; Silver et al., 2017; Kalashnikov et al., 2018; Vinyals et al., 2019). In this framework, the "task" is frequently expressed as maximizing the cumulative reward of trajectories produced by deploying the learning policy in the corresponding environment. However, the above RL formulation presupposes two critical conditions for decision policy training: 1) an interactable environment, and 2) a pre-specified reward function. Regrettably, online interactions with the environment can be both expensive and hazardous (Mihatsch & Neuneier, 2002; Hans et al., 2008; Garcia & Fernández, 2015), while developing a suitable reward function typically necessitates considerable human effort. Additionally, the heuristic rewards often employed may be insufficient to express the true intent (Hadfield-Menell et al., 2017).

To address these challenges, prior research has explored two approaches. First, some works have focused on the offline RL formulation (Fujimoto et al., 2019), where the learner has access to fixed offline trajectories along with a reward signal for each transition (or limited expert demonstrations). Second, others have considered the (online) preference-based RL formulation, where the task objective is conveyed to the learner through preferences of a human annotator between two trajectories rather than rewards for each transition. In pursuit of further advancements in this setting, we propose a novel approach that relaxes both of these requirements and advocates for offline preference-based RL (PbRL).

In the context of offline preference-based reinforcement learning (PbRL), where access to an offline dataset and labeled preferences between the offline trajectories is available, a common approach is to combine previous online PbRL methods with off-the-shelf offline RL algorithms (Shin & Brown, 2021). This two-step strategy, as illustrated in Fig. 1 (left), typically involves training a reward function using the Bradley-Terry model (Bradley & Terry, 1952) in a supervised manner, followed by training the policy with any offline RL algorithm on the transitions relabeled via the learned reward function. However, the

How to avoid hallucinated citations with RAG systems

Find scientific papers researching system prompts for code-related applications

Okay, this is a great question, and it's a growing area of research! System prompts, used strategically, are showing a lot of promise in software engineering. Here's a breakdown of scientific papers that have systematically explored system prompts for various software engineering tasks, categorized by task and with increasing levels of complexity. I'm including links where available and brief explanations of their approach. **Please read the "Caveats" section at the bottom – the field is evolving very rapidly.**

I. Code Generation & Completion

- **"Large Language Models Encode Human Preferences for Code" (Chen et al., 2023):** This is a foundational paper. They used system prompts to guide code generation based on different "personas" (e.g., "experienced developer," "junior developer") and observed how these personas affected code quality and style. They systematically varied the system prompt to explore these influences. <https://arxiv.org/abs/2309.04008>
- **"Program Synthesis with Large Language Models" (Austin et al., 2023):** Explores using system prompts to guide the generation of code snippets from natural language descriptions, with a focus on program synthesis. <https://arxiv.org/abs/2305.16217>
- **"Refining Code Generation with Large Language Models through System Prompt Engineering" (Wang et al., 2023):** Focuses directly on system prompt engineering to improve the quality and relevance of generated code. They experiment with different prompt templates and content. (This one is a bit older and may use earlier LLMs, but the principles still apply.) https://www.researchgate.net/publication/369712093/Refining_Code_Generation_with_Large_Language_Models_through_System_Prompt_Engineering



Evaluating and aligning codellms on human preference

[J Yang, J Yang, K Jin, Y Miao, L Zhang, L Yang...](#) - arXiv preprint arXiv ..., 2024 - arxiv.org
... nearly 40 **large language models** (LLMs) using CodeArena reveals significant performance differences between **code**-execution-based benchmarks and our **humancurated** benchmark. ...
☆ Save 📄 Cite Cited by 19 Related articles All 3 versions 🔗

[PDF] arxiv.org

A survey on large language models for code generation

[J Jiang, F Wang, J Shen, S Kim, S Kim](#) - arXiv preprint arXiv:2406.00515, 2024 - arxiv.org
... Transformer introduces an absolute positional **encoding** method to supplement this positional ... the use of reinforcement learning or **human preference modeling** [328]. The QuRating ...
☆ Save 📄 Cite Cited by 494 Related articles All 5 versions 🔗

[PDF] arxiv.org

A systematic evaluation of large language models of code

[FF Xu, U Alon, G Neubig, VJ Helleendoorn](#) - Proceedings of the 6th ACM ..., 2022 - dl.acm.org
... and evaluation **settings** between PolyCoder, open-source **models**, and Codex. Second, we ... (Figure 2, right) An encoderdecoder **model** first uses an encoder to **encode** an input se...
☆ Save 📄 Cite Cited by 901 Related articles All 6 versions 🔗

[PDF] acm.org

Aligning large language models with human: A survey

[Y Wang, W Zhong, L Li, F Mi, X Zeng, W Huang...](#) - arXiv preprint arXiv ..., 2023 - arxiv.org
... of SFT into: Online **human preference** training, Offline **human preference** training and ... **code** generated by chatgpt really correct? rigorous evaluation of **large language models** for **code** ...
☆ Save 📄 Cite Cited by 442 Related articles All 3 versions 🔗

[PDF] arxiv.org

Do LLMs Understand Code Preference? Training Code Preference Models via Synthetic Code Evolution

[J Liu, TV Nguyen, M Shang, H Ding, X Li, Y Yu...](#) - ... Deep Learning for **Code** - openreview.net
... Efficient memory management for **large language model** serving with pagedattention. In Proceedings of the 29th Symposium on Operating Systems Principles, pp. 611–626, 2023. ...
☆ Save 📄 Cite Related articles 🔗

[PDF] openreview.net

A survey on human preference learning for large language models

[R Jiang, K Chen, X Bai, Z He, J Li, M Yang...](#) - arXiv preprint arXiv ..., 2024 - arxiv.org
... **human preference** learning for LLMs from a **preference**-centered perspective, covering the sources and formats of **preference** feedback, the **modeling** and usage of **preference** ... **human** ...
☆ Save 📄 Cite Cited by 34 Related articles All 3 versions 🔗

[PDF] arxiv.org

Do people prefer "natural" code?

[C Casalnuovo, K Lee, H Wang, P Devanbu...](#) - arXiv preprint arXiv ..., 2019 - arxiv.org
... that this **preference** for more familiar forms is manifest in **large** ... Using a **language model** trained on Java **code**, if we perform ... We test if the surprisal scores predict **human preference**. ...
☆ Save 📄 Cite Cited by 8 Related articles All 5 versions 🔗

[PDF] arxiv.org

Evaluating large language models trained on code

[M Chen, J Tworek, H Jun, Q Yuan, HPDO Pinto...](#) - arXiv preprint arXiv ..., 2021 - arxiv.org
... The **largest** source of inefficiency arises from **encoding** ... Inspired by similar work in **language modeling**, we find that ... (k) for various temperature **settings**. Higher temperatures are ...
☆ Save 📄 Cite Cited by 6313 Related articles All 2 versions 🔗

[PDF] arxiv.org

Learning code preference via synthetic evolution

[J Liu, T Nguyen, M Shang, H Ding, X Li, Y Yu...](#) - arXiv preprint arXiv ..., 2024 - arxiv.org
... Efficient memory management for **large language model** serving with pagedattention. In Proceedings of the 29th Symposium on Operating Systems Principles, pp. 611–626, 2023. ...
☆ Save 📄 Cite Cited by 12 Related articles All 4 versions 🔗

[PDF] arxiv.org

Large language model-aware in-context learning for code generation

[J Li, C Tao, J Li, G Li, Z Jin, H Zhang...](#) - ACM Transactions on ..., 2025 - dl.acm.org
... (referred to as **preferences**) of LLMs. In this section, we propose a **Large language model-Aware** ... Next, We apply the retriever to **encode** the requirements of these examples and acquire ...
☆ Save 📄 Cite Cited by 48 Related articles All 4 versions 🔗

[PDF] acm.org



How to avoid hallucinated citations with RAG systems

Counterexample to the Néron–Ogg–Shafarevich criterion for Calabi–Yau threefolds

Tymoteusz Chmiel and Marcin Oczko

ABSTRACT

For any prime $p > 5$ we construct a Calabi-Yau threefold X defined over a finite extension K of \mathbb{Q}_p such that every model of X over $\text{Spec } \mathcal{O}_K$ has singular special fiber, yet the Galois action on the ℓ -adic cohomology group $H_{\text{ét}}^3(X, \mathbb{Q}_\ell)$ is unramified for $\ell \neq p$ and crystalline for $\ell = p$. This provides a counterexample to the analogue of the Néron-Ogg-Shafarevich criterion in dimension three.

1. Introduction

A *Calabi-Yau manifold* is a smooth projective variety X with trivial canonical bundle and vanishing cohomology groups $H^i(X, \mathcal{O}_X)$ for $0 < i < \dim X$. Usually one considers Calabi-Yau manifolds defined over \mathbb{C} . On the other hand, in this paper we are concerned with arithmetic properties of Calabi-Yau manifolds and the field of definition is either p -adic or finite.

One-dimensional Calabi-Yau manifolds are *elliptic curves* and two-dimensional ones are called *K3 surfaces*. This paper is concerned with *Calabi-Yau threefolds* and a phenomenon that does not appear in lower dimensions: the étale cohomology fails to detect bad reduction. Thus the natural generalization of the Néron-Ogg-Shafarevich criterion does not hold in dimension 3.

Néron-Ogg-Shafarevich criterion states that an elliptic curve E over a p -adic field has good reduction if and only if the Galois representation $H_{\text{ét}}^1(E, \mathbb{Q}_\ell)$ is unramified for $\ell \neq p$ and crystalline for $\ell = p$ (see [Ogg67]). A version of this criterion holds for K3 surfaces: K3 surface S admitting a semi-stable model has potentially good reduction if and only if the Galois representation $H_{\text{ét}}^2(S, \mathbb{Q}_\ell)$ is unramified for $\ell \neq p$ and crystalline for $\ell = p$ (see [LM18, CLL17]).

It means that for Calabi-Yau varieties of dimension ≤ 2 good reduction is detected by the Galois action on the middle cohomology. We show it is not the case in dimension three:

MAIN THEOREM. *Let $p > 5$ be a prime, $K := \mathbb{Q}_p[\sqrt{p}]$ and denote by G_K the absolute Galois group of K . There exists smooth Calabi-Yau threefold Y_p defined over K such that:*

- the Galois representation

$$G_K \rightarrow \text{Aut}(H_{\text{ét}}^3(Y_p, \mathbb{Q}_\ell))$$

2020 *Mathematics Subject Classification* 11G25, 14F30, 14D06, 14J32

Keywords: Calabi-Yau threefolds, bad reduction, crystalline cohomology.

The first author was supported by the National Science Center grant no. 2023/49/N/ST1/04089. The second named author was partially supported by National Science Centre, Poland, grant No. 2020/39/B/ST1/03358.

Okay, this
strategica
papers th
categoriz
brief expl
is evolvin

I. Code G

• "Larg

a fou

"pers

perso

expl

• "Pro

sys

with

• "Refi

Engi

the q

and c

https

e_La

papers researching
s for code-related

scientific

asks,

ble and

e field

This is

ferent

t to

ing

ptions,

improve

nplates

ply.)

h_Larg



antoniomastropaolo.com



aura-se-lab.github.io



How to avoid hallucinated citations with RAG systems

Find scientific papers researching system prompts for code-related applications

Okay, this is a great question, and it's a growing area of research! System prompts, used strategically, are showing a lot of promise in software engineering. Here's a breakdown of scientific papers that have systematically explored system prompts for various software engineering tasks, categorized by task and with increasing levels of complexity. I'm including links where available and brief explanations of their approach. **Please read the "Caveats" section at the bottom – the field is evolving very rapidly.**

I. Code Generation & Completion

- **"Large Language Models Encode Human Preferences for Code" (Chen et al., 2023):** This is a foundational paper. They used system prompts to guide code generation based on different "personas" (e.g., "experienced developer," "junior developer") and observed how these personas affected code quality and style. They systematically varied the system prompt to explore these influences. <https://arxiv.org/abs/2309.04008>
- **"Program Synthesis with Large Language Models" (Austin et al., 2023):** Explores using system prompts to guide the generation of code snippets from natural language descriptions, with a focus on program synthesis. <https://arxiv.org/abs/2305.16217>
- **"Refining Code Generation with Large Language Models through System Prompt Engineering" (Wang et al., 2023):** Focuses directly on system prompt engineering to improve the quality and relevance of generated code. They experiment with different prompt templates and content. (This one is a bit older and may use earlier LLMs, but the principles still apply.) https://www.researchgate.net/publication/369712093/Refining_Code_Generation_with_Large_Language_Models_through_System_Prompt_Engineering



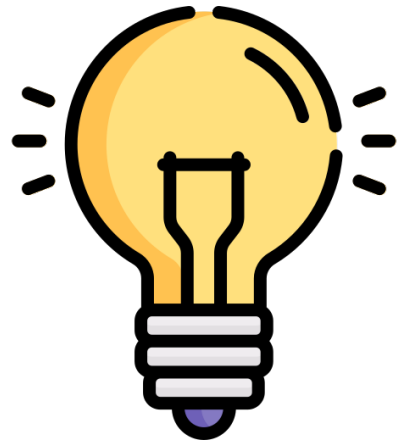
How to avoid hallucinated citations with RAG systems



Rather than asking AI to generate citations coming from its internal knowledge, retrieval-augmented generation (RAG) systems first search a verified database, then use those real sources as the basis for responses. This grounds AI outputs in factual information rather than statistical patterns.



How to avoid hallucinated citations with RAG systems



Rather than asking AI to generate citations coming from its internal knowledge, retrieval-augmented generation (RAG) systems first search a verified database, then use those real sources as the basis for responses. This grounds AI outputs in factual information rather than statistical patterns.

You are an **academic assistant** that answers **only** using information from verified source documents.

Instructions:

- 1. You will receive a user question and a set of documents or passages retrieved from a trusted database.
- 2. Before answering, identify the 2–3 most relevant excerpts from the sources.
- 3. Then, generate a clear, fact-based answer that uses only those excerpts (i.e., summary).
- 4. Cite each claim with the source title and page number (if available) in the format [Title, p.#].
- 5. If the answer is not found in the sources, reply exactly: “Not found in the provided materials.”

Never use your internal knowledge. Do not fabricate facts, estimates, or citations. Your role is to ground every response in the documents provided.



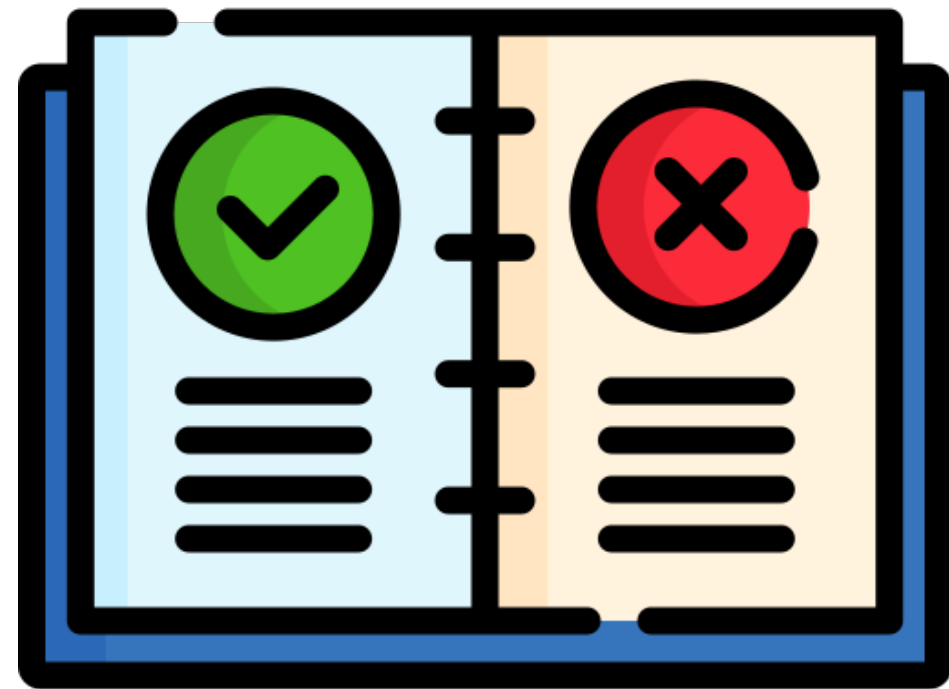
antoniomastropaolo.com



[aura-se-lab.github.io](https://github.com/aura-se-lab)



How to avoid hallucinated citations with RAG systems



How to recognize when an AI has invented a citation:

- Suspiciously generic author names — think “Smith” or “Brown.”
- Nonsense article titles — often a random jumble of words copied from your own prompt.
- Broken identifiers — the DOI doesn’t resolve, or the link goes to an unrelated (but real) article.
- Borrowed credibility — fake citations may still drop the name of a real author or a legitimate journal, even though the specific paper doesn’t actually exist.



How to avoid hallucinated citations with RAG systems



By [Benjamin Weiser](#)

May 27, 2023

The lawsuit began like so many others: A man named Roberto Mata sued the airline Avianca, saying he was injured when a metal serving cart struck his knee during a flight to Kennedy International Airport in New York.

When Avianca asked a Manhattan federal judge to toss out the case, Mr. Mata’s lawyers vehemently objected, submitting a 10-page brief that cited more than half a dozen relevant court decisions. There was *Martinez v. Delta Air Lines*, *Zicherman v. Korean Air Lines* and, of course, *Varghese v. China Southern Airlines*, with its learned discussion of federal law and “the tolling effect of the automatic stay on a statute of limitations.”

There was just one hitch: No one — not the airline’s lawyers, not even the judge himself — could find the decisions or the quotations cited and summarized in the brief.

That was because ChatGPT had invented everything.



antoniomastropaolo.com

[aura-se-lab.github.io](https://github.com/aura-se-lab)

