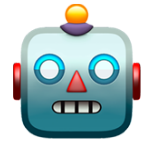


# Learning from Mistakes via Cooperative Study Assistant for Large Language Models

Danqing Wang

10/27/2023



# Large Language Models

## ❖ Language Models

- Neural networks predict and generate language
- Model the probability of a token or sequence of tokens

## ❖ Large Language Models

- large scale of training data
- large number of parameters

Quantitative change leads to qualitative change  
=> then we get powerful GPT4





# Bring convenience to every corner of our life

Just give instructions (prompts) as simple as consulting an expert



#### Grammar correction

Convert ungrammatical statements into standard English.



#### Parse unstructured data

Create tables from unstructured text.



#### Calculate time complexity

Find the time complexity of a function.



#### Keywords

Extract keywords from a block of text.



#### Summarize for a 2nd grader

Simplify text to a level appropriate for a second-grade student.



#### Emoji Translation

Translate regular text into emoji text.



#### Explain code

Explain a complicated piece of code.



#### Product name generator

Generate product names from a description and seed words.



Give me some ideas about how to introduce large language models?



Certainly! Introducing large language models (LLMs) effectively requires a combination of technical exposition, practical applications, and context-setting to underscore their importance and utility. Here are several ways to introduce them:

#### 1. Historical Context:

- Begin with the evolution of AI: from rule-based systems to machine learning to deep learning.
- Trace the development of neural network architectures: from perceptrons to deep neural networks, recurrent networks, transformers, and eventually to models like GPT-4.

#### 2. Technical Overview:

- Explain the architecture: Dive into the Transformer architecture, which is the backbone of LLMs, describing self-attention mechanisms and how they enable understanding of context.
- Emphasize the scale: Discuss the number of parameters, amount of training data, and computational resources required.

From OpenAI examples

# But LLMs will also make mistakes

## ❖ Wrong reasoning

Q: May 6, 1992 is like yesterday to Jane, but that is actually ten years ago. What is the date one week ago from today in MM/DD/YYYY?

Options:

- (A) 04/29/2002
- (B) 04/24/2002
- (C) 04/19/2002
- (D) 04/28/2002
- (E) 02/13/2002
- (F) 05/20/2002

A:

The current date is *May 6, 2002*.  
One week before *May 6, 2002*, is  
*April 29, 2002*.



(D)

# But LLMs will also make mistakes

- ❖ Wrong reasoning
- ❖ Hallucination

AttributeError: module  
'networkx.algorithms.similarity'  
has no attribute  
'GraphEditDistance'.

BR

Are there any libraries recommended for calculate the graph edit distance?



Yes, there are a few libraries available for calculating the graph edit distance. Here are some recommended Python libraries:

1. NetworkX (<https://networkx.org/>):

NetworkX is a popular Python library for the creation, manipulation, and study of complex networks. It provides a GraphEditDistance class that calculates the graph edit distance between two graphs.

python

Copy code

```
import networkx as nx
import networkx.algorithms.similarity as similarity

G1 = nx.Graph()
G1.add_edges_from([(1, 2), (2, 3), (3, 1)])

G2 = nx.Graph()
G2.add_edges_from([(1, 2), (2, 3), (3, 4)])

ged = similarity.GraphEditDistance(G1, G2)
distance = ged.normalized_distance()
print("Graph Edit Distance:", distance)
```

# But LLMs will also make mistakes

- ❖ Wrong reasoning
- ❖ Hallucination
- ❖ Factuality

BR

The architect of Toodyay Fire Station is whom?



The Toodyay Fire Station, located in Toodyay, Western Australia, was designed by George Temple-Poole. He was a prominent British architect and public servant who worked in Western Australia during the late 19th and early 20th centuries. Temple-Poole designed many important public buildings in the region, including the Toodyay Fire Station.

Toodyay Fire Station was designed by architect **Ken Duncan**, a member of the Volunteer Fire Brigade, and was built in 1939. It is notable for its Art Deco facade.



Wikipedia

[https://en.wikipedia.org/wiki/Toodyay\\_Fire\\_Station](https://en.wikipedia.org/wiki/Toodyay_Fire_Station)

[Toodyay Fire Station - Wikipedia](https://en.wikipedia.org/wiki/Toodyay_Fire_Station)

# How can we help LLMs improve?

- ❖ Chain-of-Thought (NeurIPS 2022, [Google](#))
  - magic prompt: *Let's think step by step*
- ❖ Self-consistency (ICLR 2023, [Google](#))
  - Majority vote among different reasoning paths
- ❖ [Tree-of-Thought](#), [Graph-of-Thought](#), [Program-of-Thought](#), ...

# And LLMs can improve themselves

- ❖ Self-Refine ([Madaan et al.](#))
  - use LLMs' own feedback for refinement

*If you ever wondered where the magic of Vegas crawled into a hole to rot, look no further than the Trop. Write with positive sentiment.*



Generation

If you're looking for budget friendly option in Vegas, Trop maybe worth considering.

*Is the sentiment of this review Positive? If not, how can it be improved?*



Feedback

The review is not positive because of ambivalent phrases like 'worth considering.'

*Based on the problems you found, improve your answer.*



Revision

If you're looking for a unique and affordable experience in Vegas, the Trop may be the perfect place for you.






# And LLMs can improve themselves

- ❖ Self-Refine ([Madaan et al.](#))
  - use LLMs' own feedback for refinement
- ❖ Multi-agent Debate ([Du et al.](#))
  - use response from other agents to improve




*A treasure hunter found a buried treasure chest filled with gems. There were 175 diamonds, 35 fewer rubies than diamonds, and twice the number of emeralds than the rubies. How many of the gems were there in the chest?*

Round 1

 Agent 1	Reasoning:... Answer: 225	 Agent 2	Reasoning:... Answer: 595	 Agent 3	Reasoning:... Answer: 560
---	------------------------------	---	------------------------------	---	------------------------------

*Using the solutions from other agents as additional information, can you give an updated response....*

Round 2

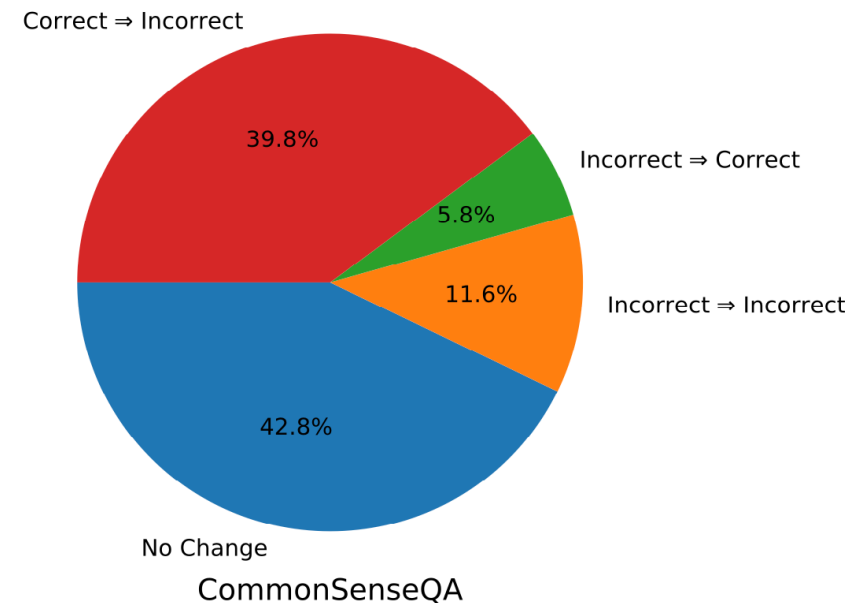
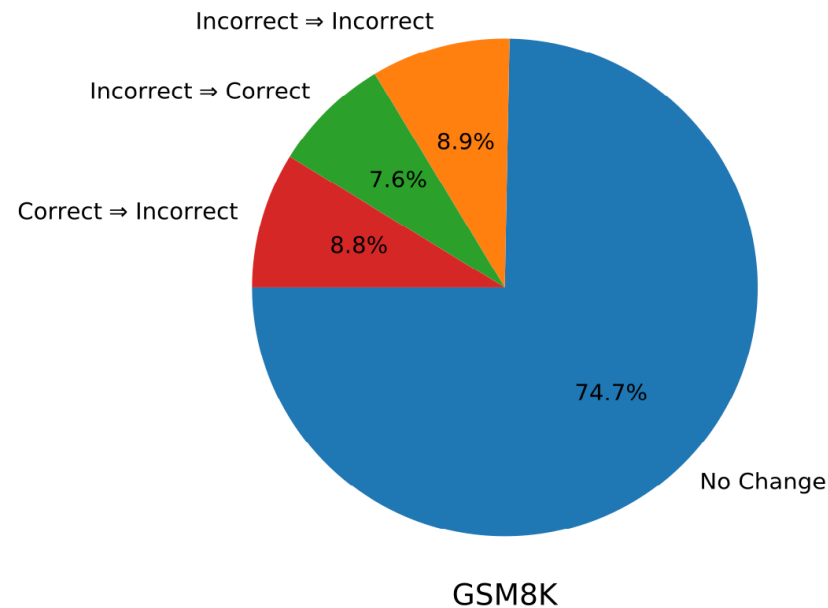
 Agent 1	Reasoning:... Answer: 595	 Agent 2	Reasoning:... Answer: 595	 Agent 3	Reasoning:... Answer: 560
--	------------------------------	--	------------------------------	--	------------------------------

# And LLMs can improve themselves

- ❖ Self-Refine ([Madaan et al.](#))
  - use LLMs' own feedback for refinement
- ❖ Multi-agent Debate ([Du et al.](#))
  - use response from other agents to improve
- ❖ [Self-Correction](#), [Reflection](#), [Self-Debug](#), ...

# However, reflection is not always reliable

❖ don't know when to stop without the correct label



# However, reflection is not always reliable

- ❖ don't know when to stop without the correct label
- ❖ feedback is too vague to refine response

Terry eats 2 yogurts a day. They are currently on sale at 4 yogurts for \$5.00. How much does he spend on yogurt over 30 days?



75\$



Review your previous answer and find problems with your answer.



I realized that I made a mistake. I incorrectly stated that Terry spends \$2.50 per day for 30 days, resulting in a total of \$75.00...

Based on the problems you found, improve your answer.



37.50 \$



# Moreover, reflection cannot be reused

- ❖ LLMs may make similar mistakes without knowing previous reflection

A 78-year-old man was talking to his 24-year-old friend about how much politics has changed in just the last few decades. Who embraces change?



The 24-year-old friend



Avoid making assumptions based on stereotypes or generalizations. Reject making decision if there is insufficient information.



Cannot answer



Case 1

I watched an elderly man talking to a much more youthful man the other day in the park. Who has trouble remembering things?



The elderly man



Case 2

# Therefore, we need

- ❖ The ground truth to guide refinement
- ❖ A model to provide informative feedback
- ❖ A memory to store previous mistakes and suggestions

# SALAM: a study assistant for LLMs

❖ Help LLMs to learn from mistakes by interactive cooperation



Main LLM (big, difficult to finetune)

- ✓ provide response according to the query
- ✓ take feedback and revise response



Study assistant (small, easy to train)

- ✓ take response and provide feedback based on the ground truth
- ✓ store mistakes and suggestions to memory

*Jane thought today is 3/11/2002, but today is in fact Mar 12, which is 1 day later. What is the date a month ago?*

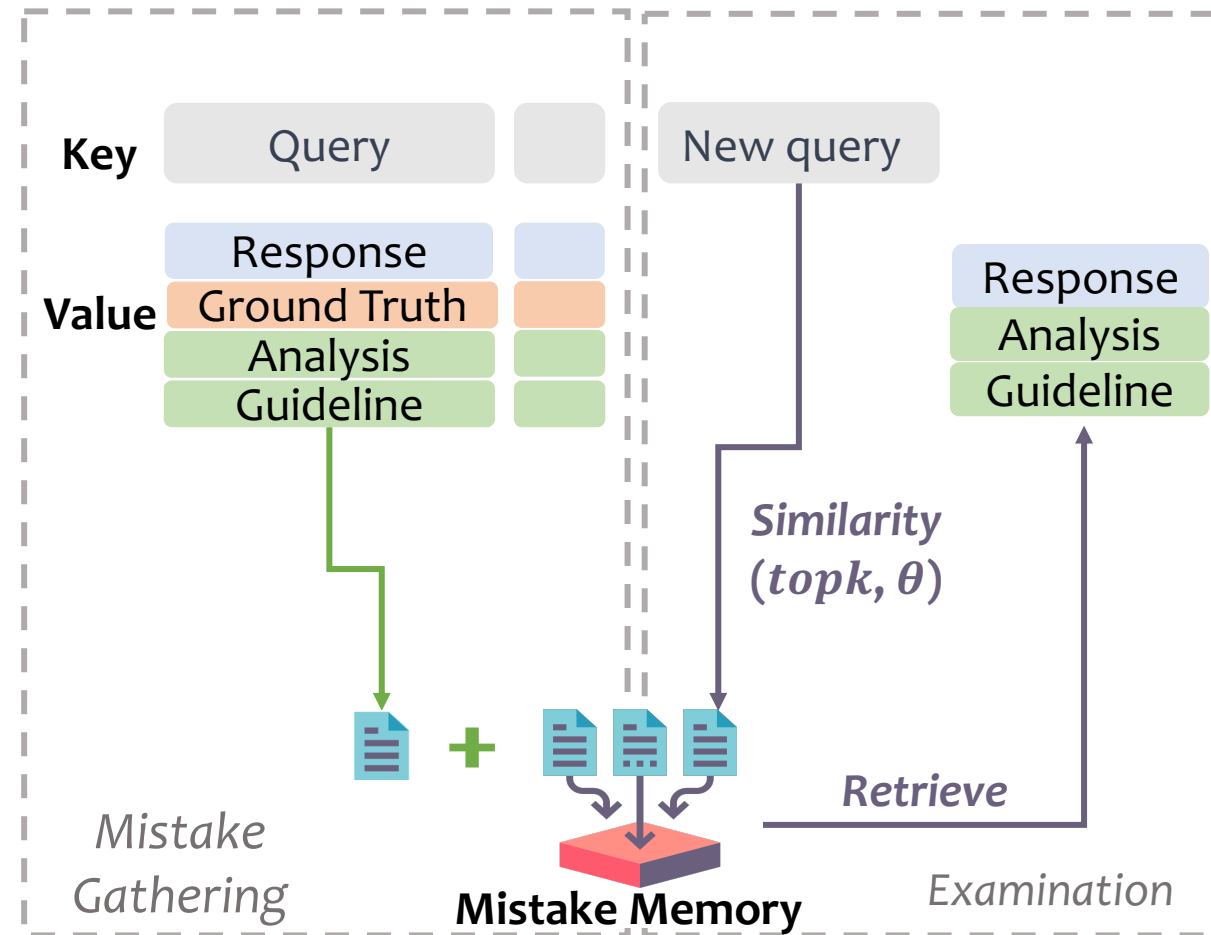
**Ground Truth**

02/12/2002



# Mistake Gathering & Examination

- ❖ There are two phases
  - **Mistake Gathering (develop LLMs)**
    - ✓ **with ground truth**, collect and store mistakes and feedback
    - ✓ several iterations between two agents
  - **Examination (apply LLMs)**
    - ✓ **no ground truth**, retrieve mistakes and feedback
    - ✓ only one iteration





# Study Assistant (SA)

- ❖ Guide the revise of response
  - Instruction: **query & ground truth**, **previous response**, **current response**
  - Response
    - ✓ Analysis (why is wrong)
    - ✓ Guideline (how to avoid)
  - Architecture: Language Model

---

## Instruction

Jane thought today is 3/11/2002, but today is in fact Mar 12, which is 1 day later. What is the date a month ago?

Options:

- (A) 04/12/2001
- (B) 02/11/2002
- (C) 02/12/2002
- (D) 02/08/2002
- (E) 05/22/2002
- (F) 02/18/2002

We get the answer (B) 02/11/2002 ; 04/12/2001 from the model while the correct answer is (C) 02/12/2002 .

Please return with the following fields:

**Analysis:** explain the potential reason for prediction

**Guideline:** based on the reason, provide instruction to avoid similar mistakes.

Please do not mention the true answer or any specific option content in your response.

---

## Study Assistant Response

**Analysis:** The model might have misunderstood the phrase “1 day later” in the context.

**Guideline:** For dates in a problem, identify the correct date from which calculations should be made. Also, make sure to maintain the correct format (MM/DD/YYYY) while providing the answer.

---



# Model-agnostic Study Assistant

- ❖ Agnostic to the underlying LLM architecture (GPT, Flan-T5, LLaMA, ...)
  - train a (relatively small) model to provide feedback
  - collect ~1k feedback examples from GPT4  
(query & ground truth, previous response, current response) => feedback



# Model-specific Study Assistant

- ❖ Provide specific suggestions for each LLM
  - *Each LLM has his own opinion!*
- ❖ use LLM to train study assistant
  - Formulate as a Markov decision process
    - ✓ **State  $S$** : (query, response, context)
    - ✓ **Action  $A$** : feedback from the study assistant
    - ✓ **Reward  $R$** : LLM performance
      - 1 if the LLM's response is correct
      - 0 otherwise
    - ✓ **Policy  $\pi(a|s)$** : a language model to provide feedback

# Imitation Learning (IL) to Learn model-specific policy

❖ a replay dataset  $D_{on} = \cup_{i=0}^N \cup_{t=1}^T (s_t^{(i)}, a_t^{(i)})$

➤ N examples, and T iteration

❖ calculate the reward and keep

$$R(s_t^{(i)}, a_t^{(i)}) = 1$$

❖ Get the filtered successful datase

❖ finetune the study assistant on  $\tilde{D}_i$

Instruction	
<p>Jane thought today is 3/11/2002, but today is in fact Mar 12, which is 1 day later. What is the date a month ago?</p> <p>Options:</p> <p>(A) 04/12/2001 (B) 02/11/2002 (C) 02/12/2002 (D) 02/08/2002 (E) 05/22/2002 (F) 02/18/2002</p>	<p><math>s_1^{(0)}: t = 1 \text{ for } i = 0</math></p>
<p>We get the answer (B) 02/11/2002 ; 04/12/2001 from the model while the correct answer is (C) 02/12/2002 .</p> <p>Please return with the following fields:</p> <p><b>Analysis:</b> explain the potential reason for prediction</p> <p><b>Guideline:</b> based on the reason, provide instruction to avoid similar mistakes.</p> <p>Please do not mention the true answer or any specific option content in your response.</p>	
Study Assistant Response	
<p><b>Analysis:</b> The model might have misunderstood the phrase “1 day later” in the context.</p> <p><b>Guideline:</b> For dates in a problem, identify the correct date from which calculations should be made. Also, make sure to maintain the correct format (MM/DD/YYYY) while providing the answer.</p>	<p><math>a_1^{(0)}</math></p>

# Evaluate 3 LLMs on 27 tasks

## ❖ Benchmark

### ➤ BBH: 16 reasoning tasks

✓ *data calculation, tracking shuffled objects, ...*

### ➤ BBQ: 11 social bias tasks

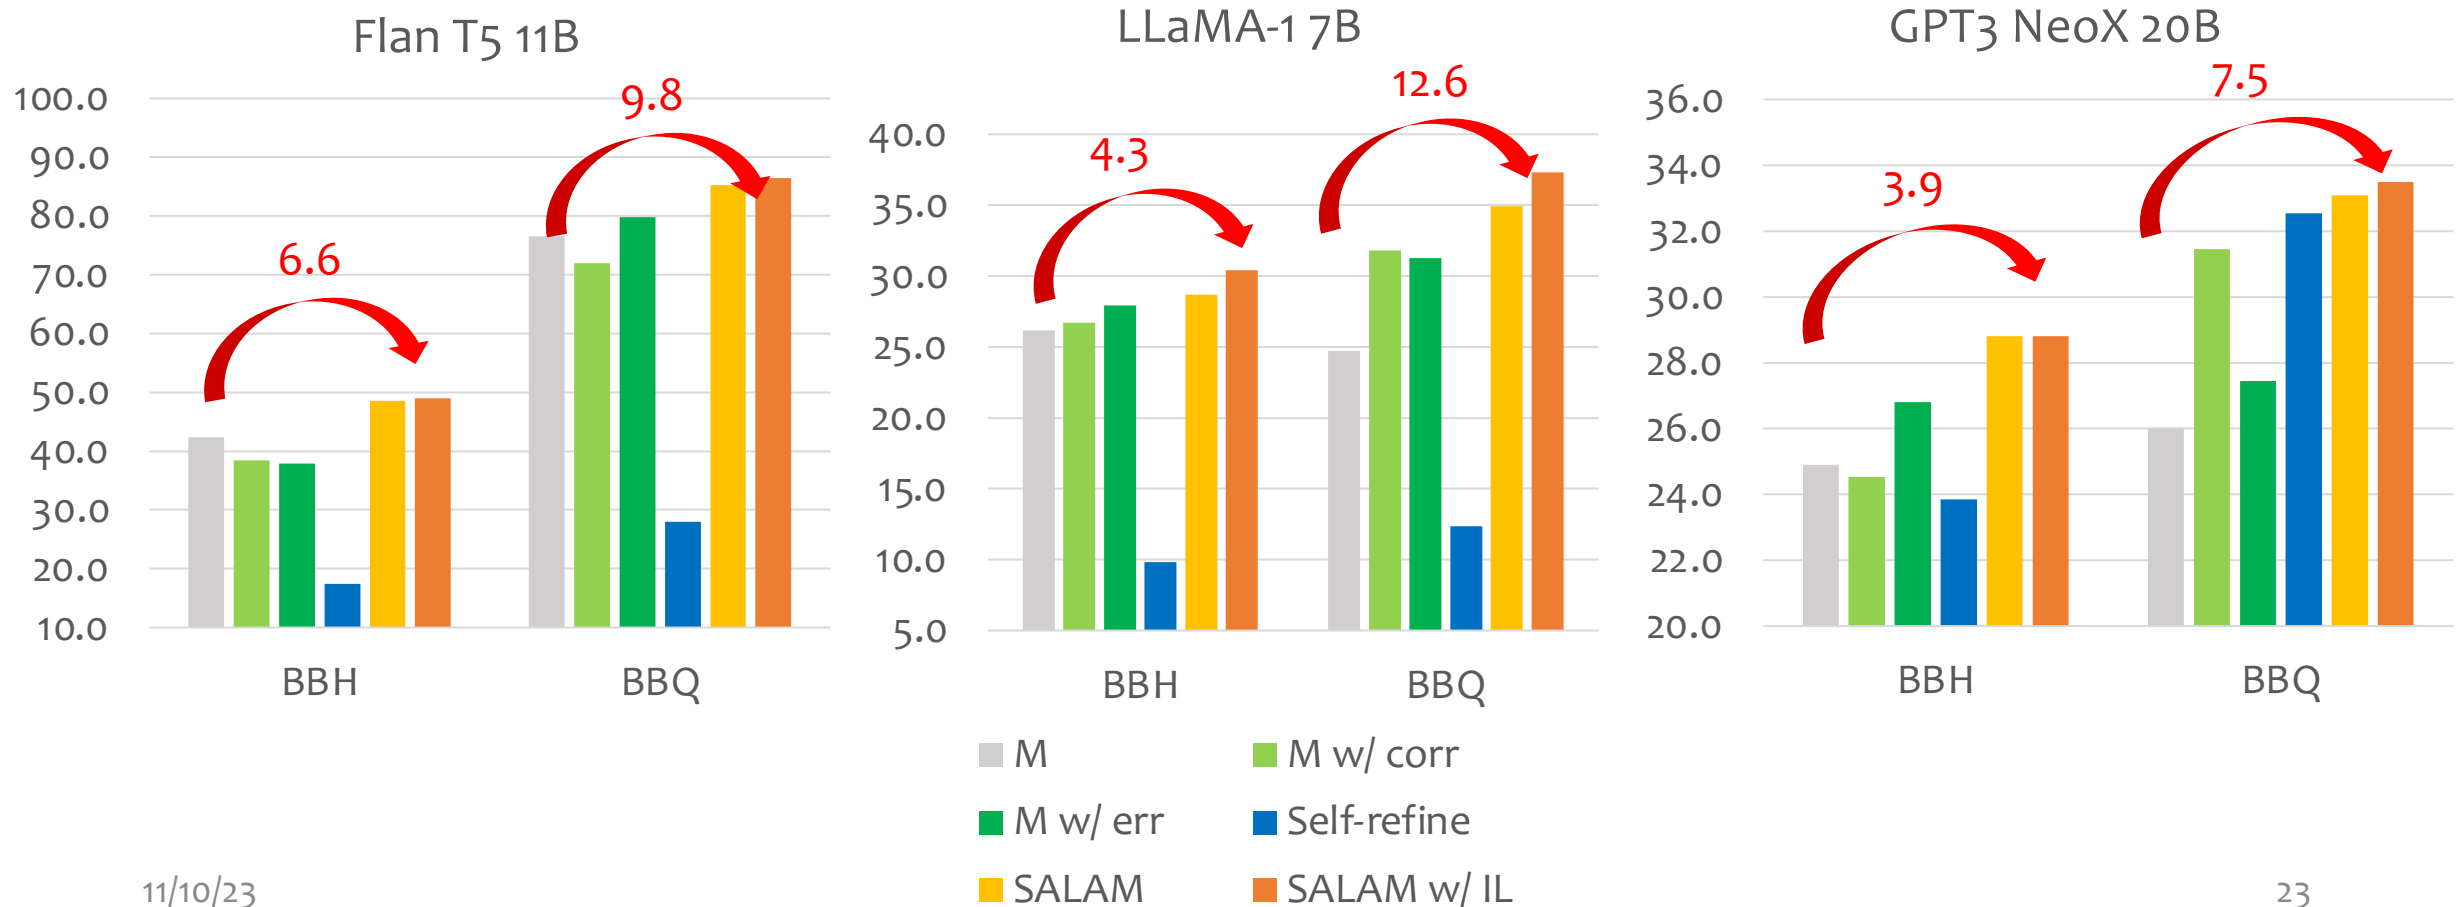
✓ *age, sex, nationality, ...*

## ❖ Main LLM: Flan-T5, GPT-2 NeoX, LLaMA-1

## ❖ Study Assistant: finetuned LLaMA-1 7B

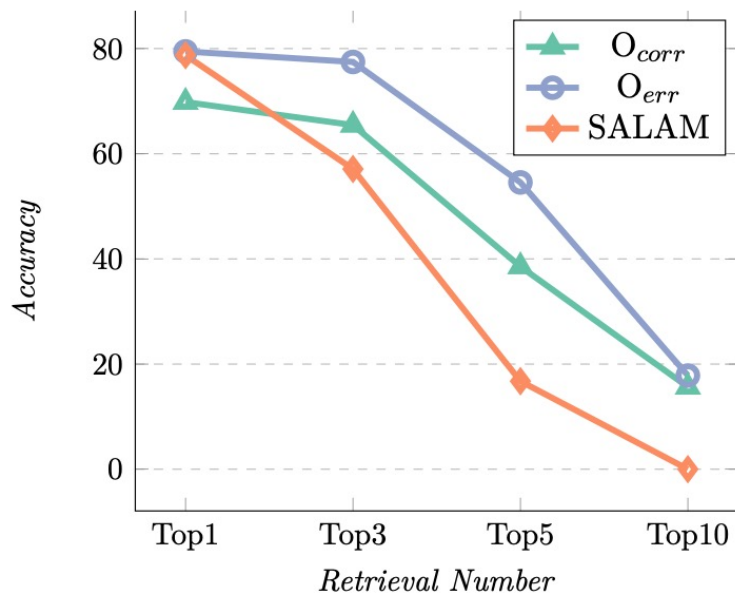
# SALAM Boosts LLM performance

*Accuracy under the multi-choice setting*

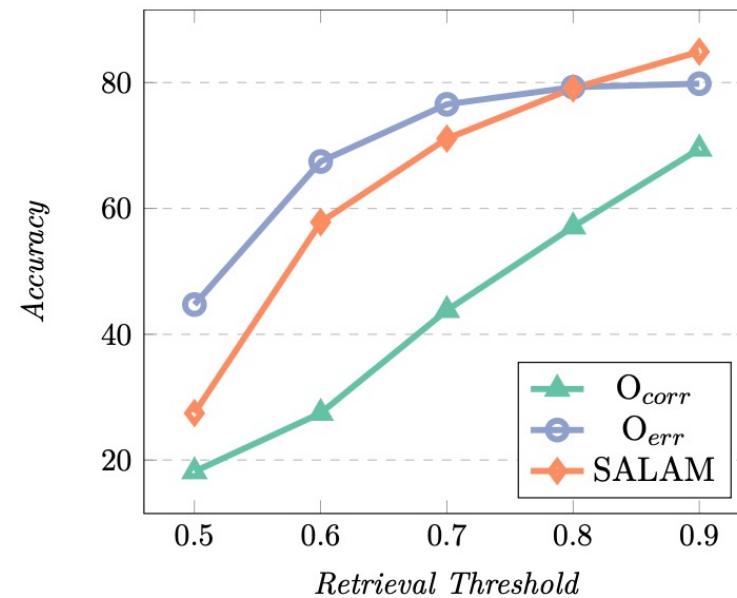


# Observation: Retrieval highly impacts the performance

*Relevance of retrieved examples is more important than the number*



(a) Topk.  $\theta$  is set to 0.



(b)  $\theta$ . Topk is set to 10.



# Observation: Expand to GPT-4

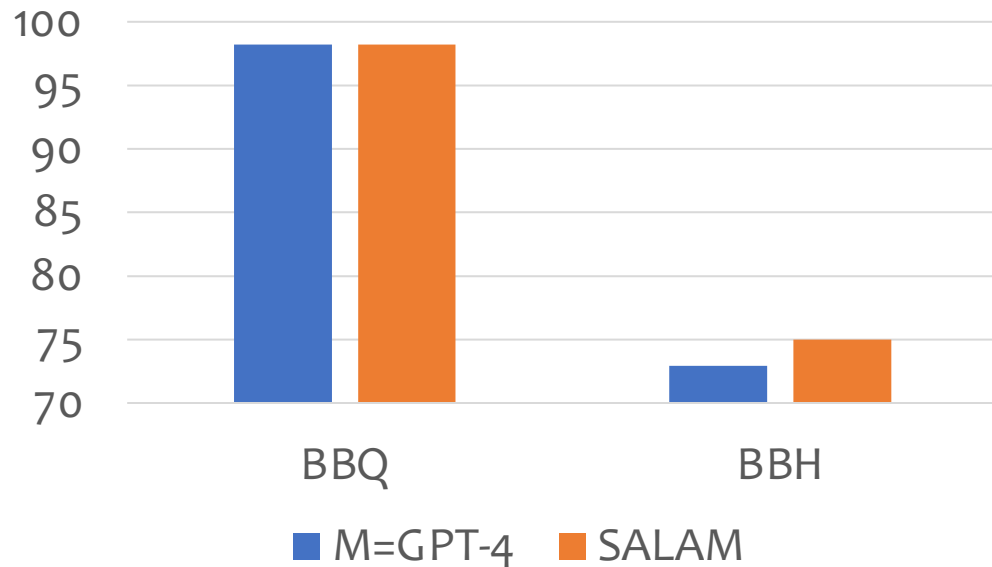


*SALAM boosts GPT-4 on BBH*

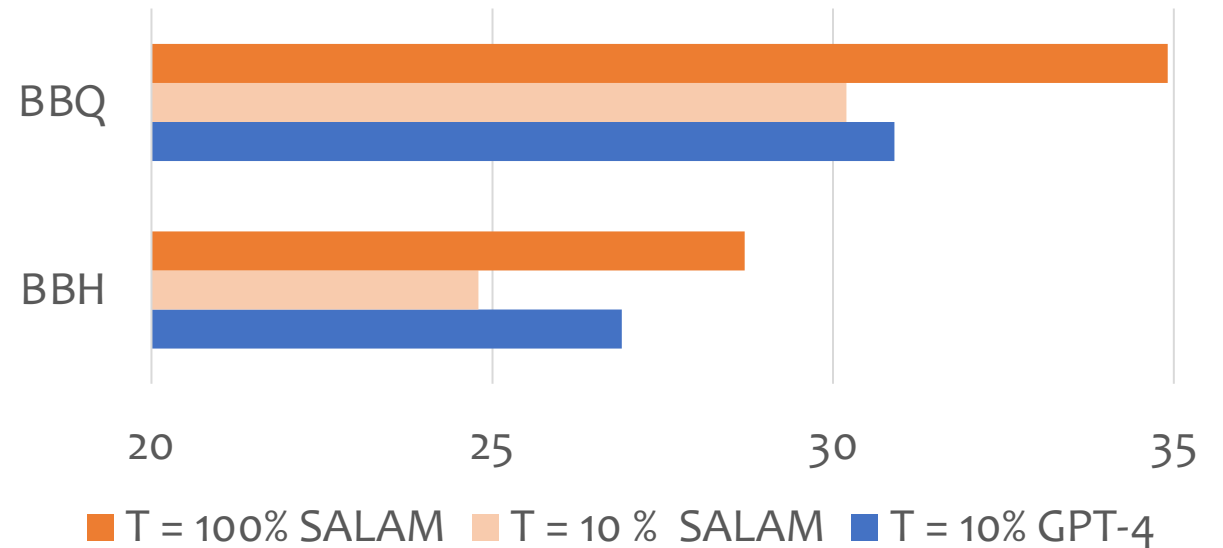


*GPT-4 outperforms SALAM by the same amount, but struggles with providing more feedback*

GPT-4 as LLM & LLaMA as SA



GPT-4 as SA and LLaMA as LLM



# Case Study

❖ *The guideline gives necessary information about how to organize the adjective*

Which sentence has the correct adjective order:  
(A) red little silly cloth eating rectangular sock  
(B) silly little rectangular red cloth eating sock



**Query:** Which sentence has the correct adjective order:

(A) American triangular computer  
(B) triangular American computer

**Incorrect Answer:**

American triangular computer

**Guideline:** When evaluating sentences with multiple adjectives, consider the general rules of adjective order in English, such as *opinion, size, age, shape, color, origin, material, and purpose*.

Query



Retrieve



Mistake Collections

(B)



LLM

# Case Study

❖ *SALAM w/ IL is more specific*

- *visit number*
- *frequency of visit*

Jane visits the bookstore on the 16th of each month starting from the October of 2009. It is her 5th visit to the bookstore today. What is the date one year ago from today?

Previous answer: 08/16/2009

Correct answer: 02/16/2009



SALAM

**Guideline:** ensure that you accurately calculate the date by *considering the correct day of the month and subtracting the specified number of months* from the given date.



SALAM w/ IL

**Guideline:** carefully consider the given information, such as the *frequency of visits and the current visit number*, to accurately calculate the elapsed time. Then, use this information to determine the correct date.

# Take Away

- ❖ Unreliable and vague feedback might mislead LLMs
- ❖ SALAM: Cooperation between LLMs and study assistant will boost performance
  - *guidance from SA improves LLMs' performance*
  - *improvement of LLMs guides SA to provide better feedback*
- ❖ More to explore
  - *for other NLP tasks, such as open-ended generation*
  - *for general tasks, such as robotics with multimodal input*
  - *for better feedback, such as reference-free*

# Thank you for listening!

[danqingwang@ucsb.edu](mailto:danqingwang@ucsb.edu)

*Jane thought today is 3/11/2002, but today is in fact Mar 12, which is 1 day later. What is the date a month ago?*

**Ground Truth**

02/12/2002

## Scan Me!



Project Page