

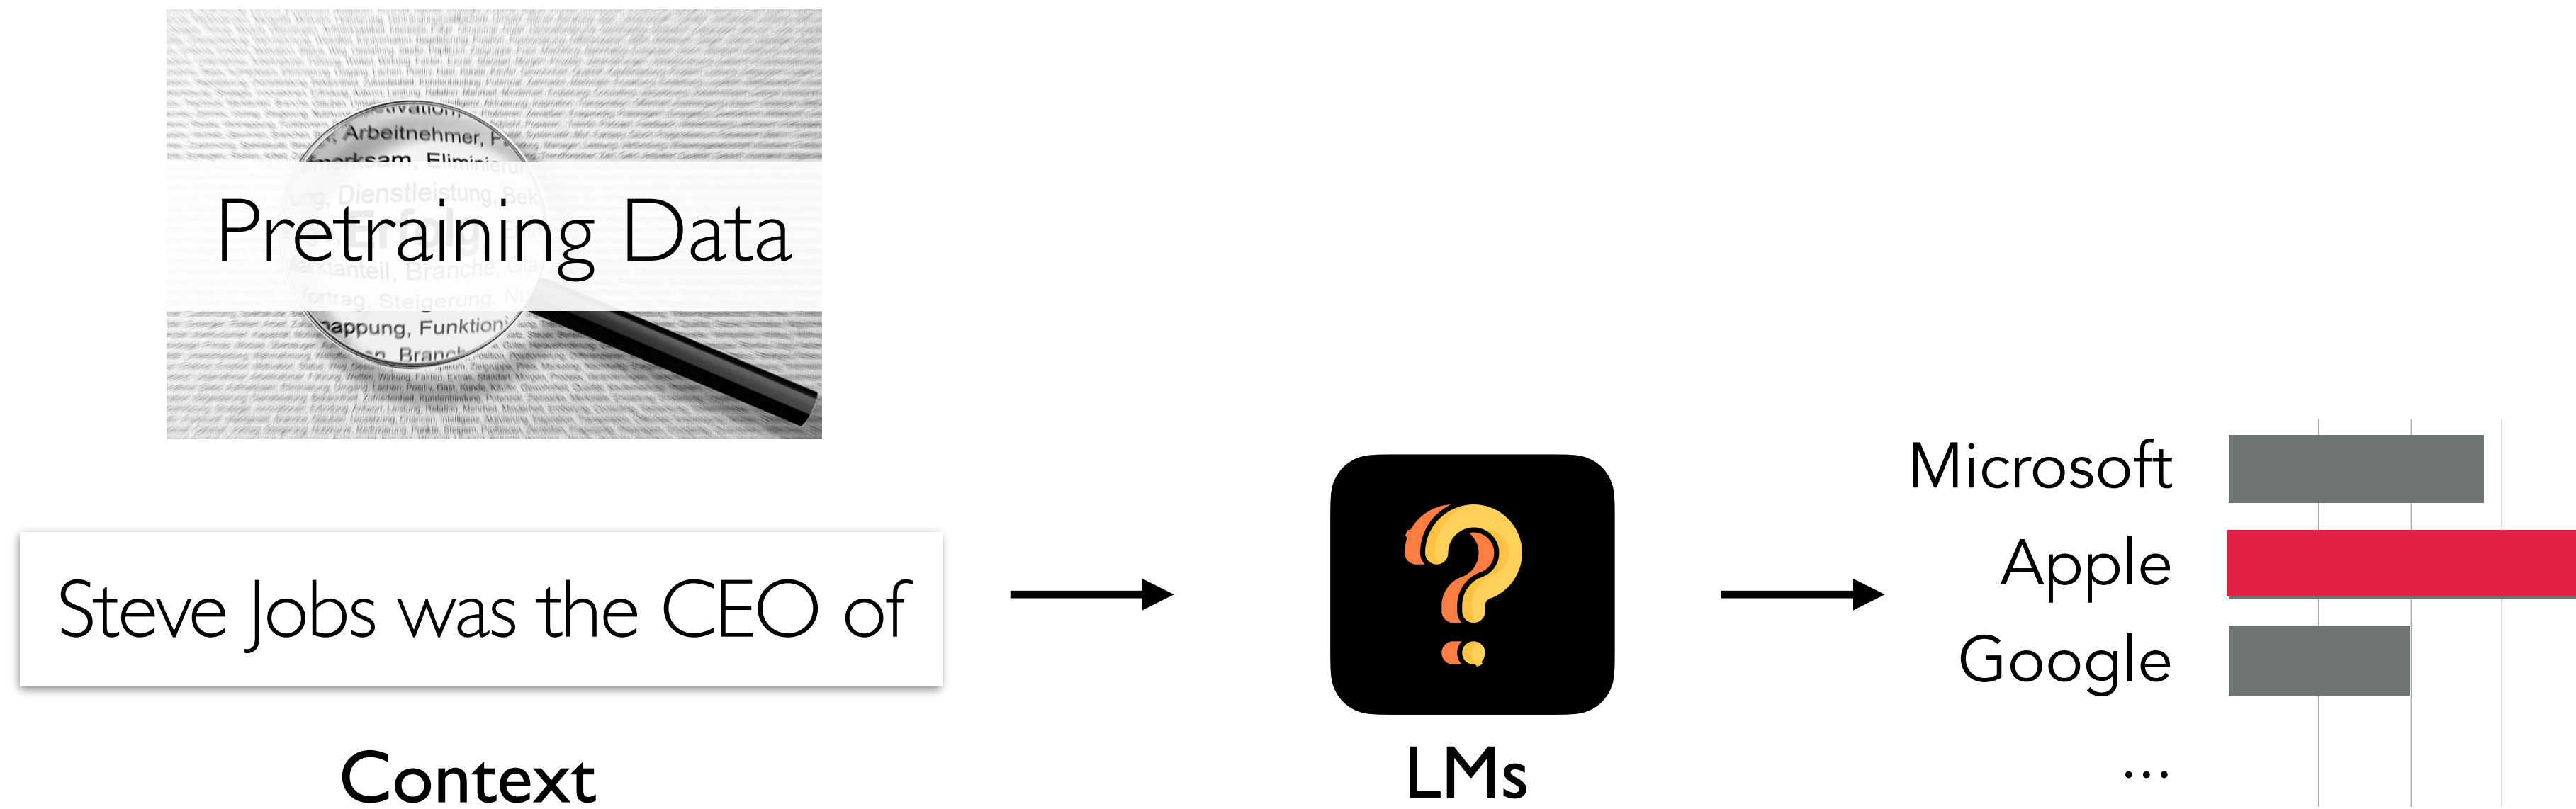
# Beyond Monolithic Language Models

Weijia Shi

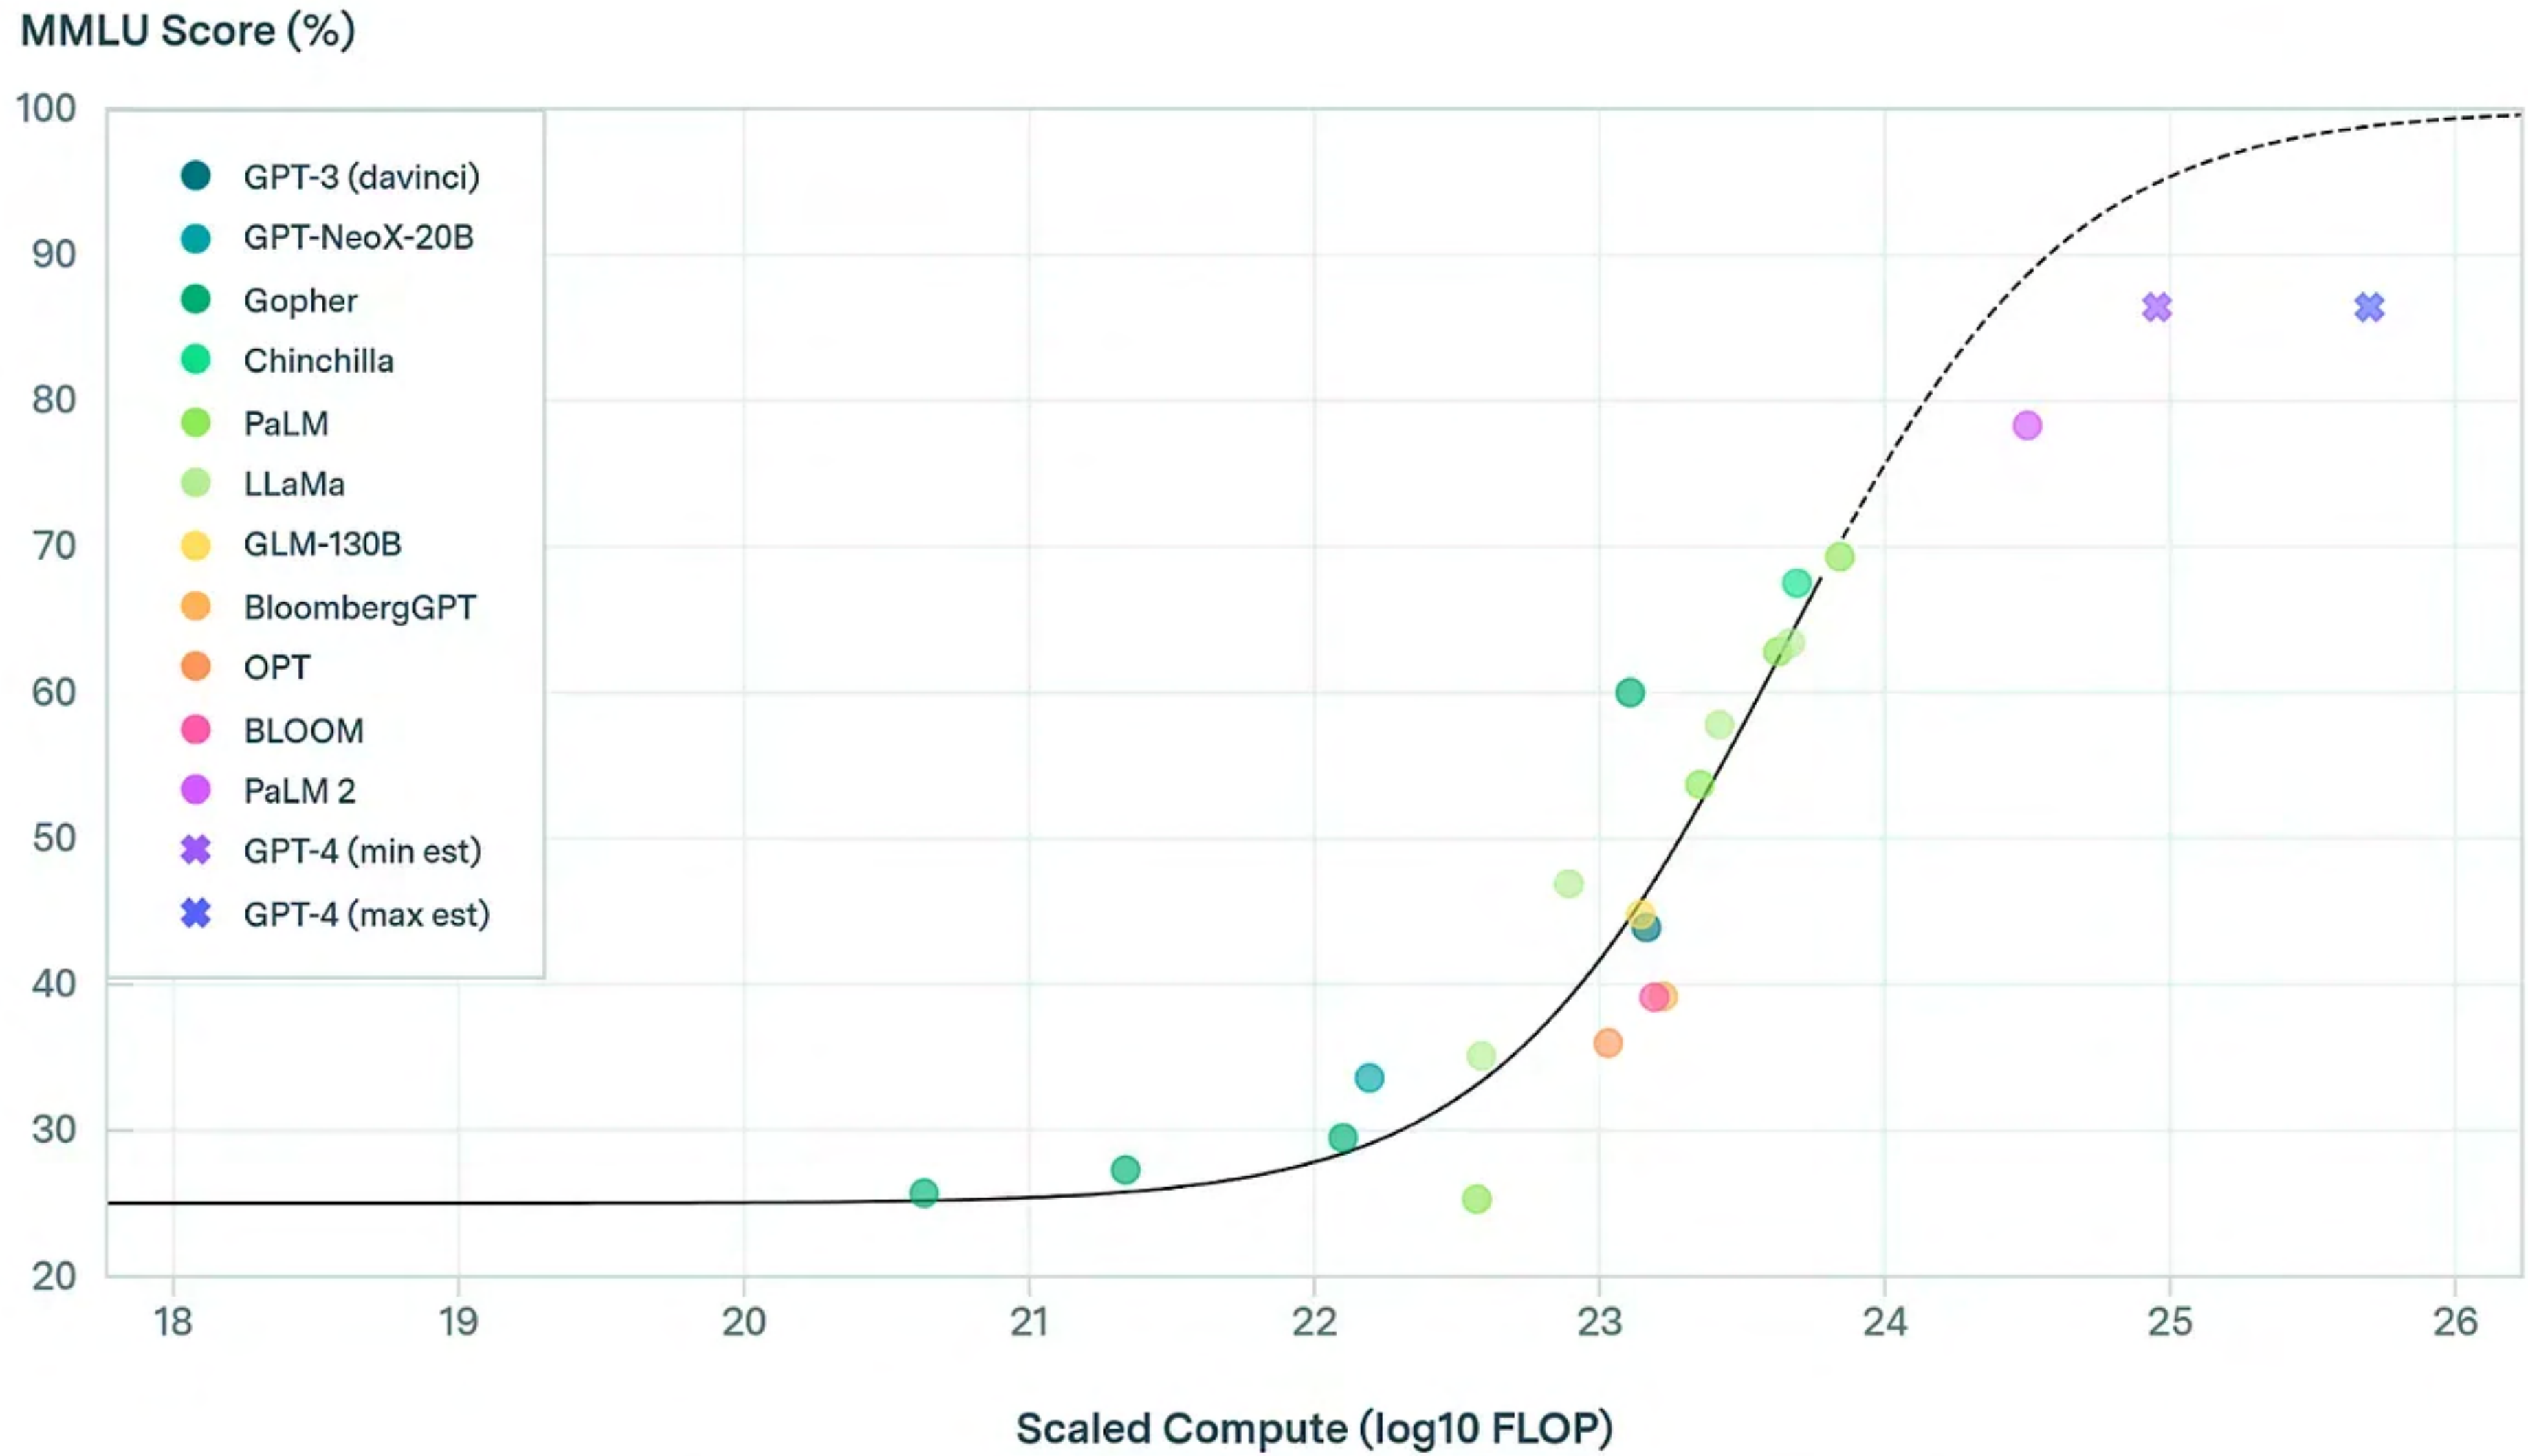
University of Washington



# Current Language Models (LMs): Monolithic



# Scaling



# *Scaling Up Monolithic LMs*

# Problems

## Hallucinations



ARTIFICIAL INTELLIGENCE / TECH / POLICY

**Former Trump lawyer Michael Cohen accidentally cited fake court cases generated by AI** / Michael Cohen provided his lawyer with made-up court cases generated by Google Bard after mistaking the chatbot for a 'a super-charged search engine.'



# Problems

Hallucinations

Hard to update new knowledge

ChatGPT o1-mini >  

What is OpenAI o1?

 Thought for a few seconds 

“OpenAI O1” does not correspond to any publicly recognized product, model, or initiative from OpenAI. It’s possible that “O1” refers to one of the following:

1. **Typographical Error:** It might be a misspelling or shorthand for another OpenAI product, such as GPT-1 (the first generation of the Generative Pre-trained Transformer models) or a different version like GPT-3 or GPT-4.

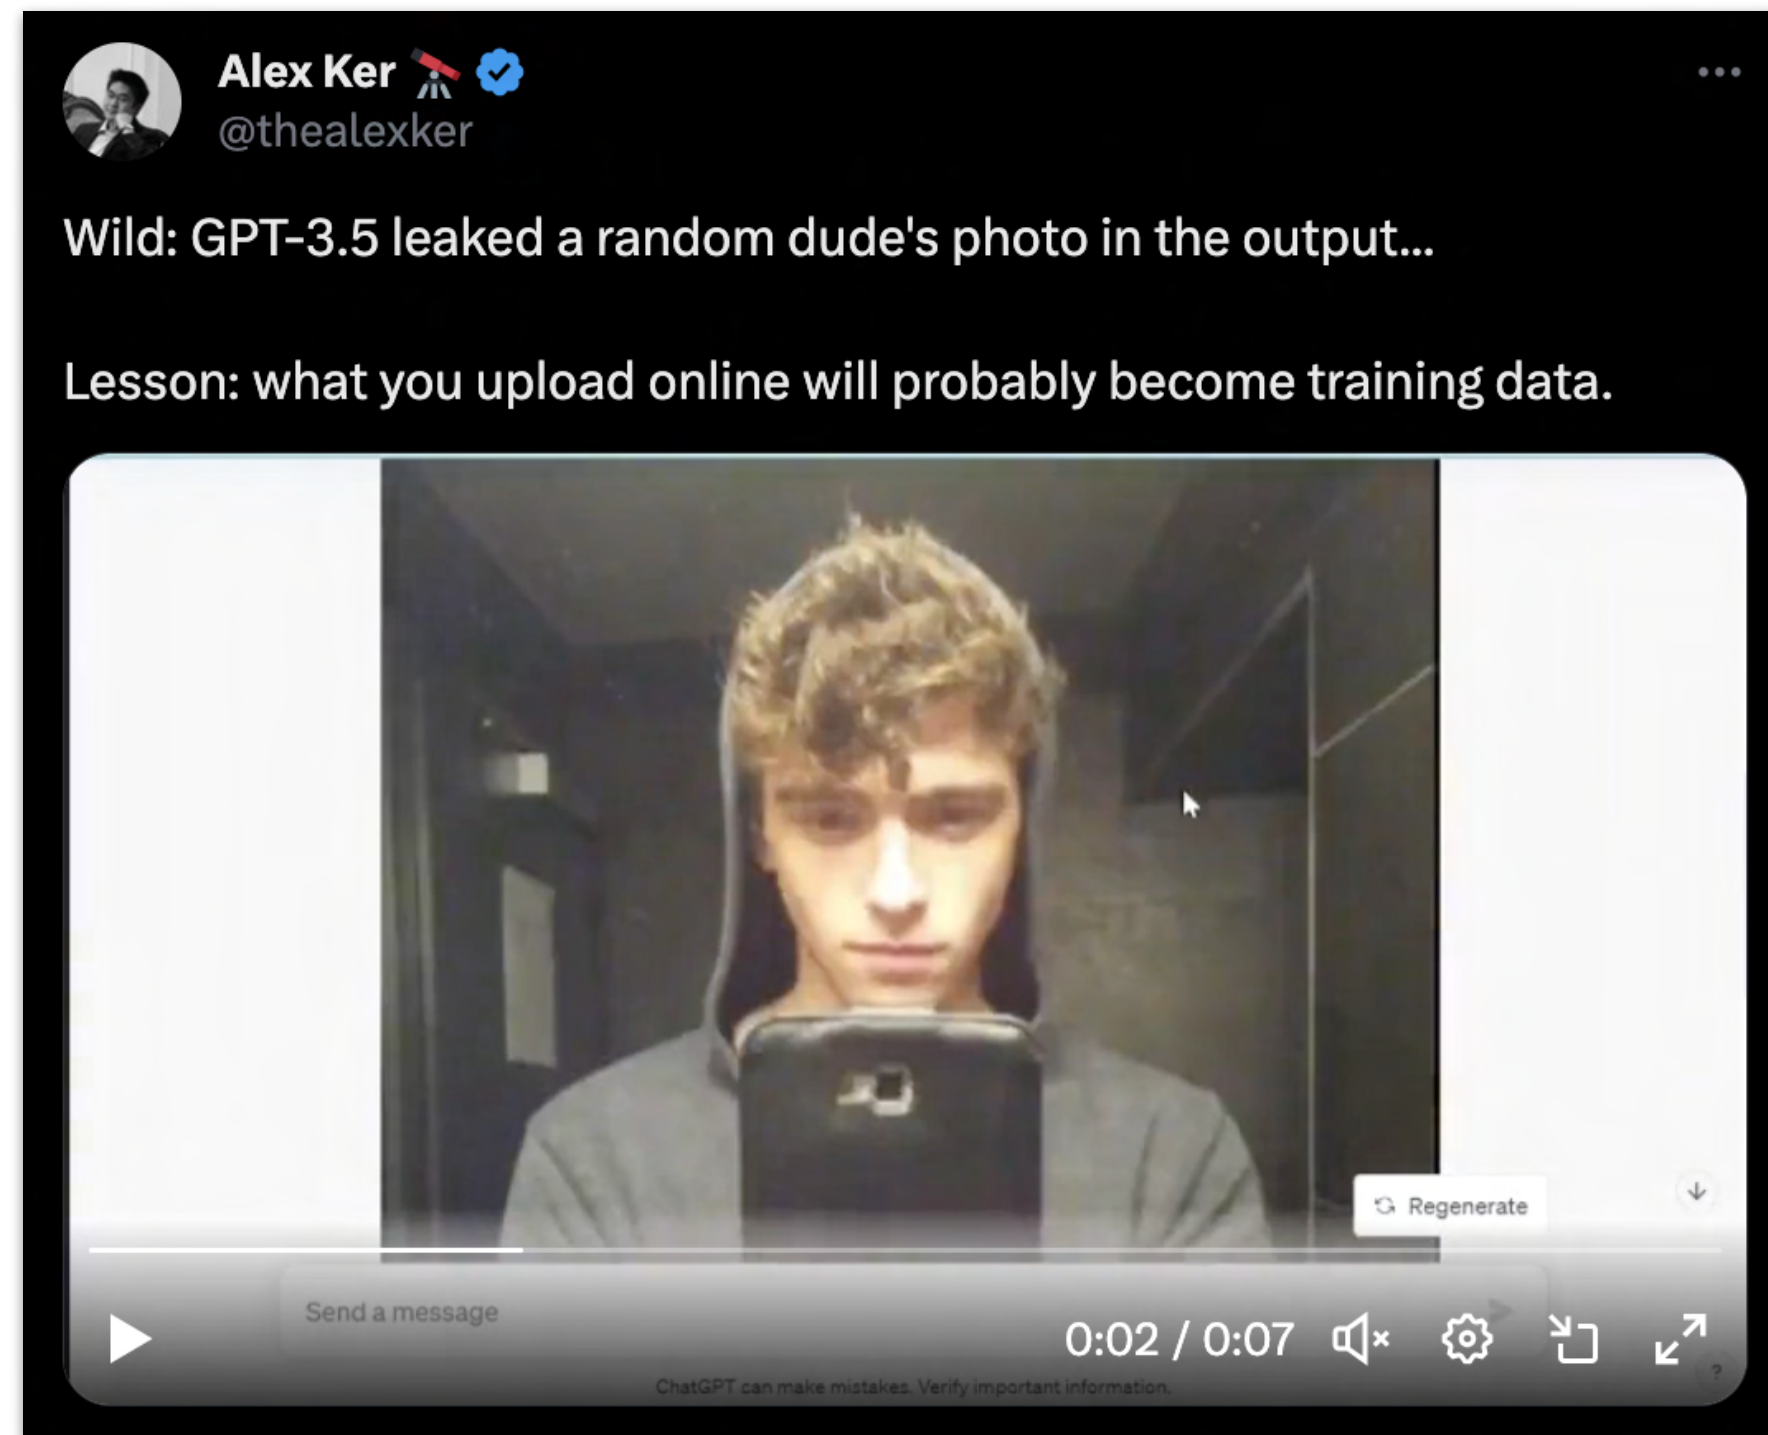


# Problems

Hallucinations

Hard to update new knowledge

Copyright and privacy risks



## The Times Sues OpenAI and Microsoft Over A.I. Use of Copyrighted Work

Millions of articles from The New York Times were used to train chatbots that now compete with it, the lawsuit said.

Dec. 27, 2023

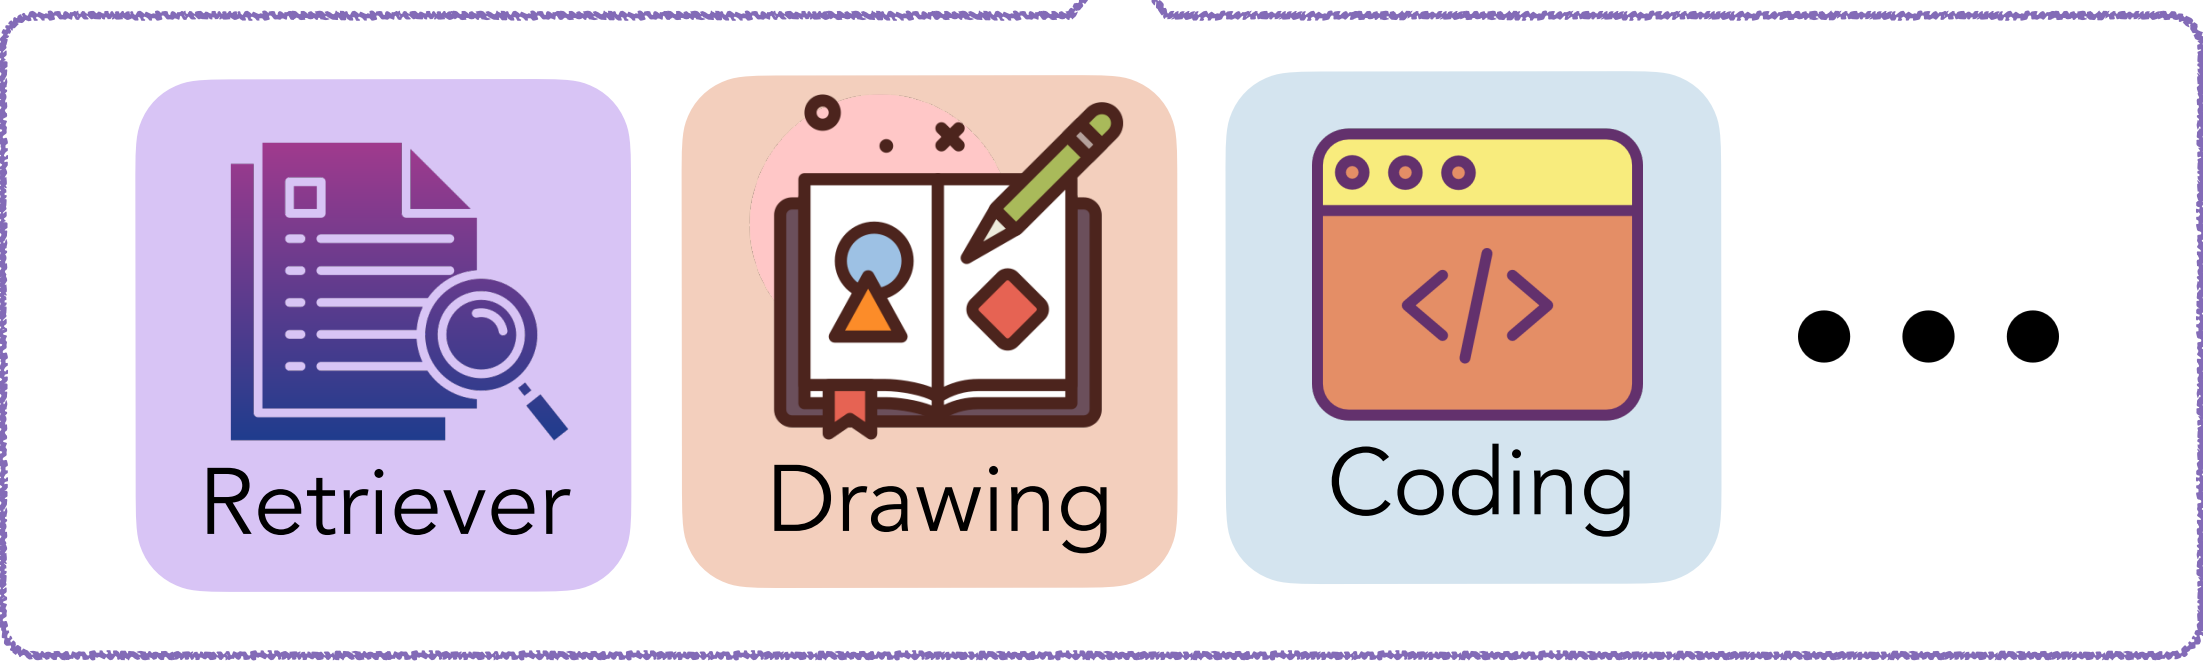
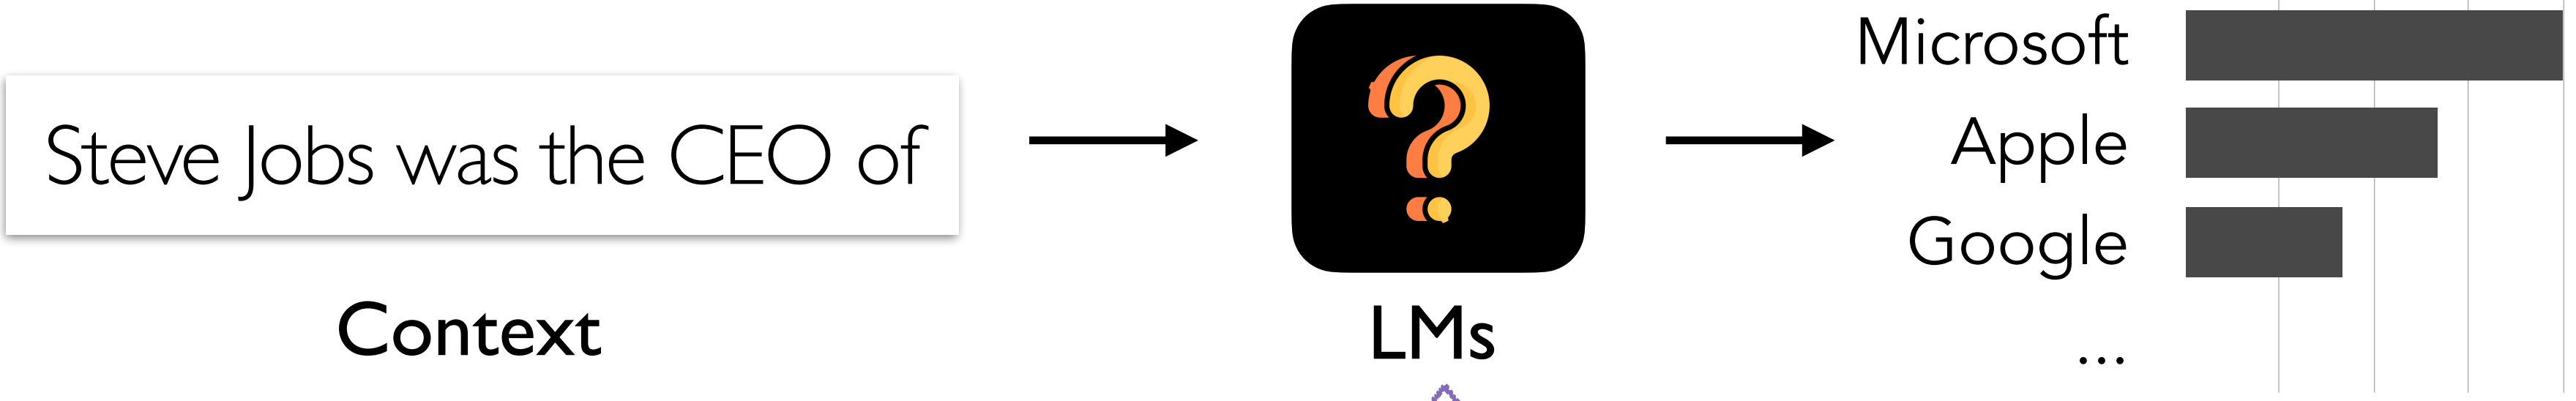
 *Scaling Up Monolithic LMs*



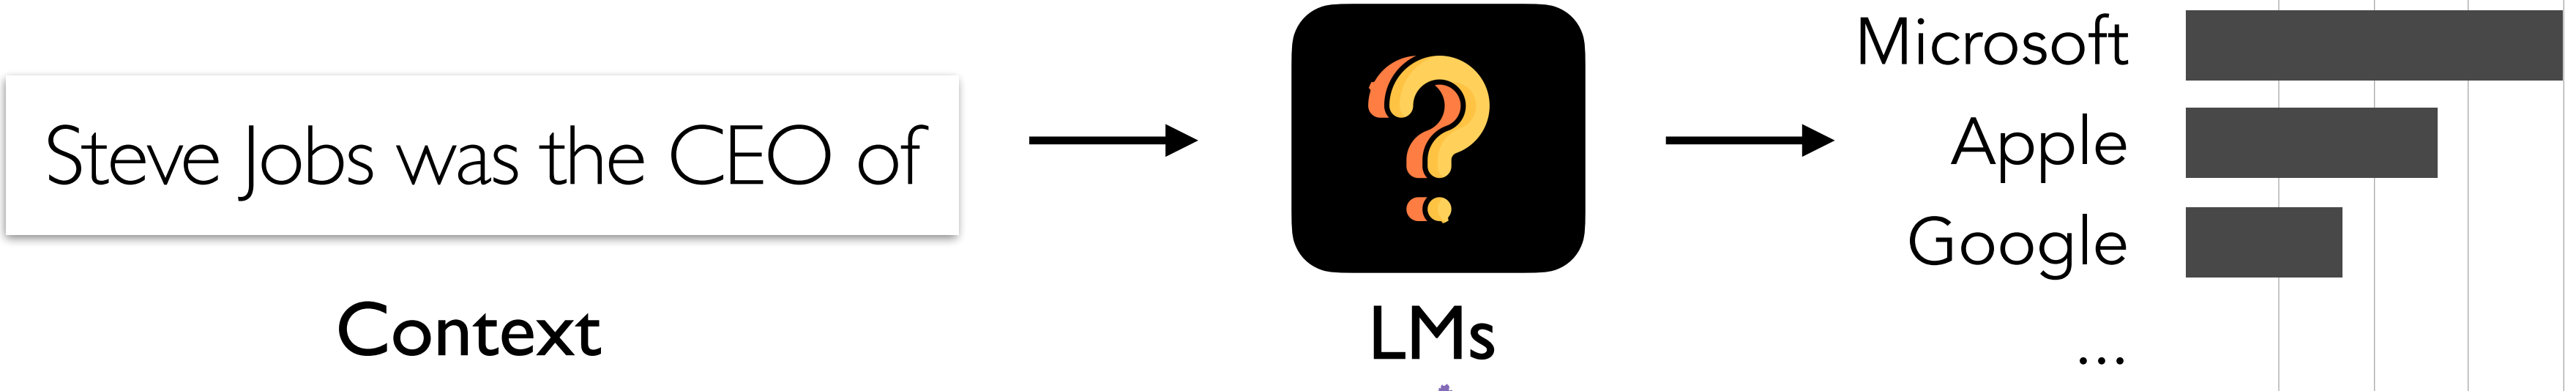
 *Scaling Up Monolithic LMs*

 *Alternative Paradigm*

# *Modularity, not Monoliths*



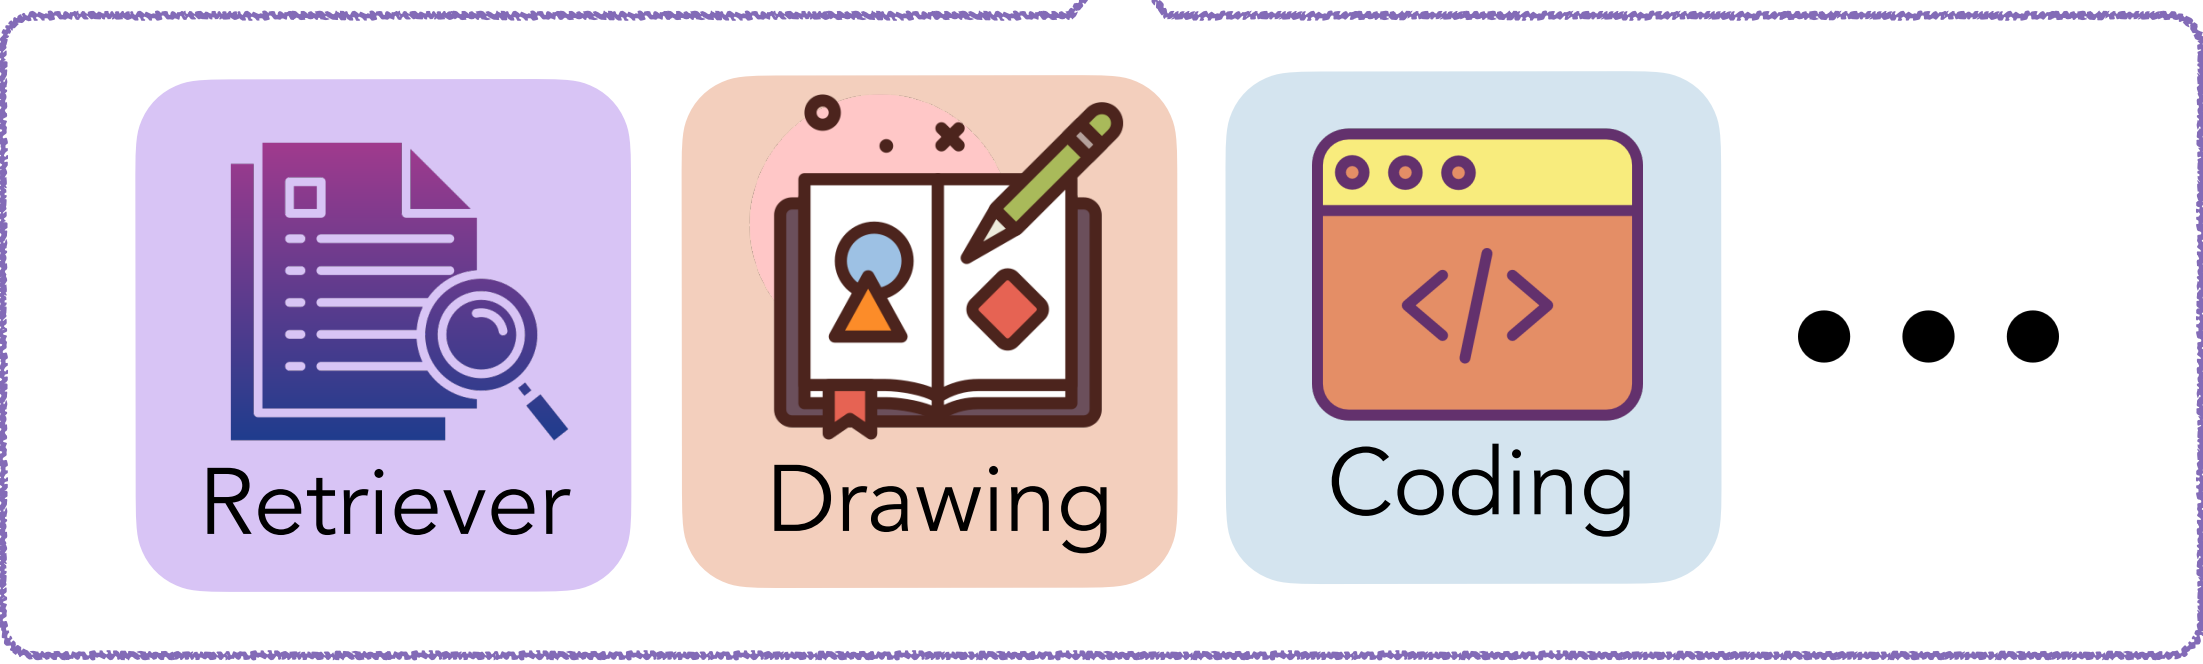
*Augmented Models*



Context

LMs

Microsoft  
Apple  
Google  
...



*Augmented Models*

# Augmented Models

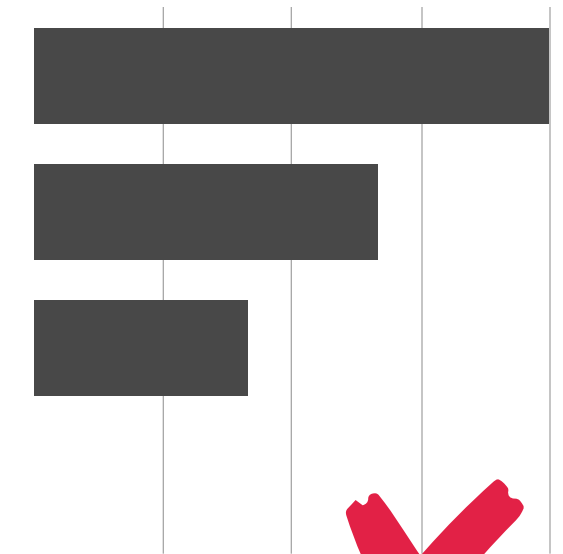
Context

Steve Jobs was the CEO of

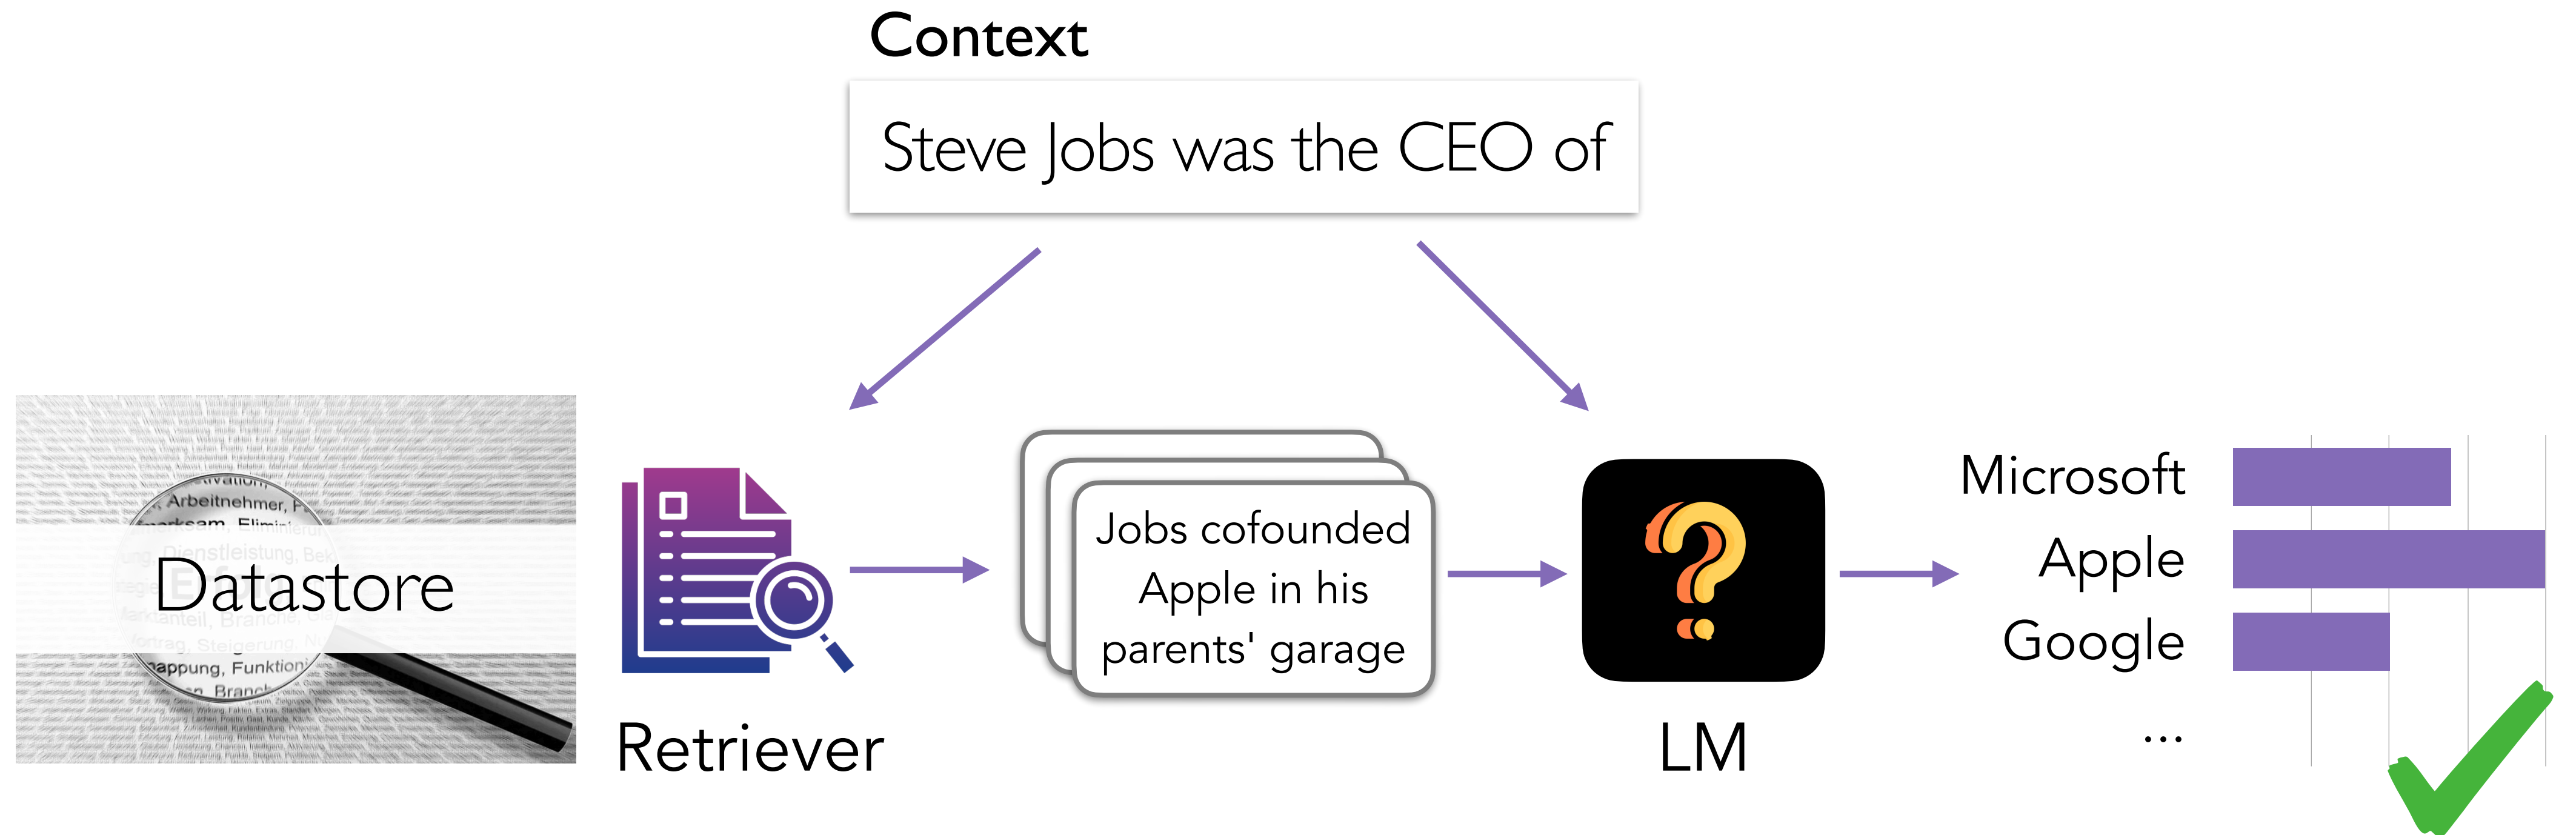


LM

Microsoft  
Apple  
Google  
...

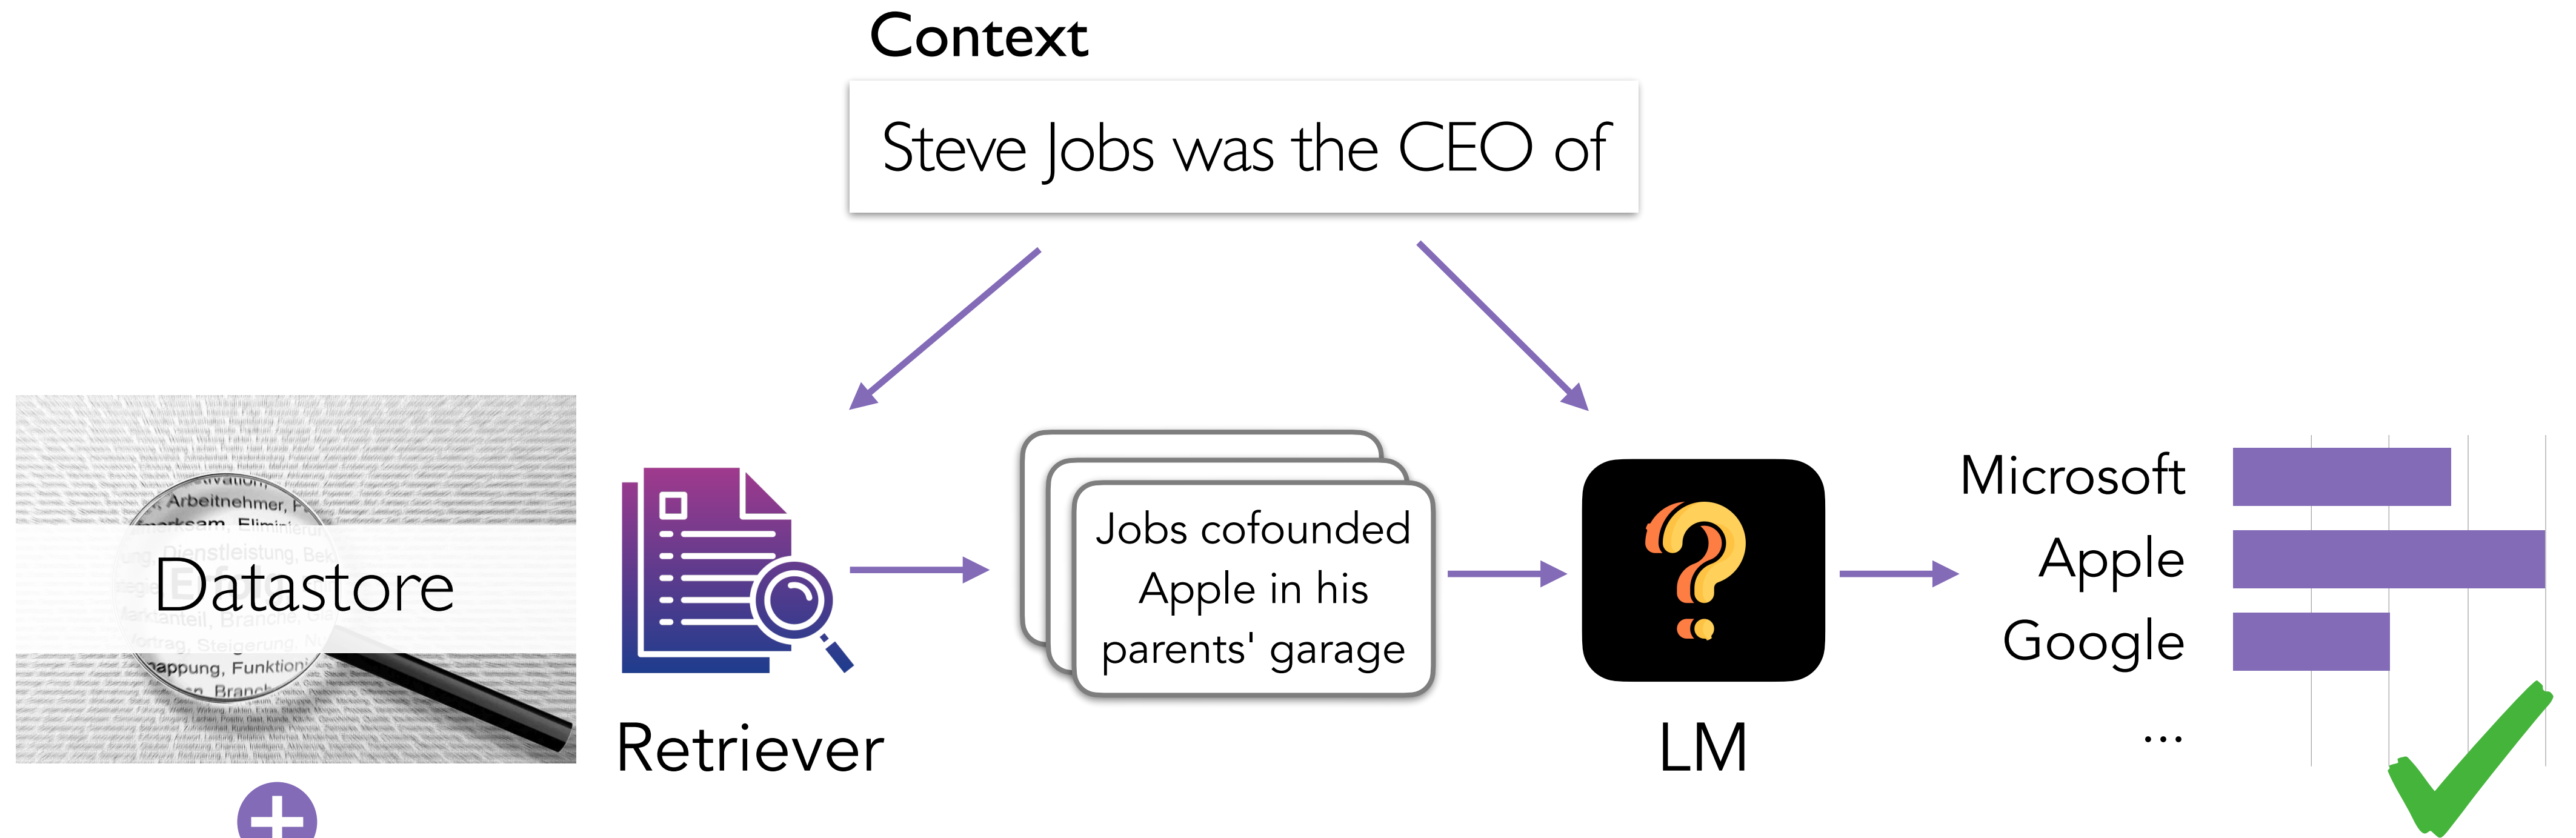


# Augmented Models



**Hallucinations**

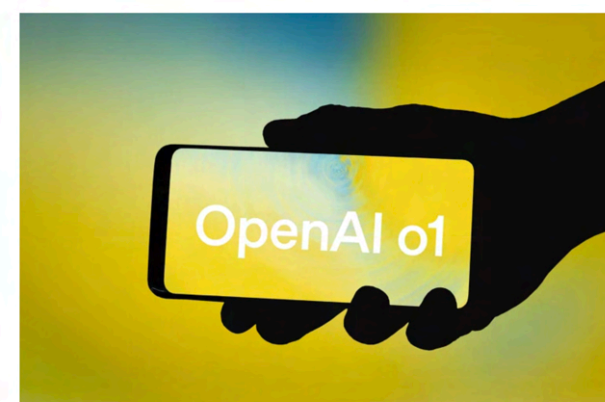
# Augmented Models



**'In awe': scientists impressed by latest ChatGPT model o1**

The chatbot excels at science, beating PhDs on a hard science test. But it might 'hallucinate' more than its predecessors.

By Nicola Jones



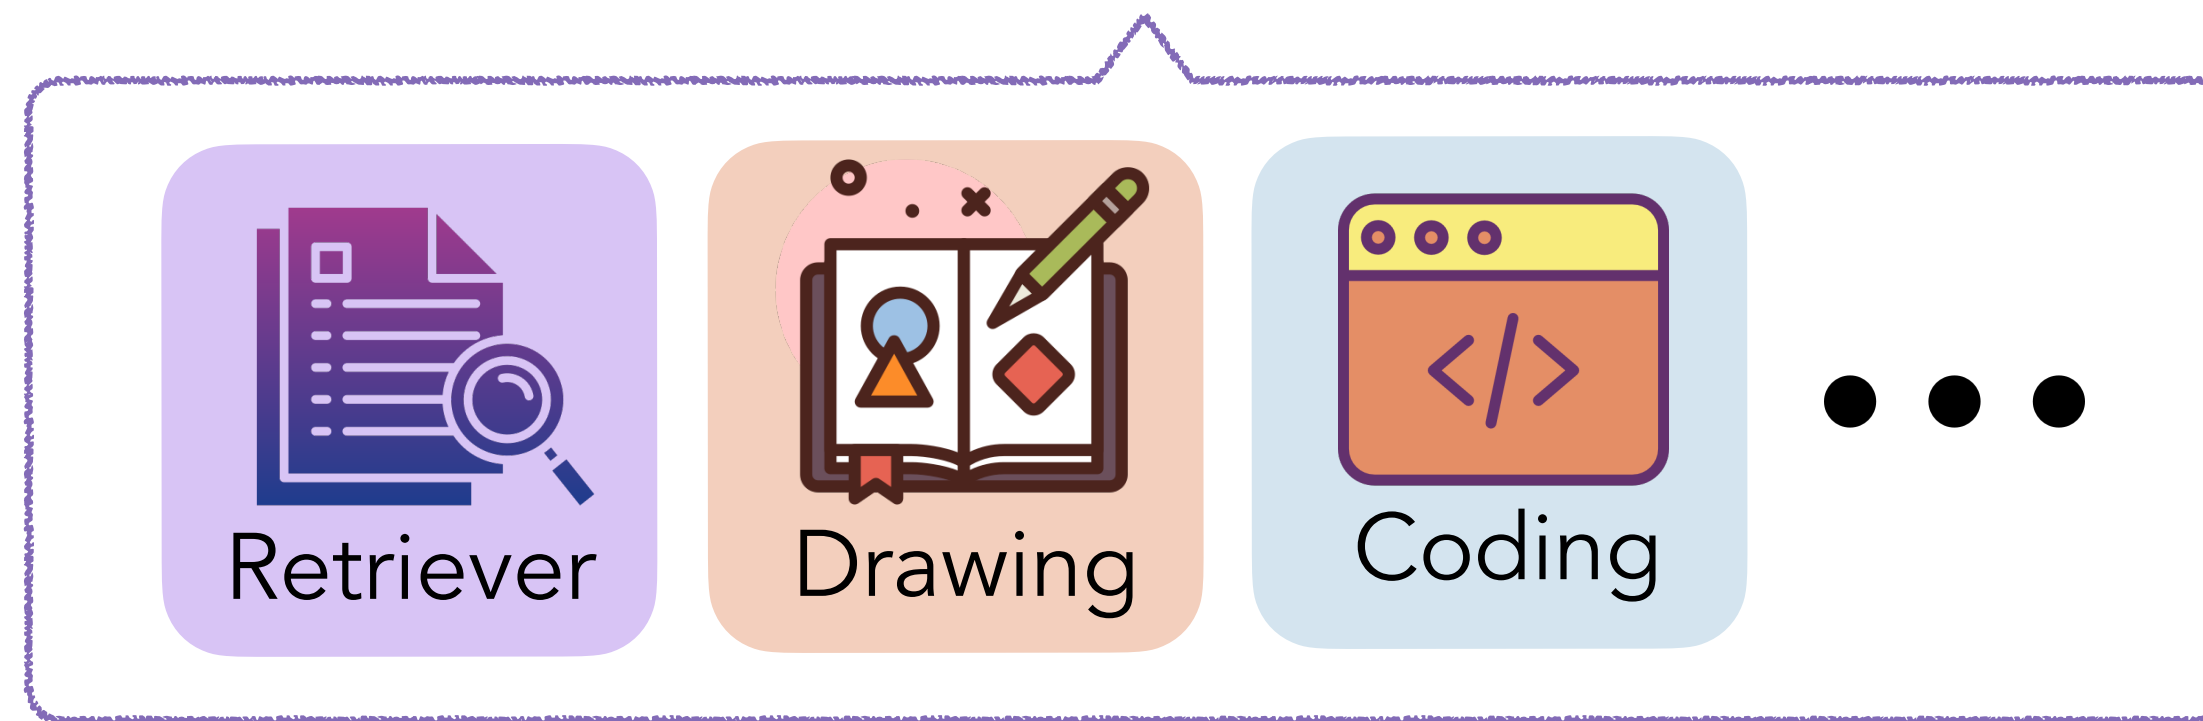
Technology firm OpenAI released a preview version of its latest chatbot, o1, last month. Credit: GK



**Hard to update new knowledge**



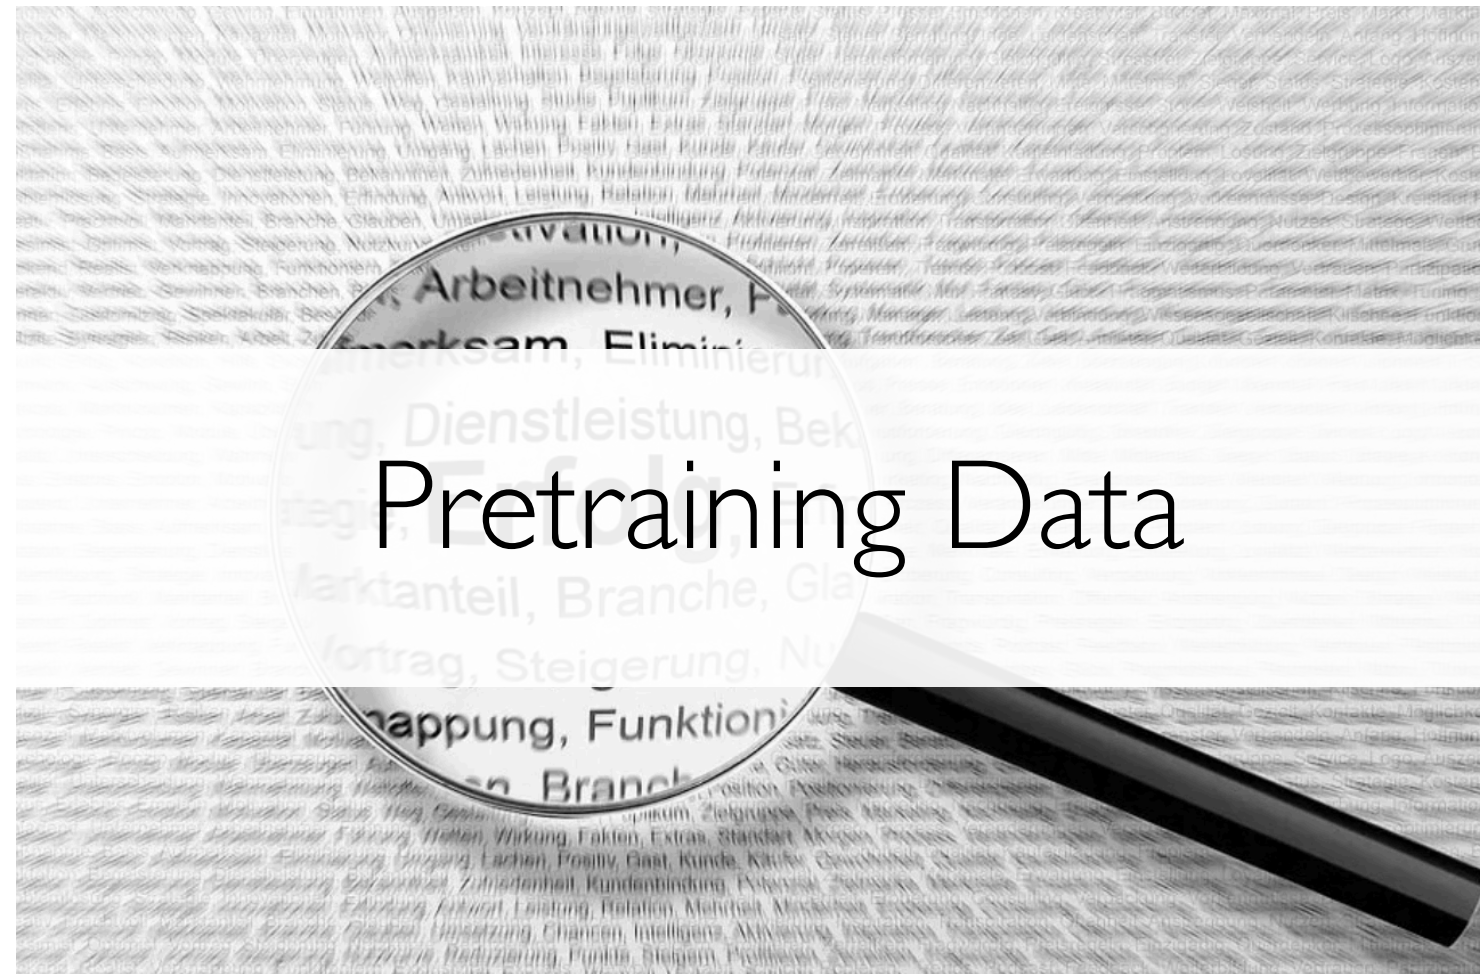
LM

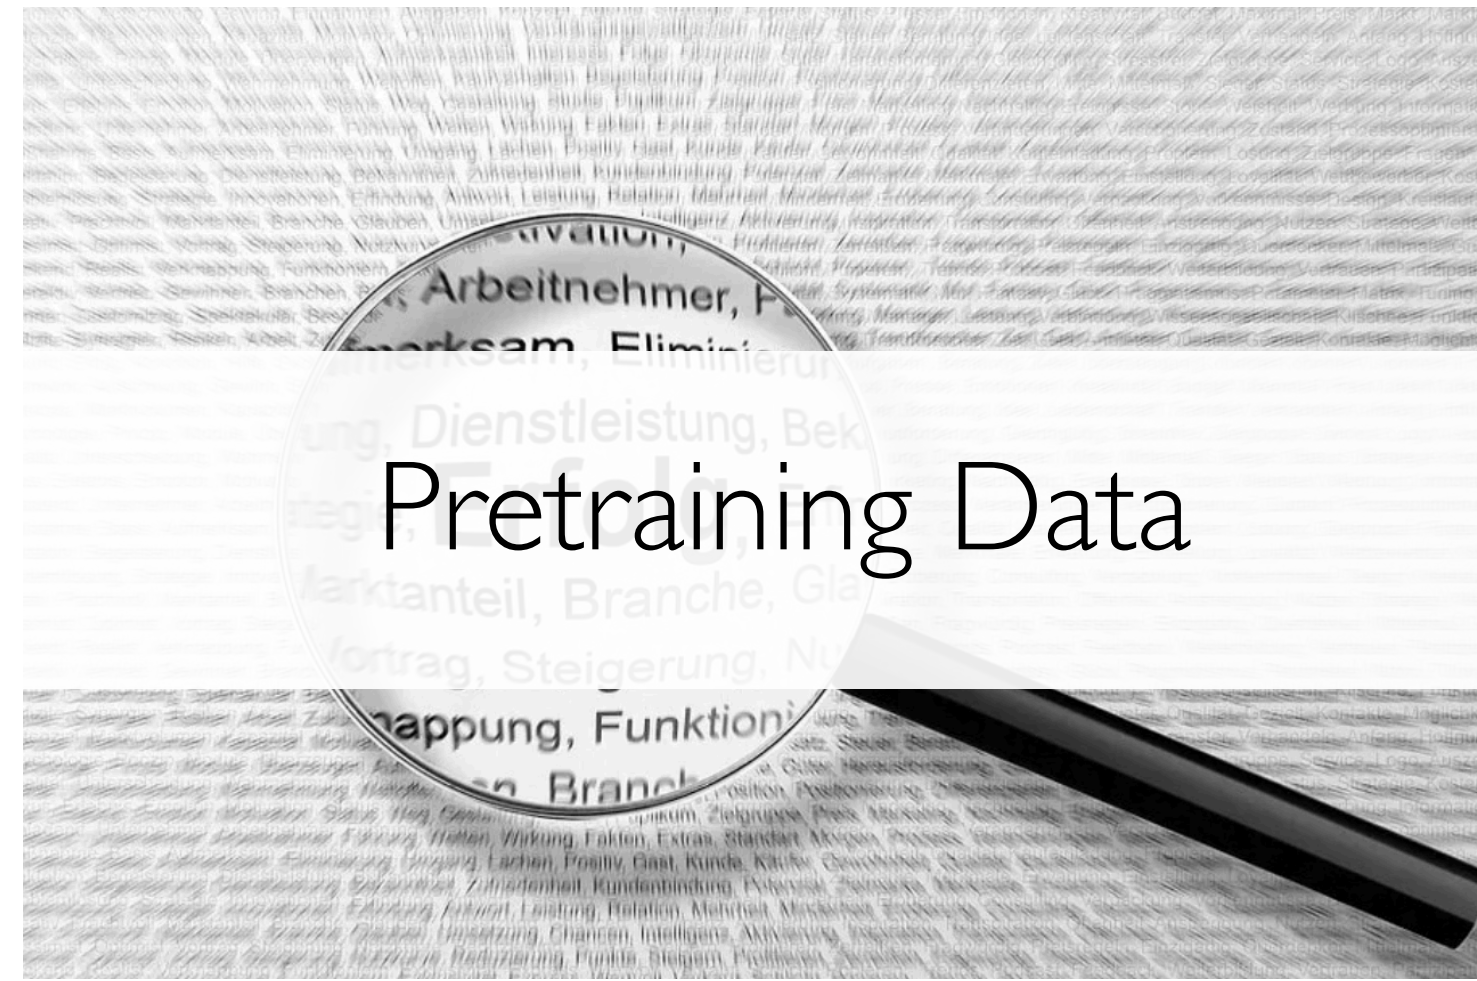


*Augmented Models*



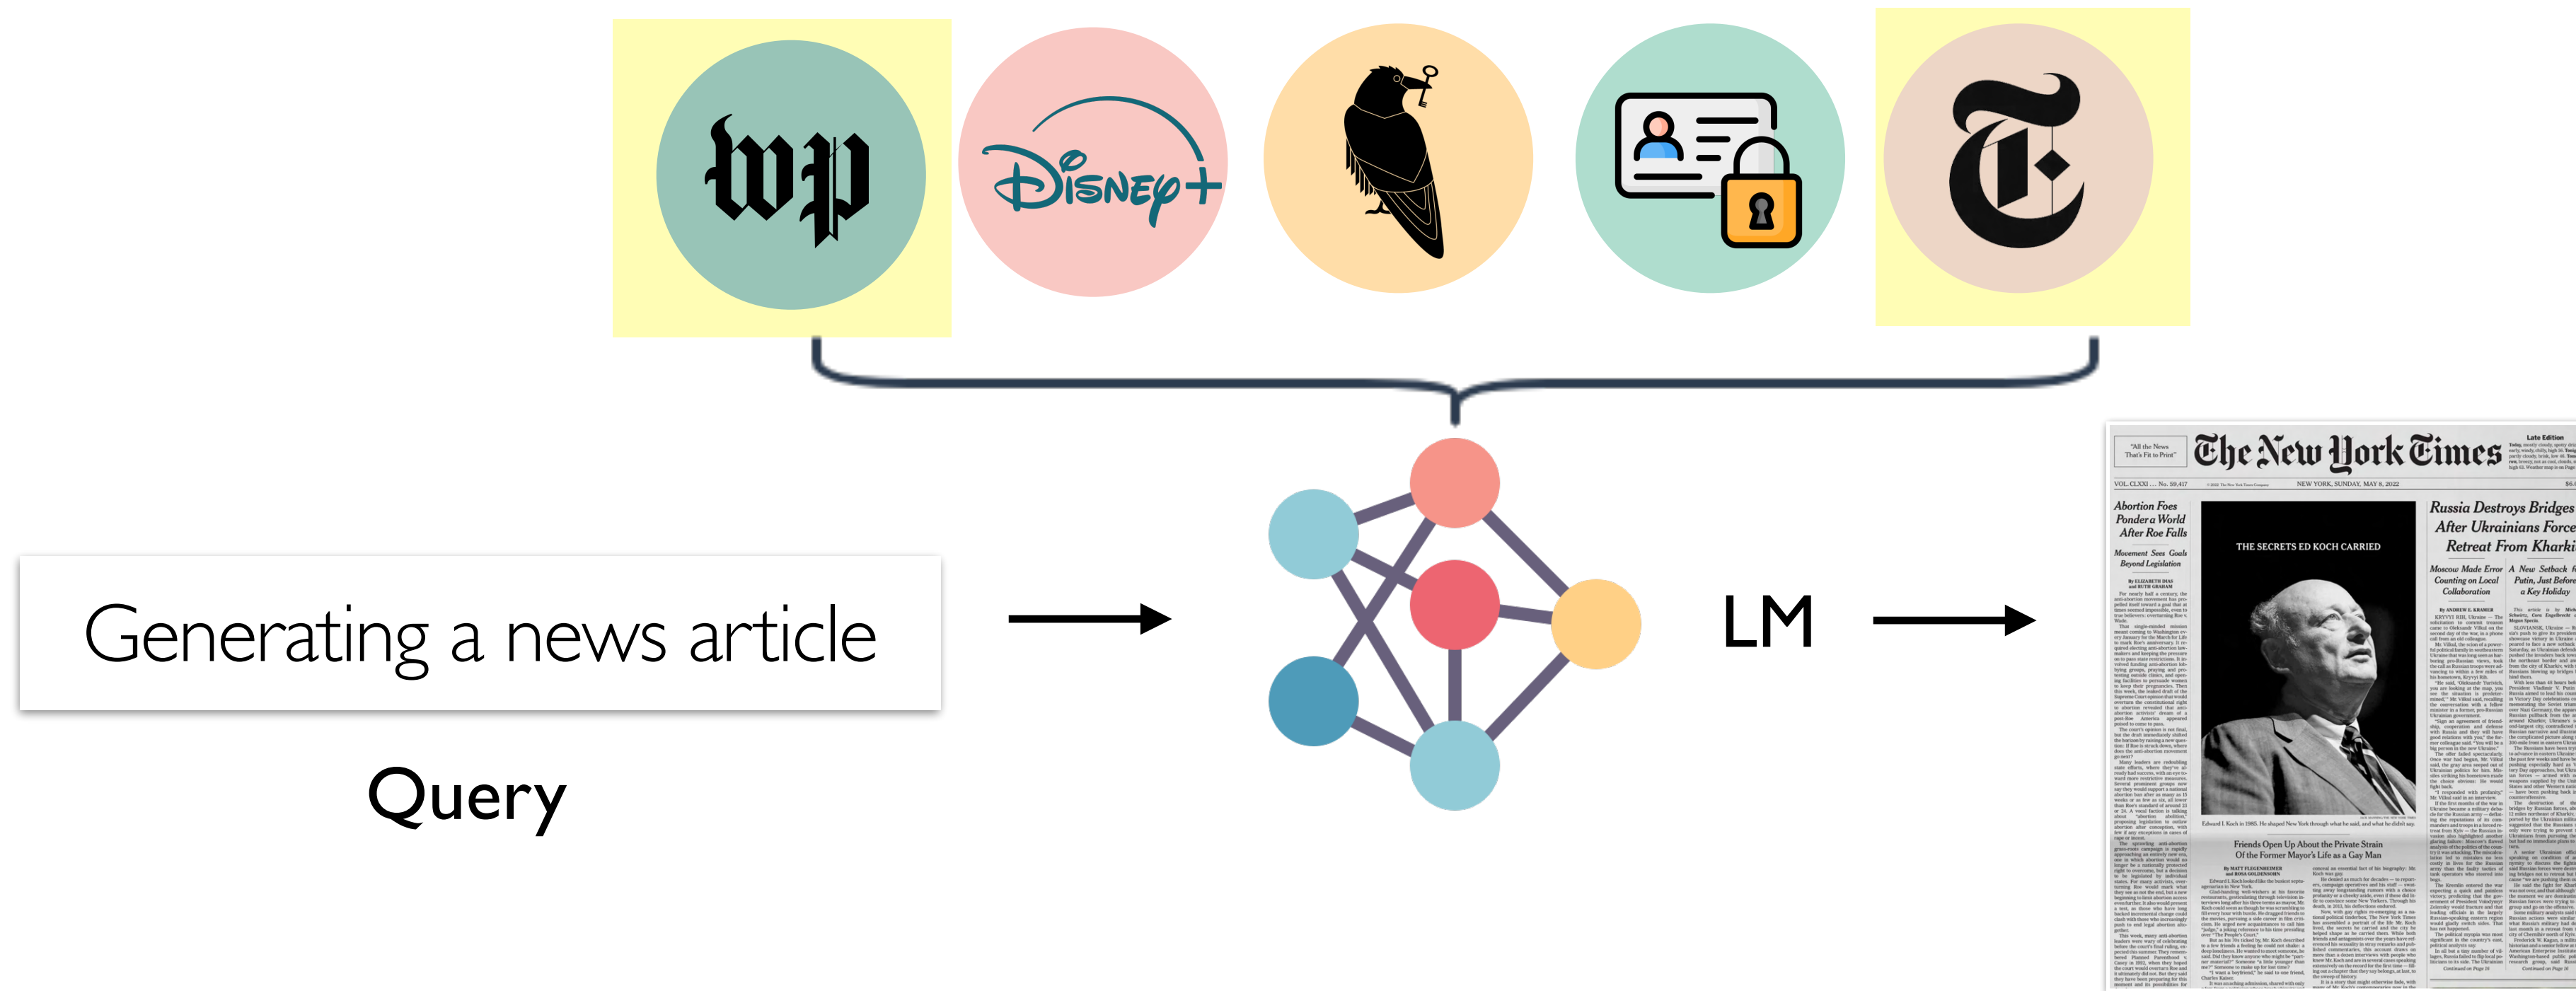
# Pretraining Data is not Monolithic





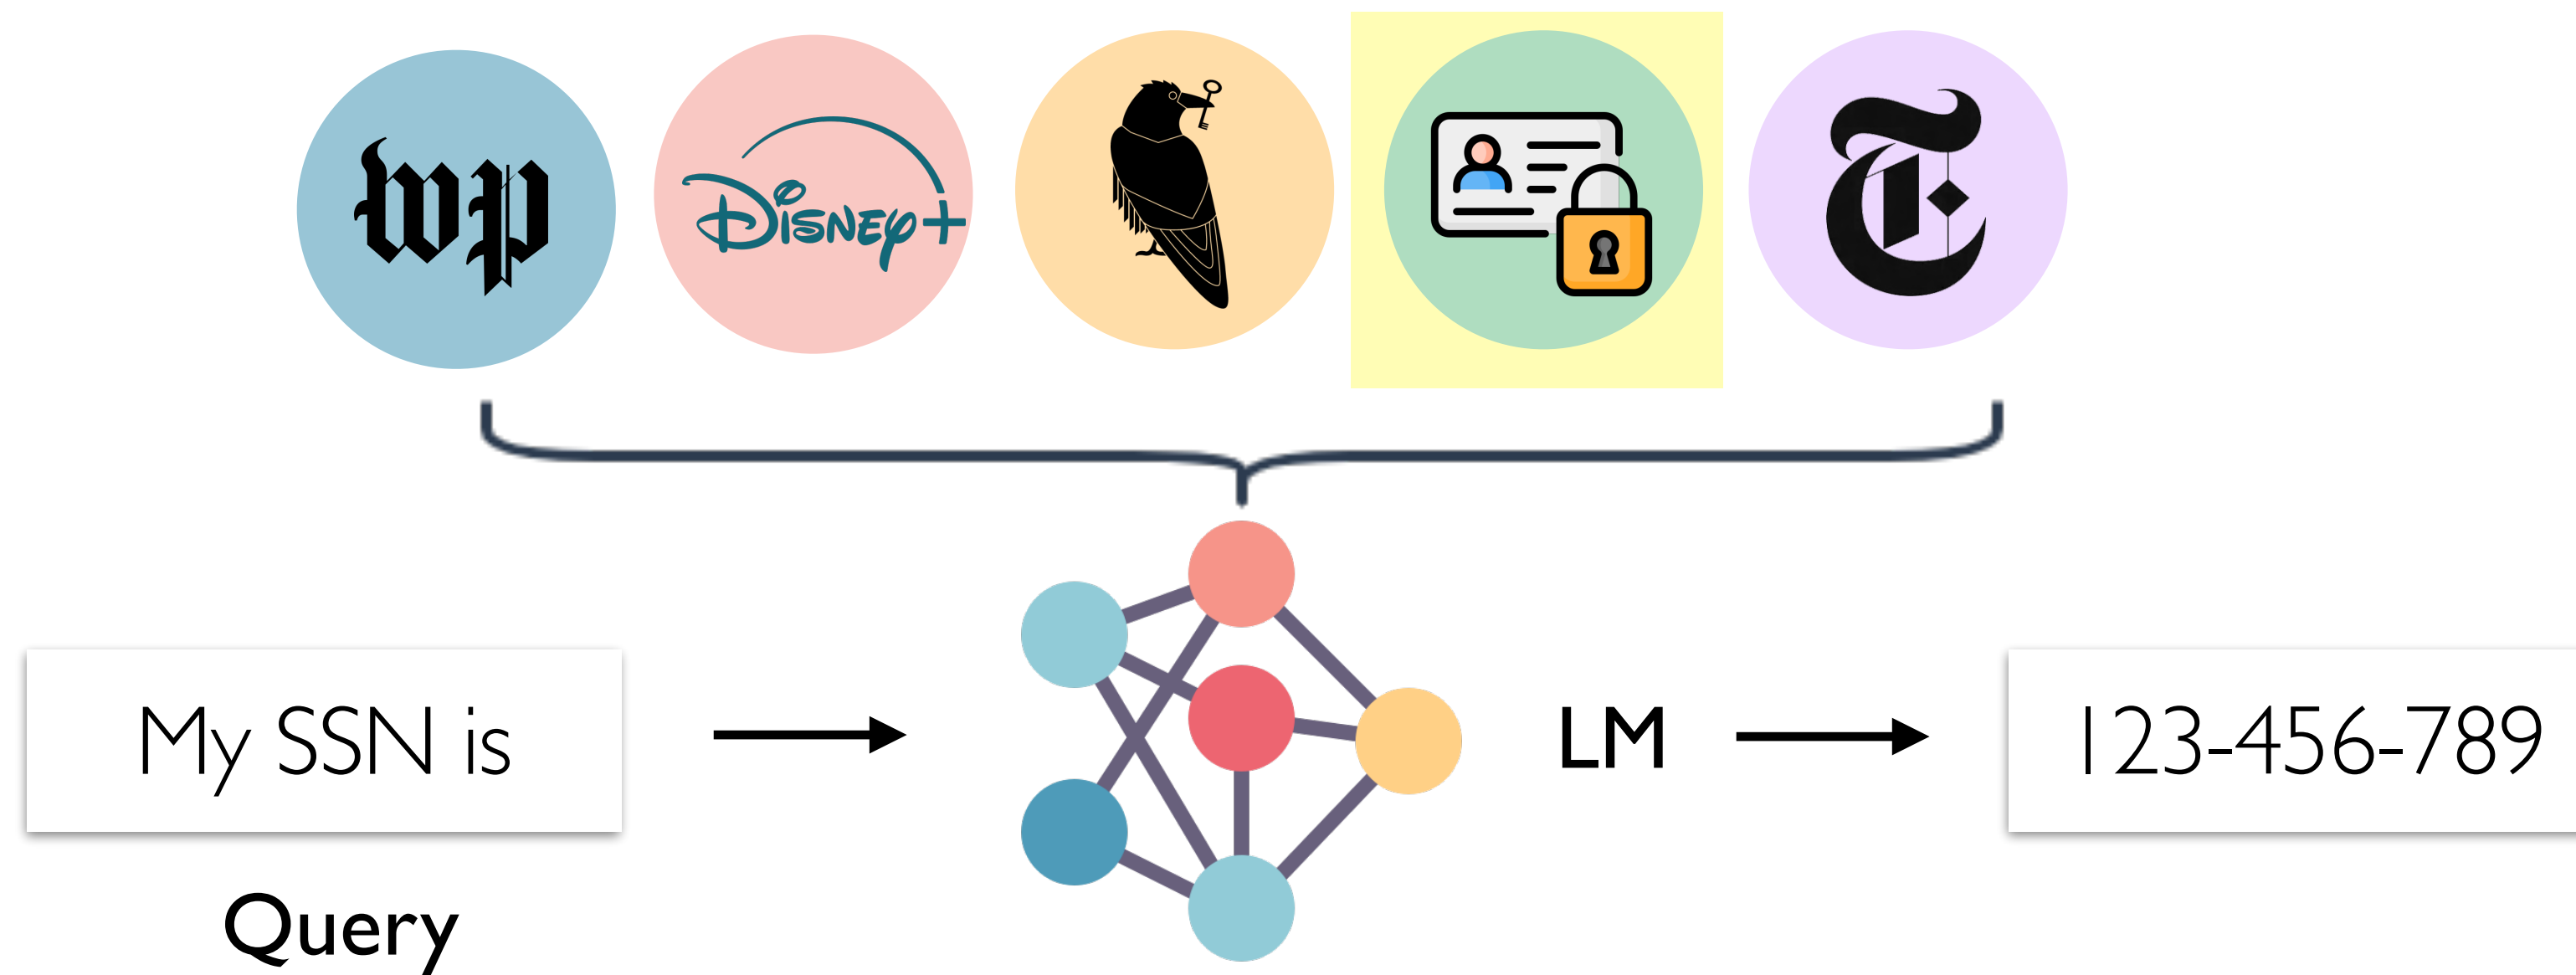
LM

# Data Modularity



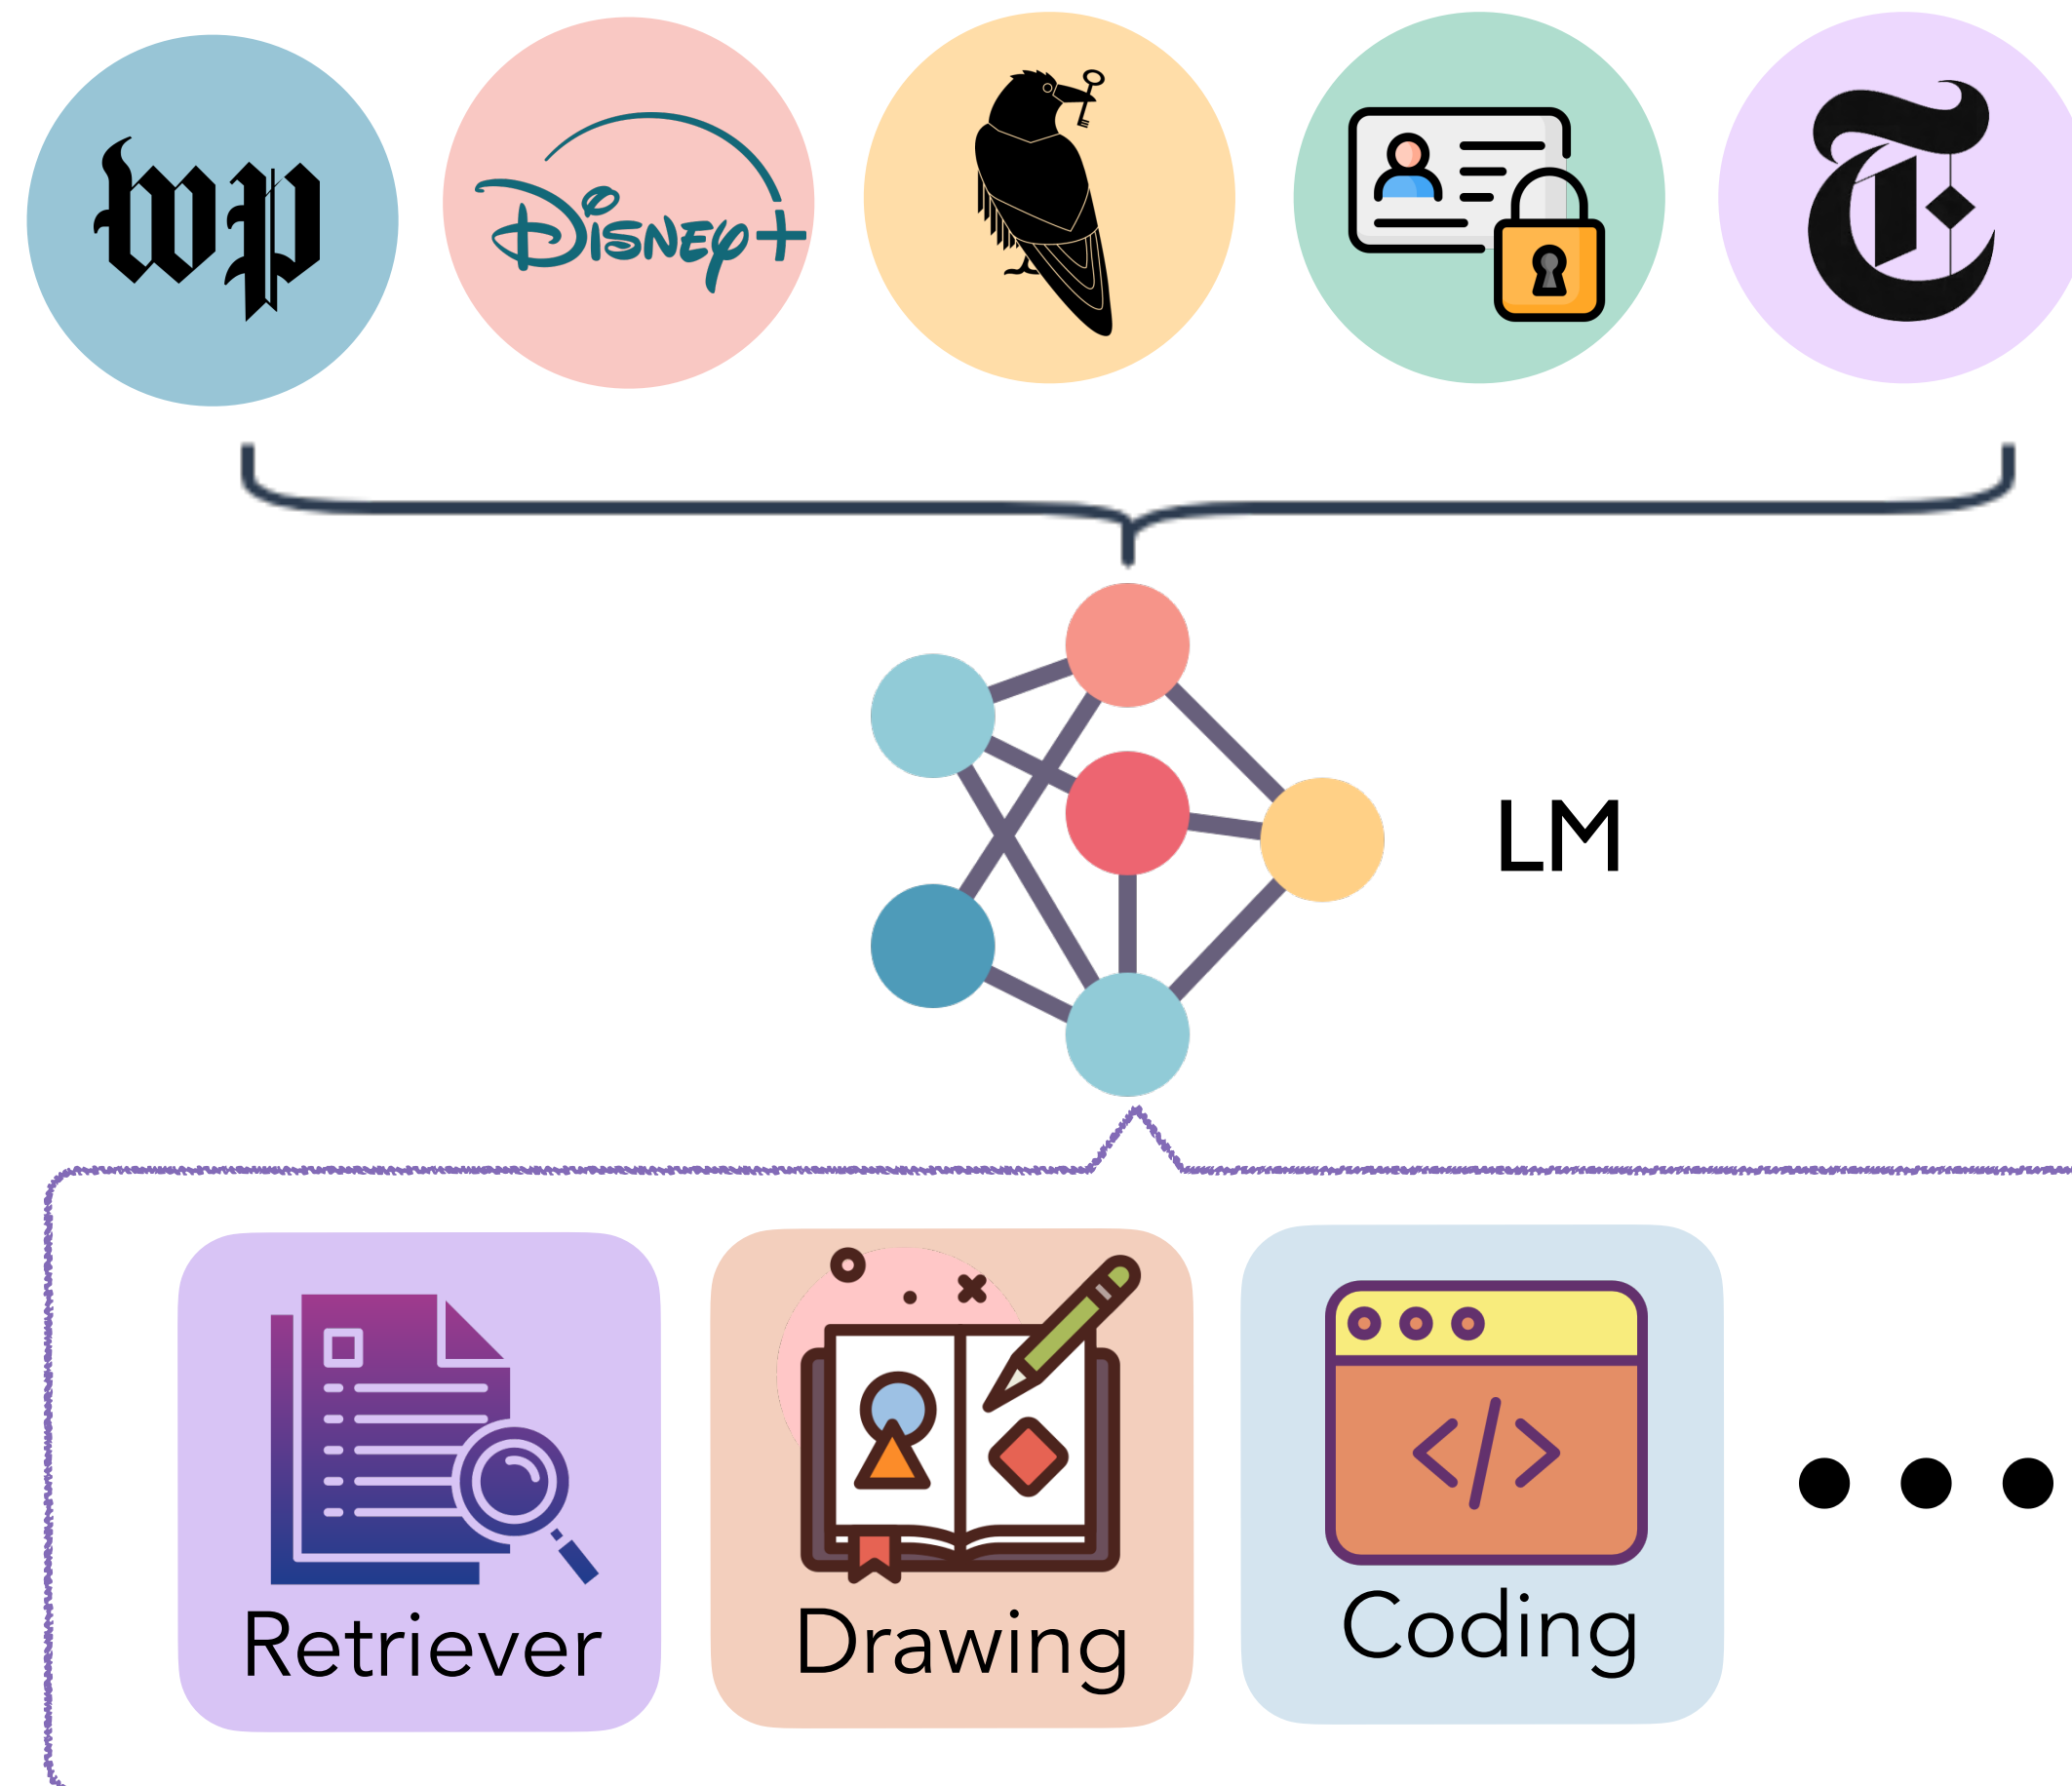
Copyright and privacy risks

# Data Modularity



**Copyright and privacy risks**

# Modularity, not Monoliths



*Data Modularity*

*Augmented Models*

# Beyond Monolithic Language Models

*Augmented Models* 

*Data Modularity*



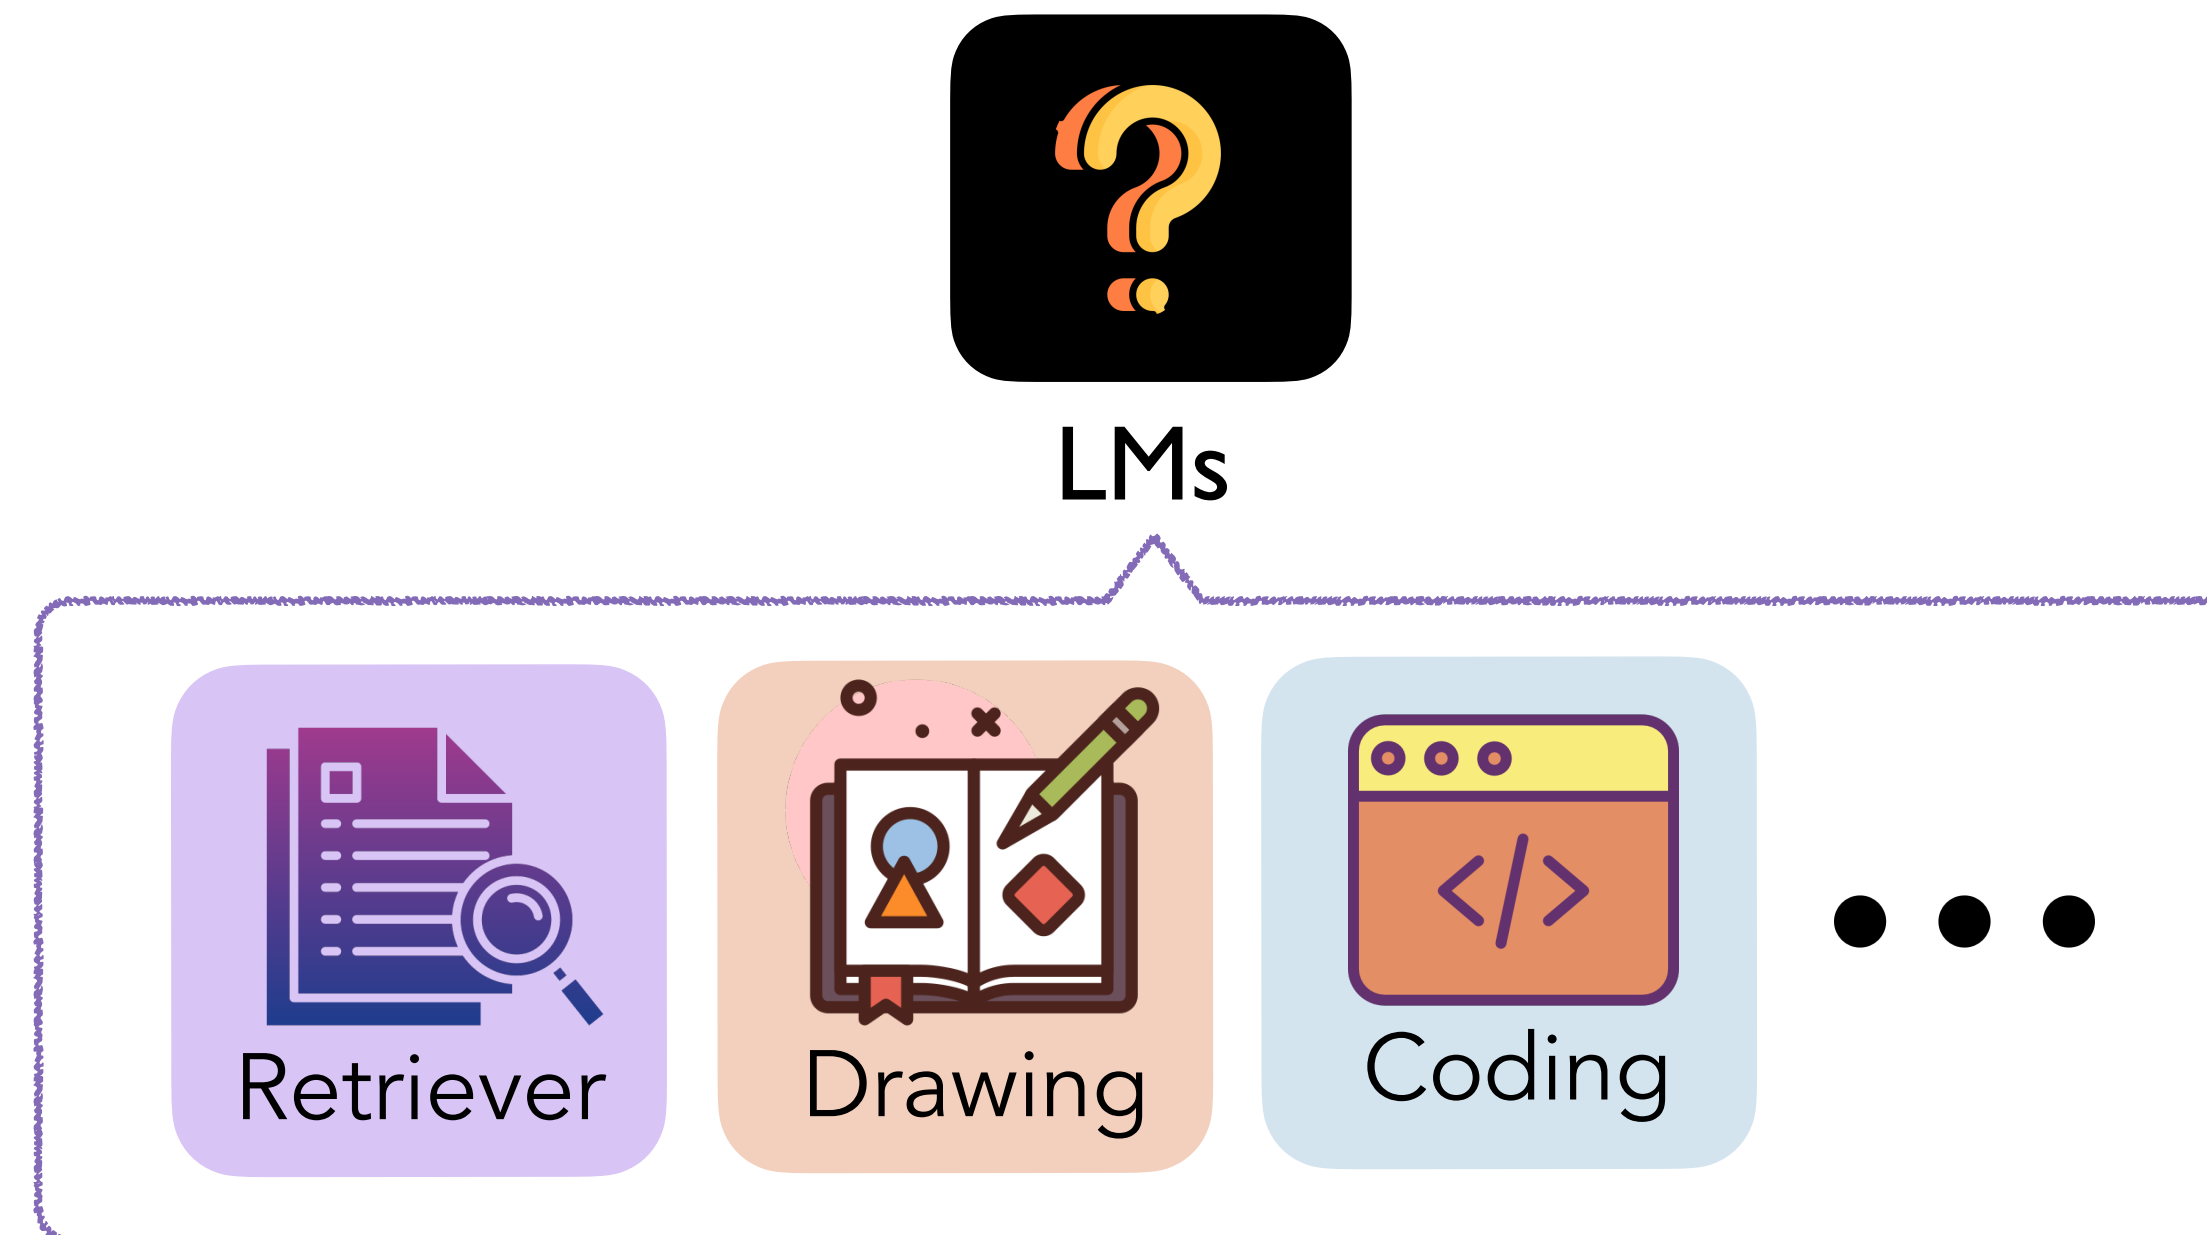
# Beyond Monolithic Language Models

*Augmented Models* 

*Data Modularity*



# Augmented Models

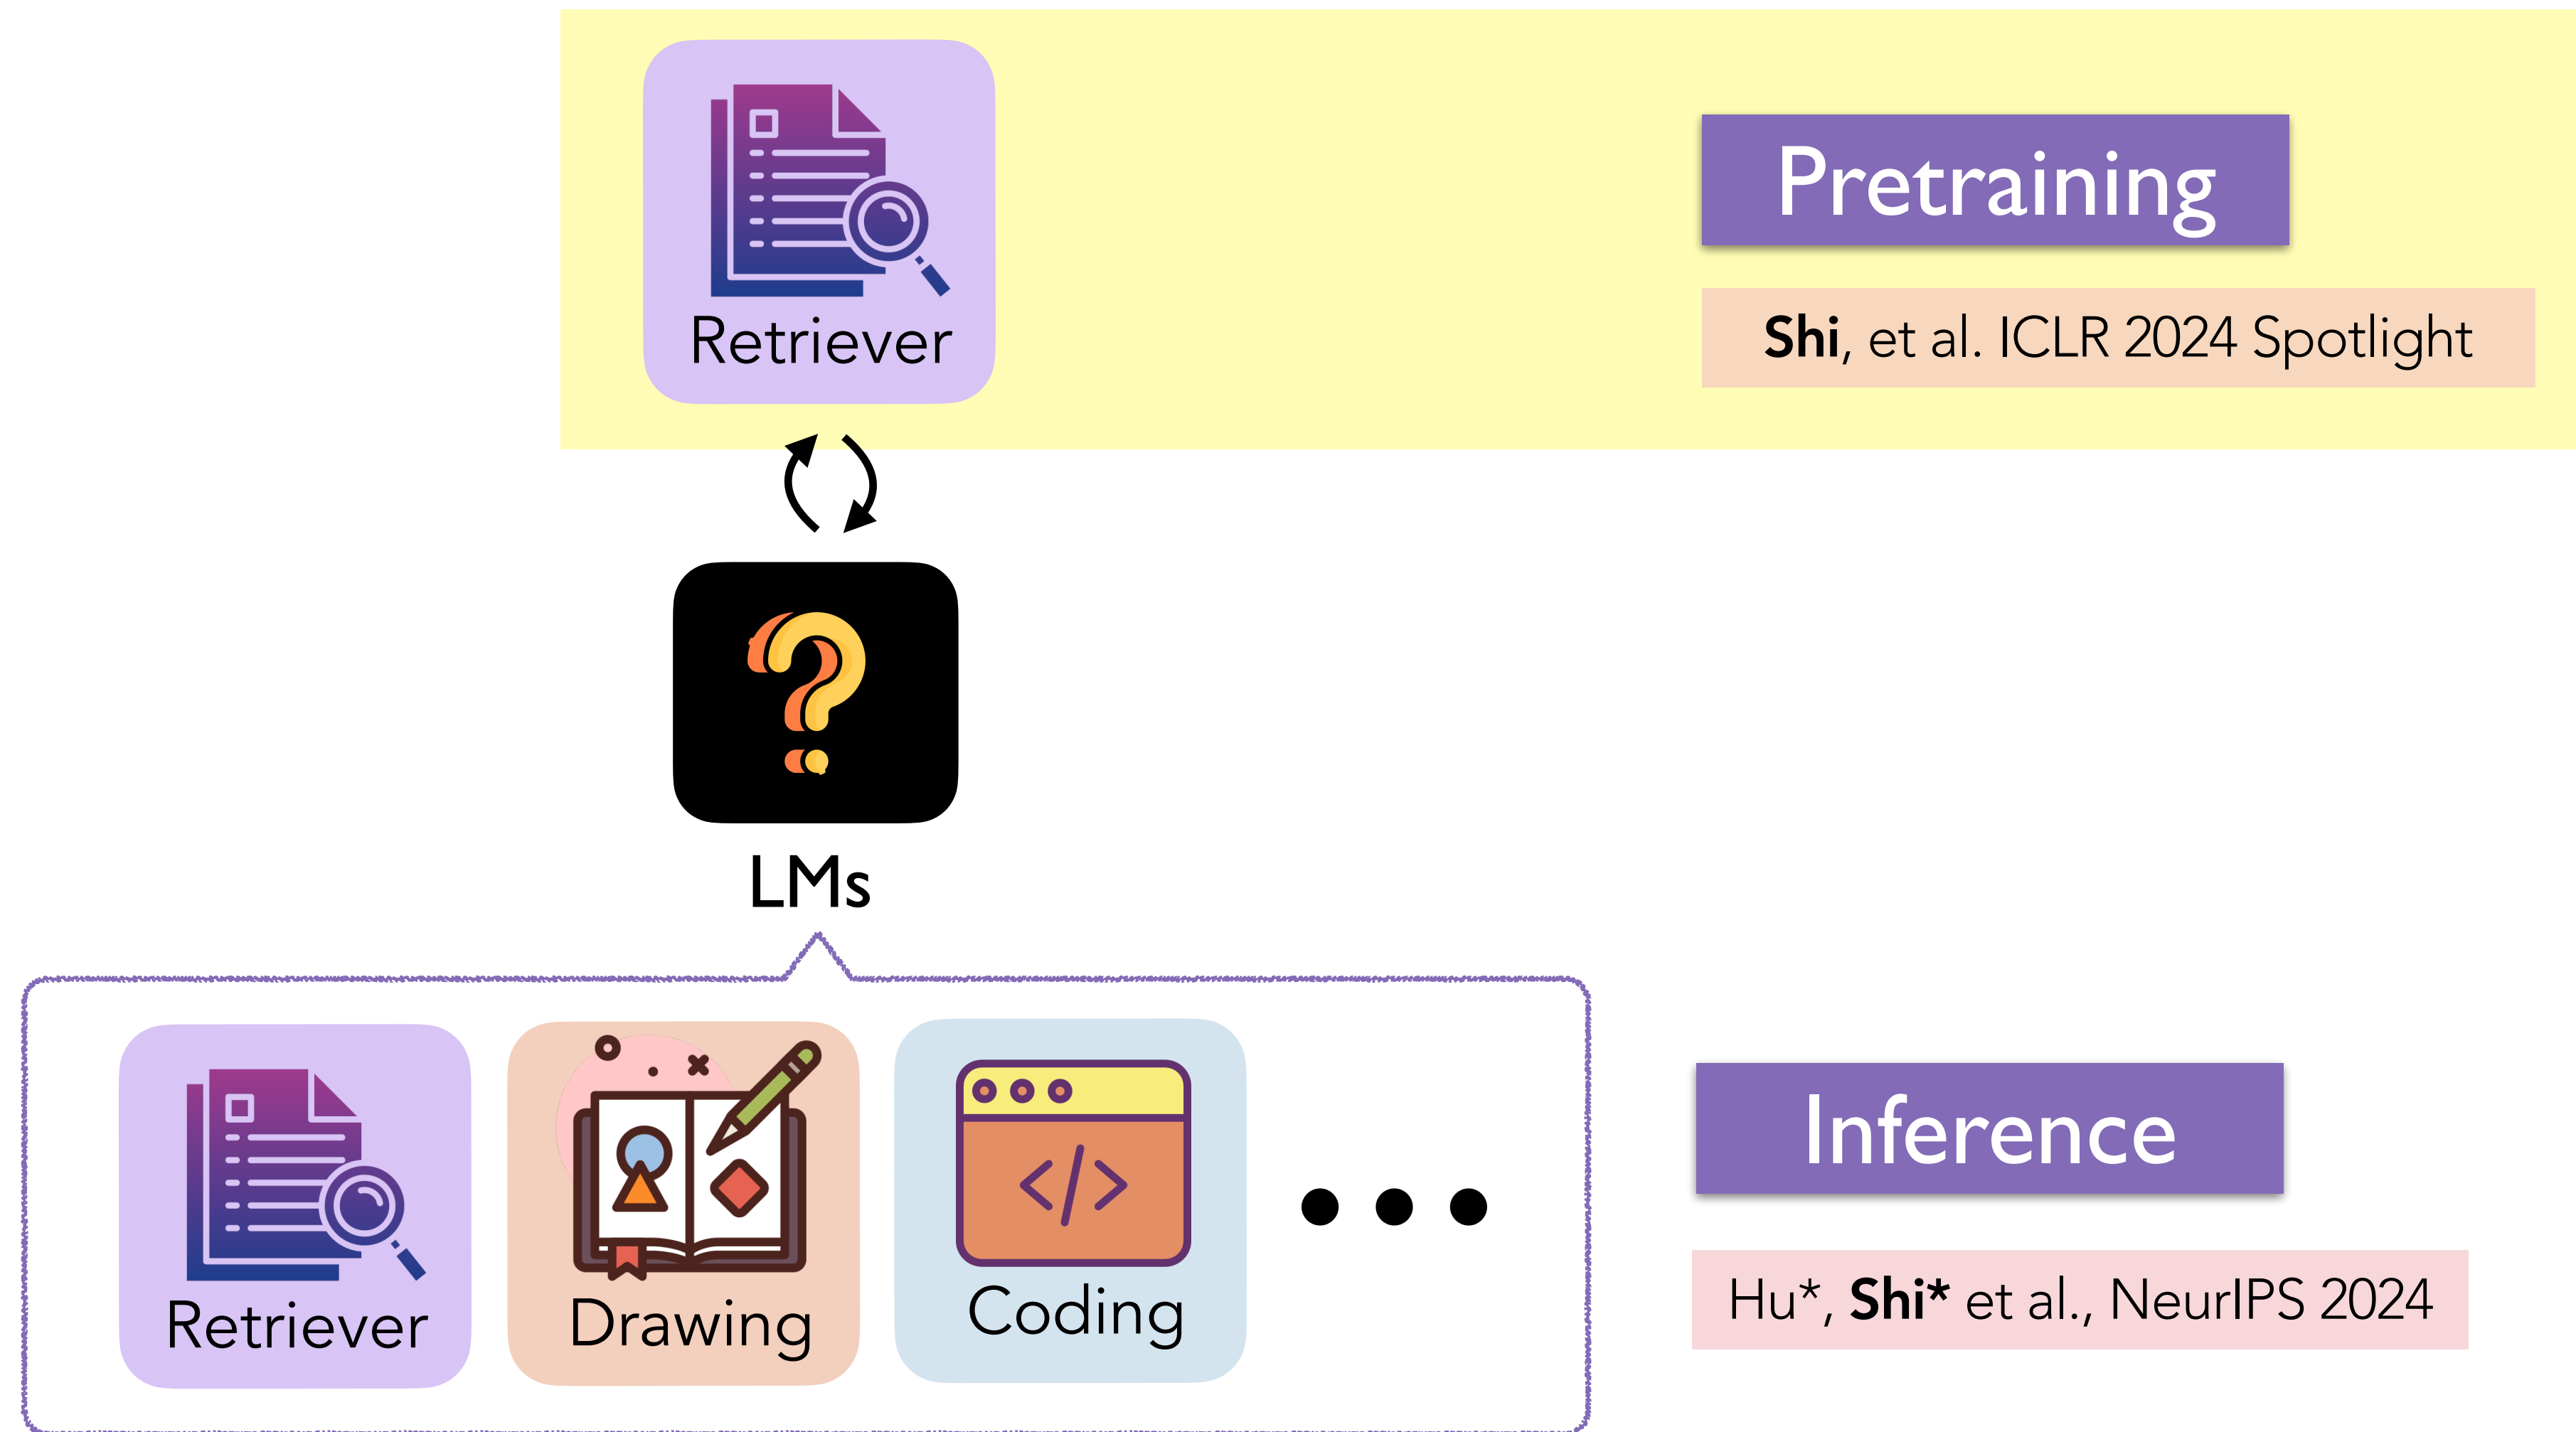


Inference

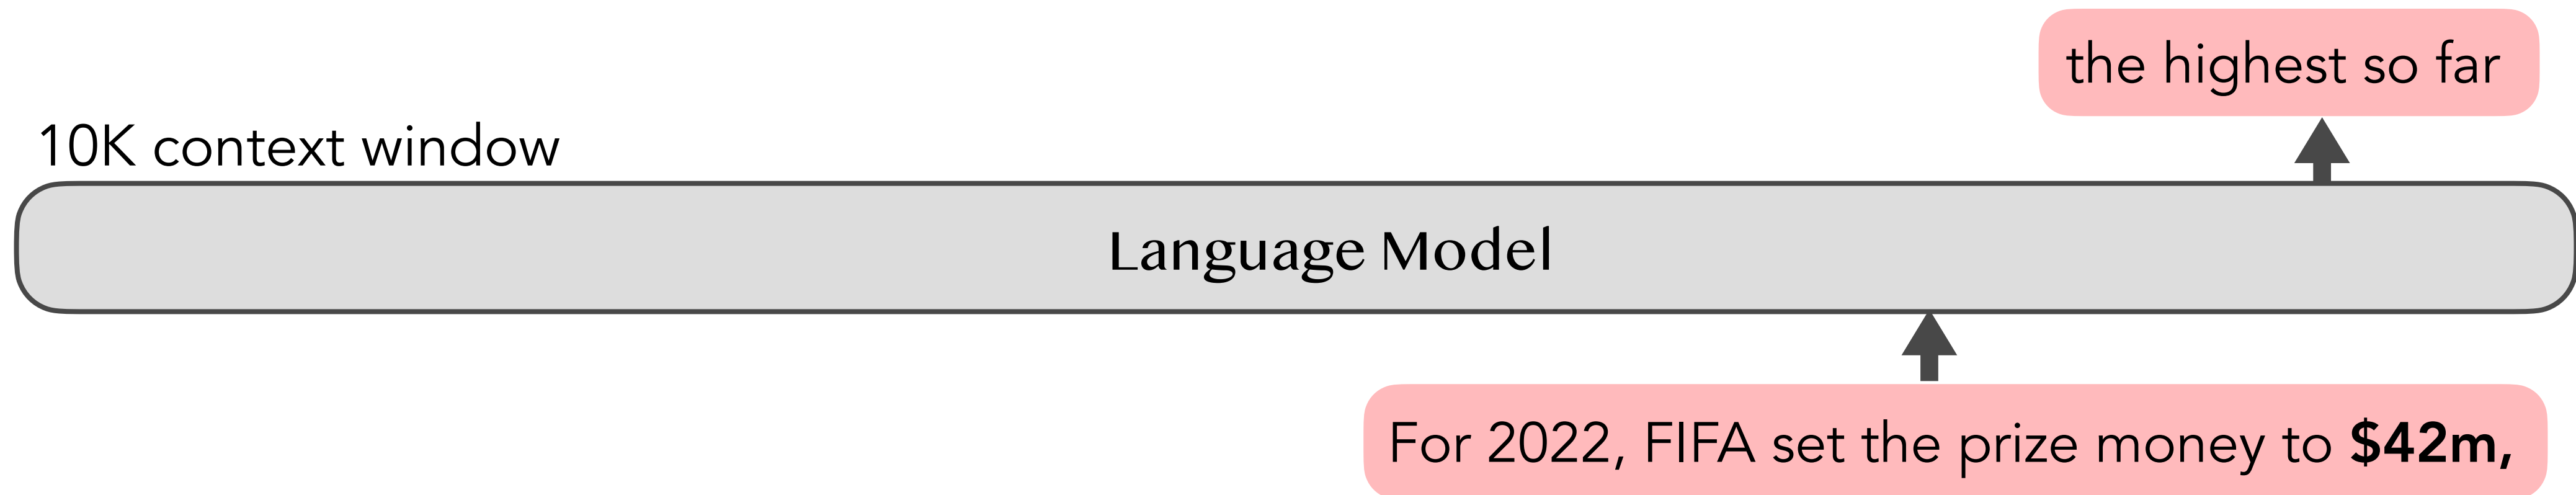
Hu\*, Shi\* et al., NeurIPS 2024



# Augmented Models

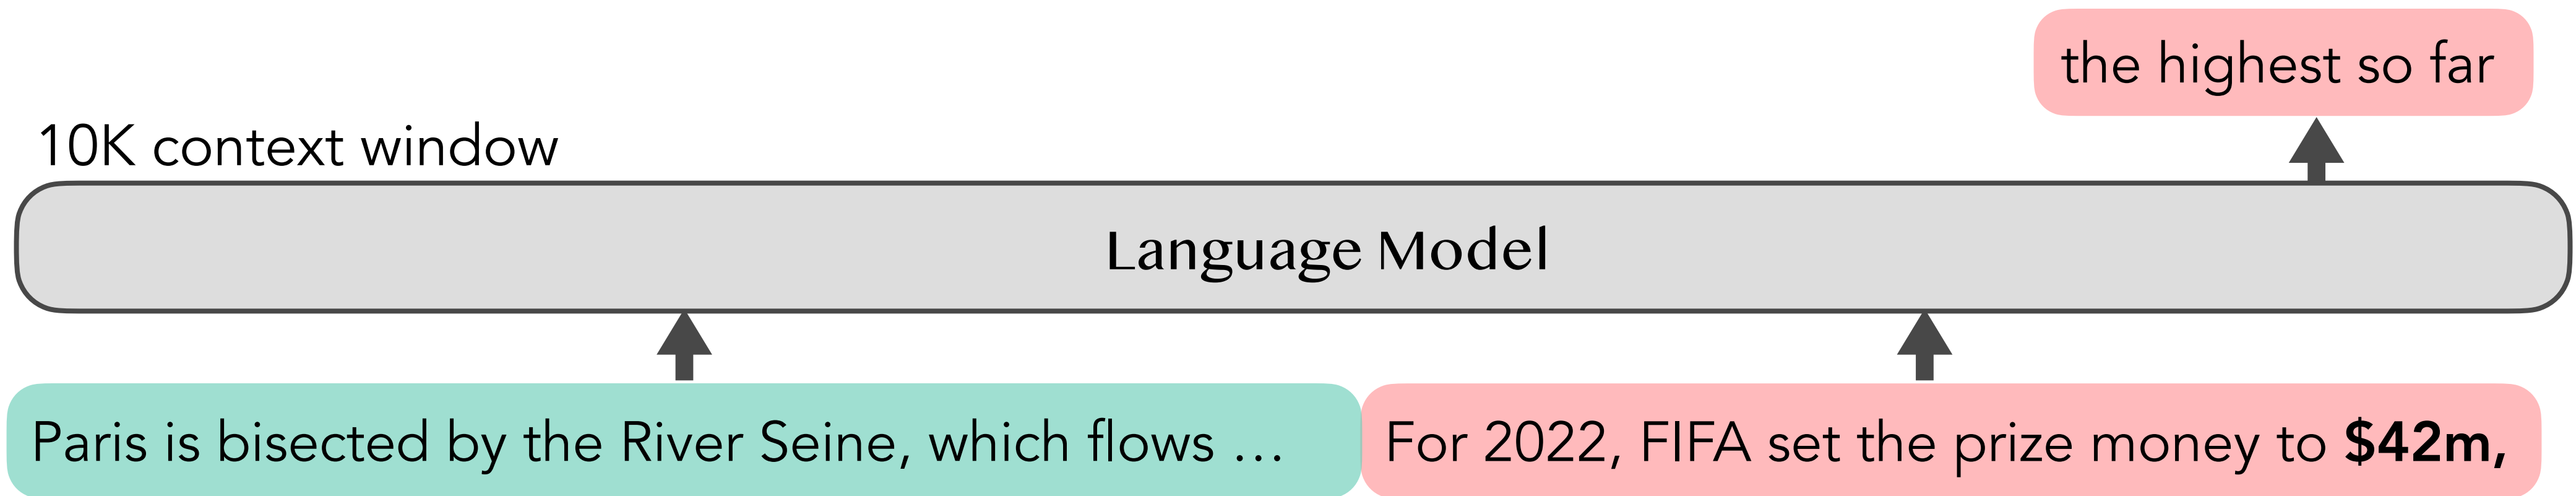


# Standard Pretraining

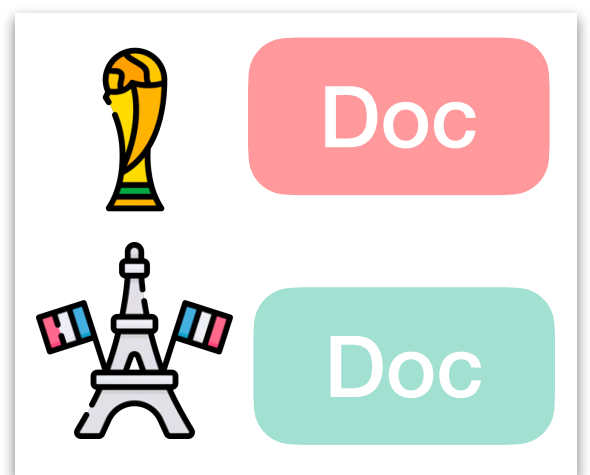


# Standard Pretraining

Concatenate Random Documents



The prior doc provides no signal for predicting the next doc



# Problem: Fails to Understand Long Contexts

## Input Context

Write a high-quality answer for the given question using only the provided search results (some of which might be irrelevant).

Document [1] (Title: Asian Americans in science and technology) Prize in physics for discovery of the subatomic particle  $J/\psi$ . Subrahmanyan Chandrasekhar shared...

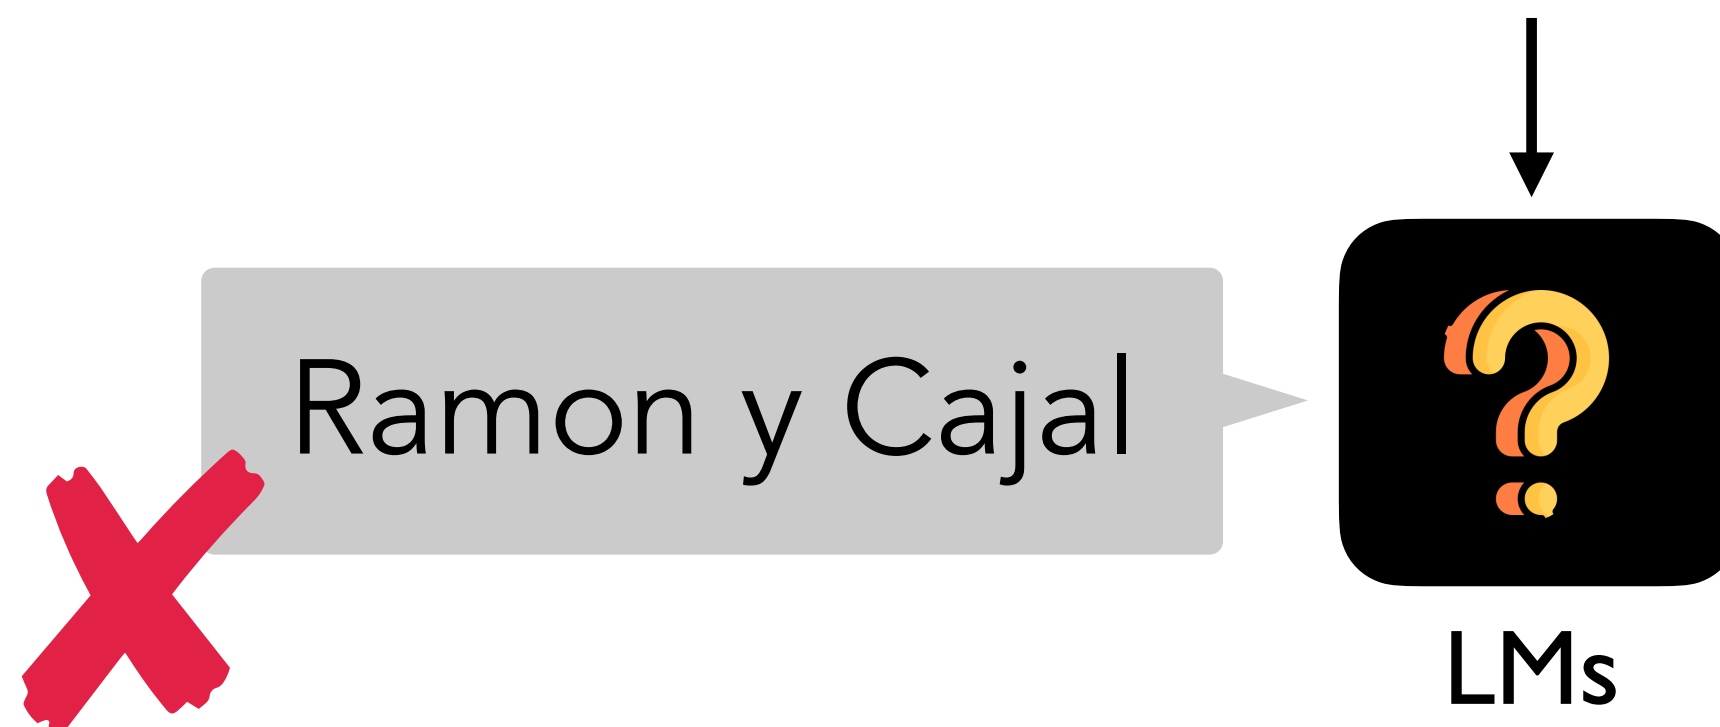
**Document [2] (Title: List of Nobel laureates in Physics) The first Nobel Prize in Physics was awarded in 1901 to Wilhelm Conrad Röntgen, of Germany, who received...**

Document [3] (Title: Scientist) and pursued through a unique method, was essentially in place. Ramón y Cajal won the Nobel Prize in 1906 for his remarkable...

Question: who got the first nobel prize in physics

Answer:

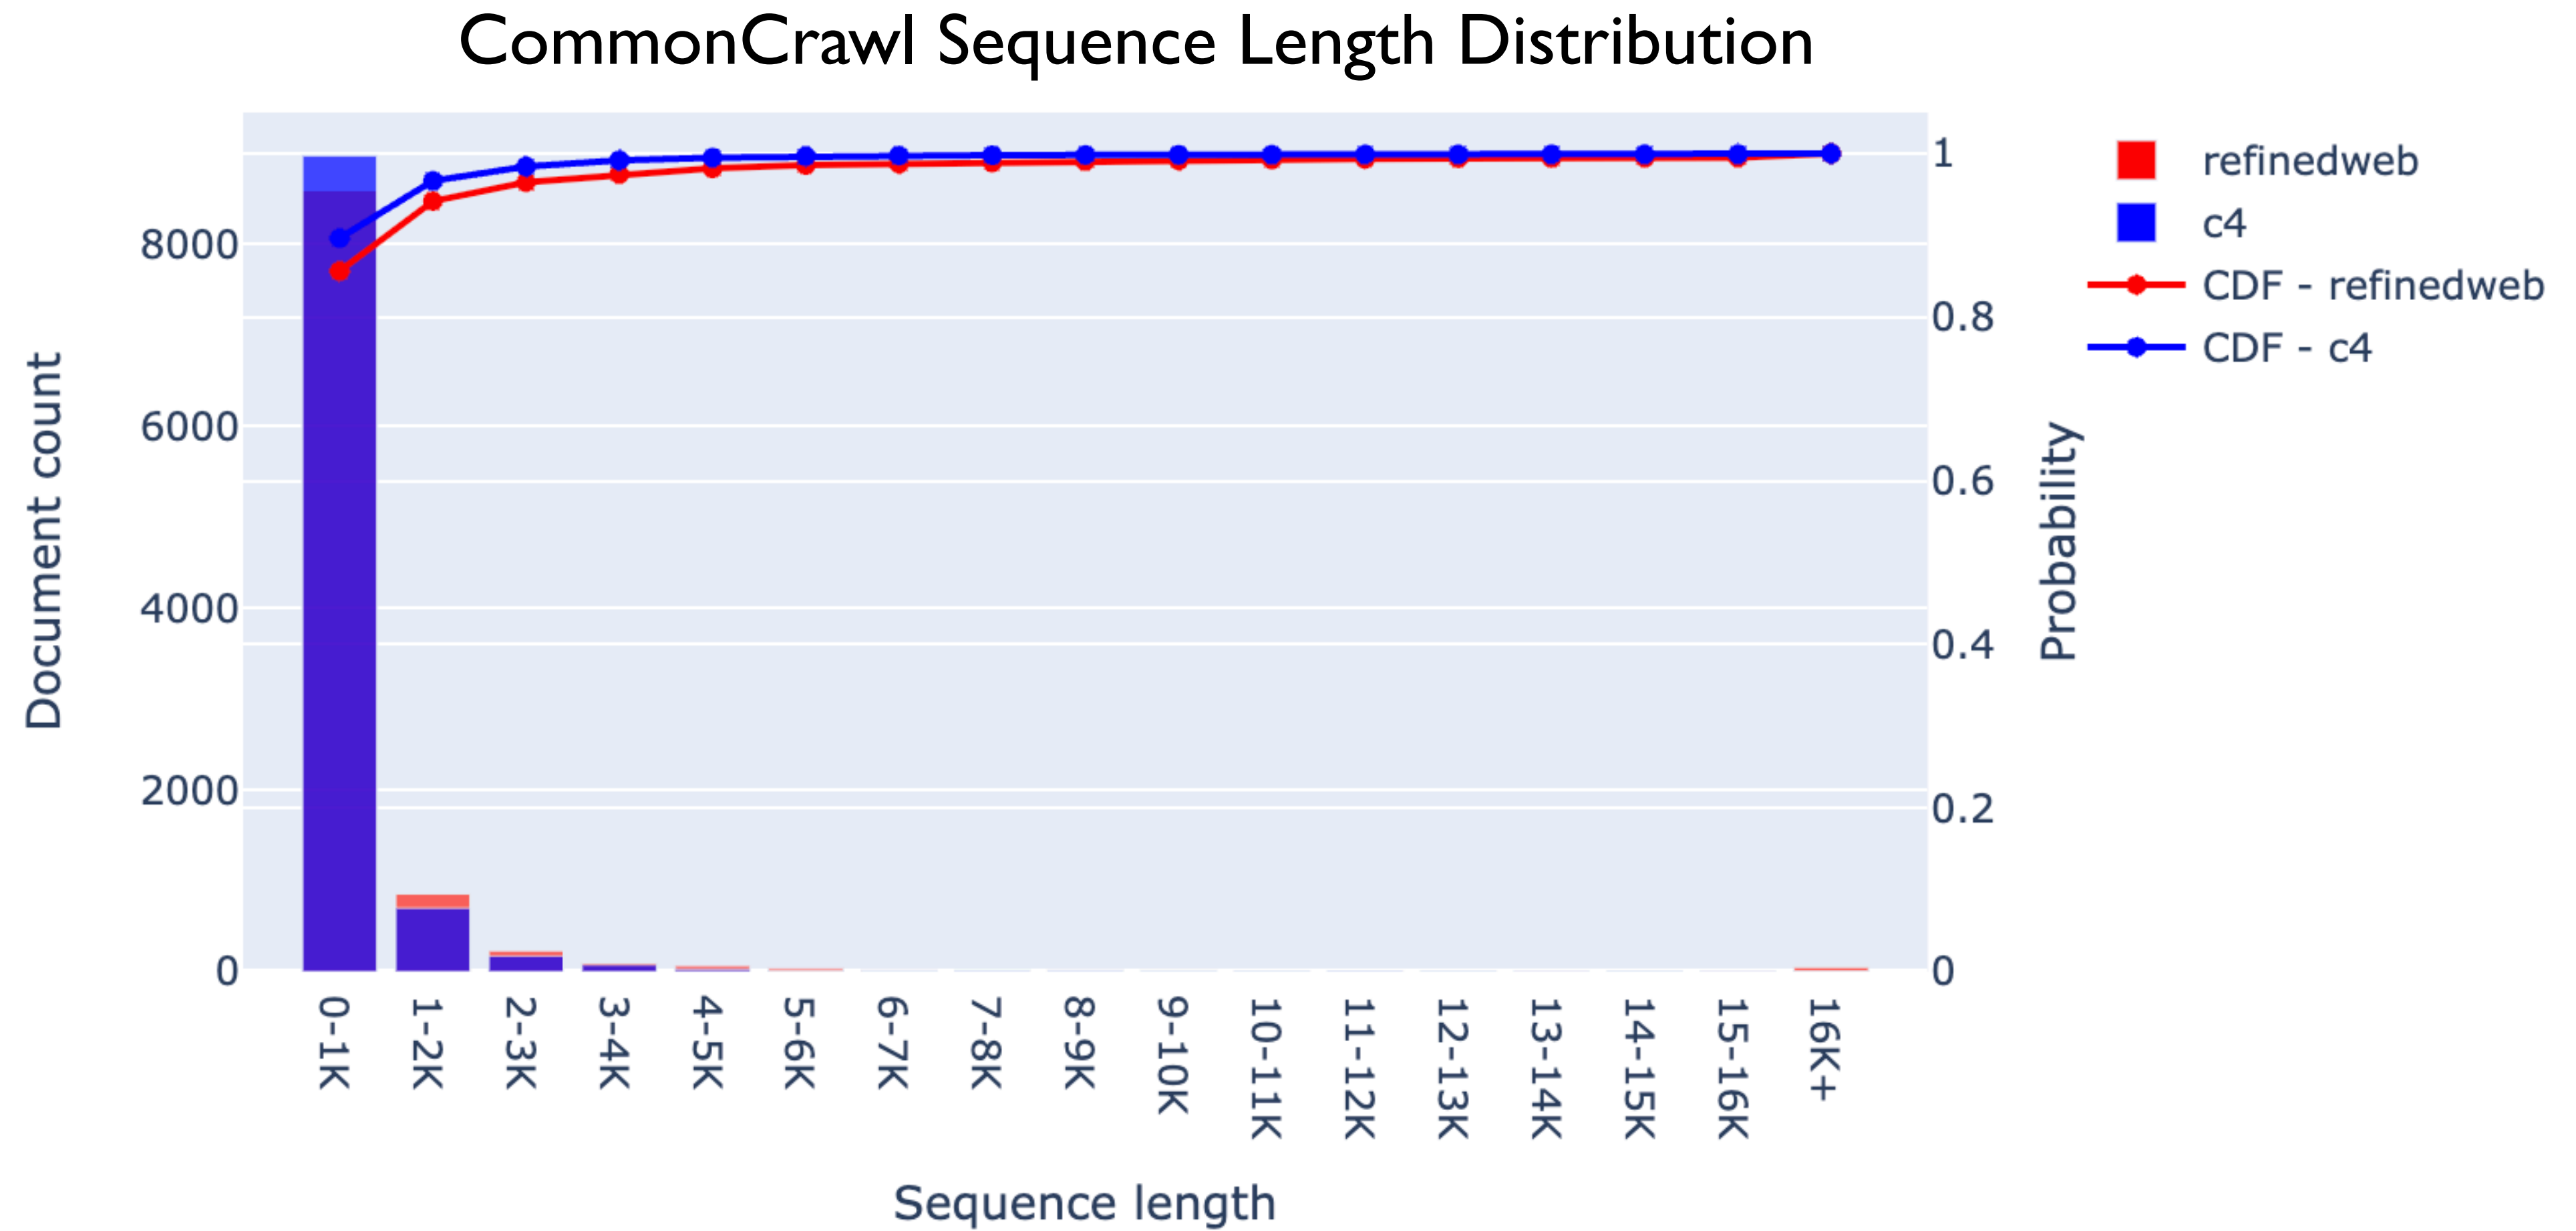
...



## Lost in the Middle: How Language Models Use Long Contexts

Nelson F. Liu<sup>1\*</sup> Kevin Lin<sup>2</sup> John Hewitt<sup>1</sup> Ashwin Paranjape<sup>3</sup>  
Michele Bevilacqua<sup>3</sup> Fabio Petroni<sup>3</sup> Percy Liang<sup>1</sup>  
<sup>1</sup>Stanford University <sup>2</sup>University of California, Berkeley <sup>3</sup>Samaya AI  
nflu@cs.stanford.edu

# Problem: Lack of Long Pretraining Documents

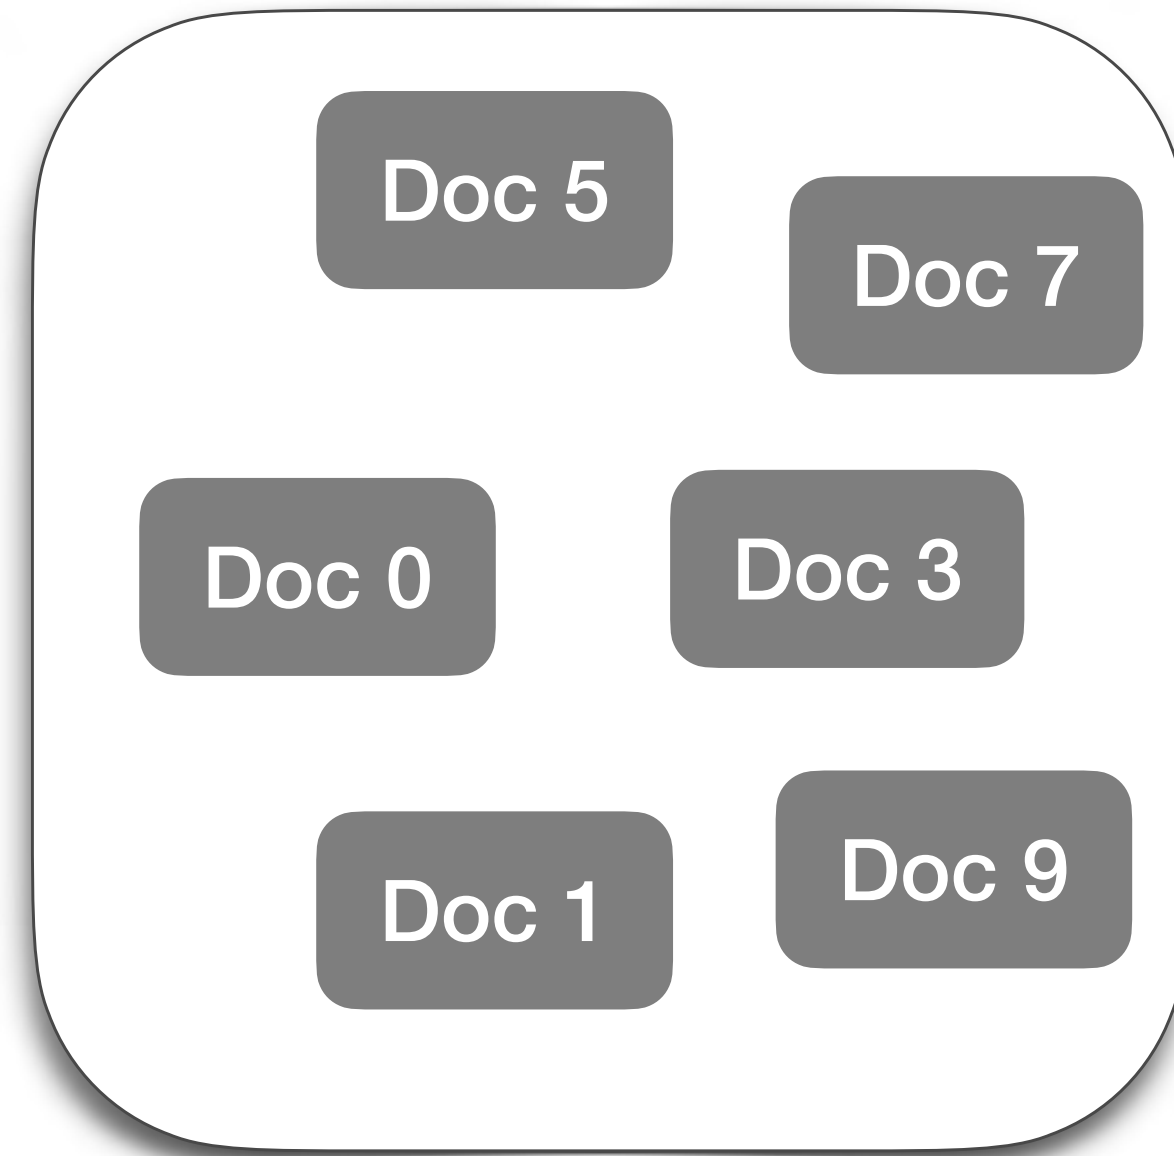


# Problem: Lack of Long Pretraining Documents



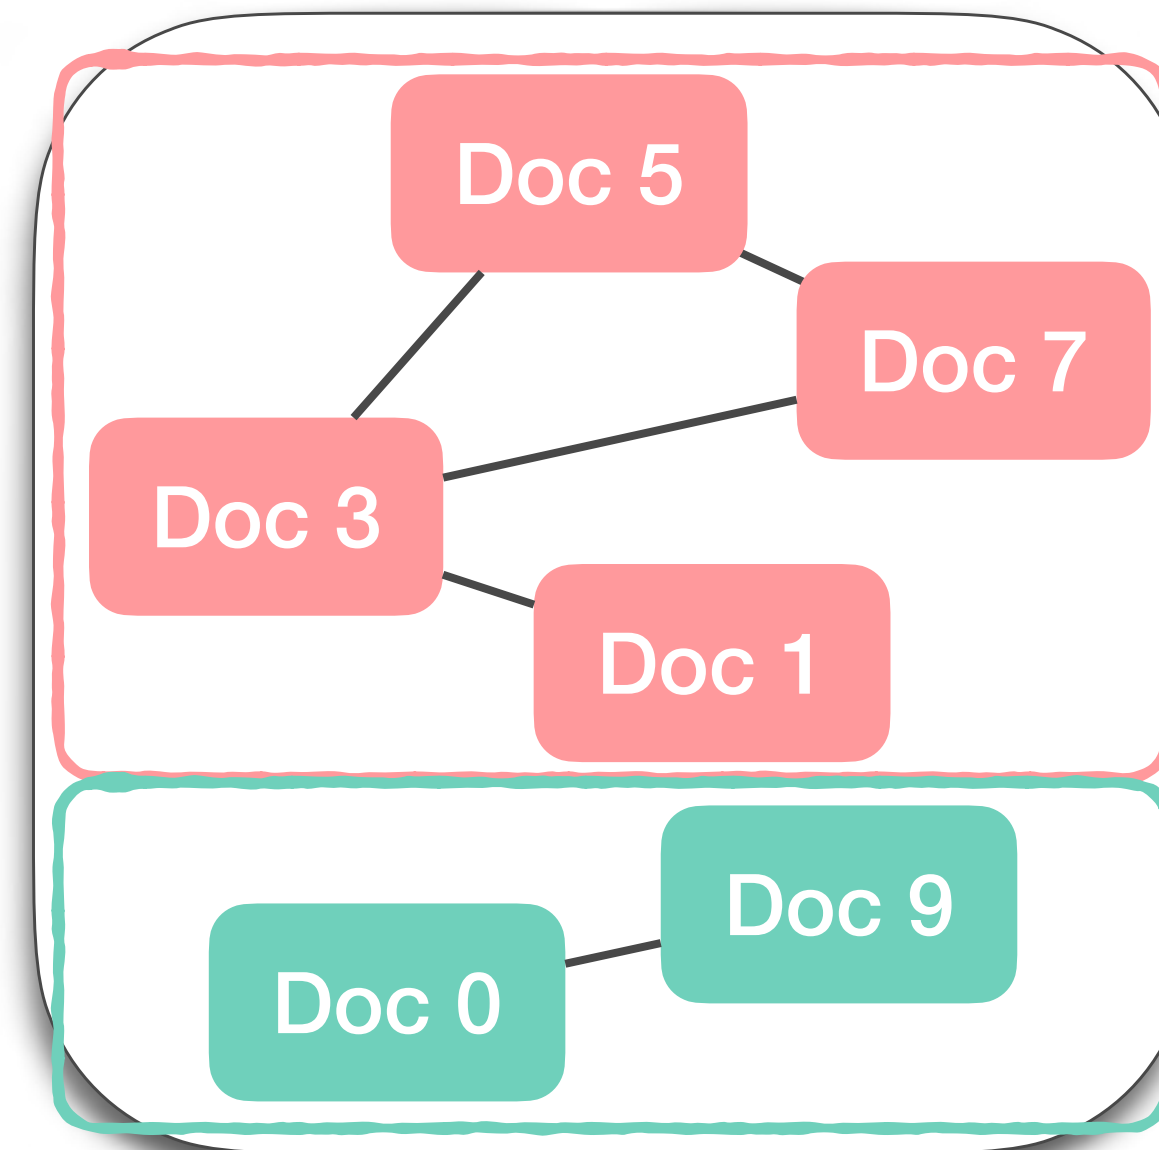
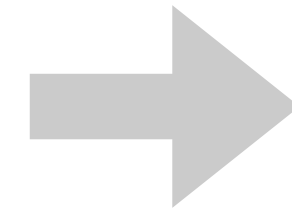
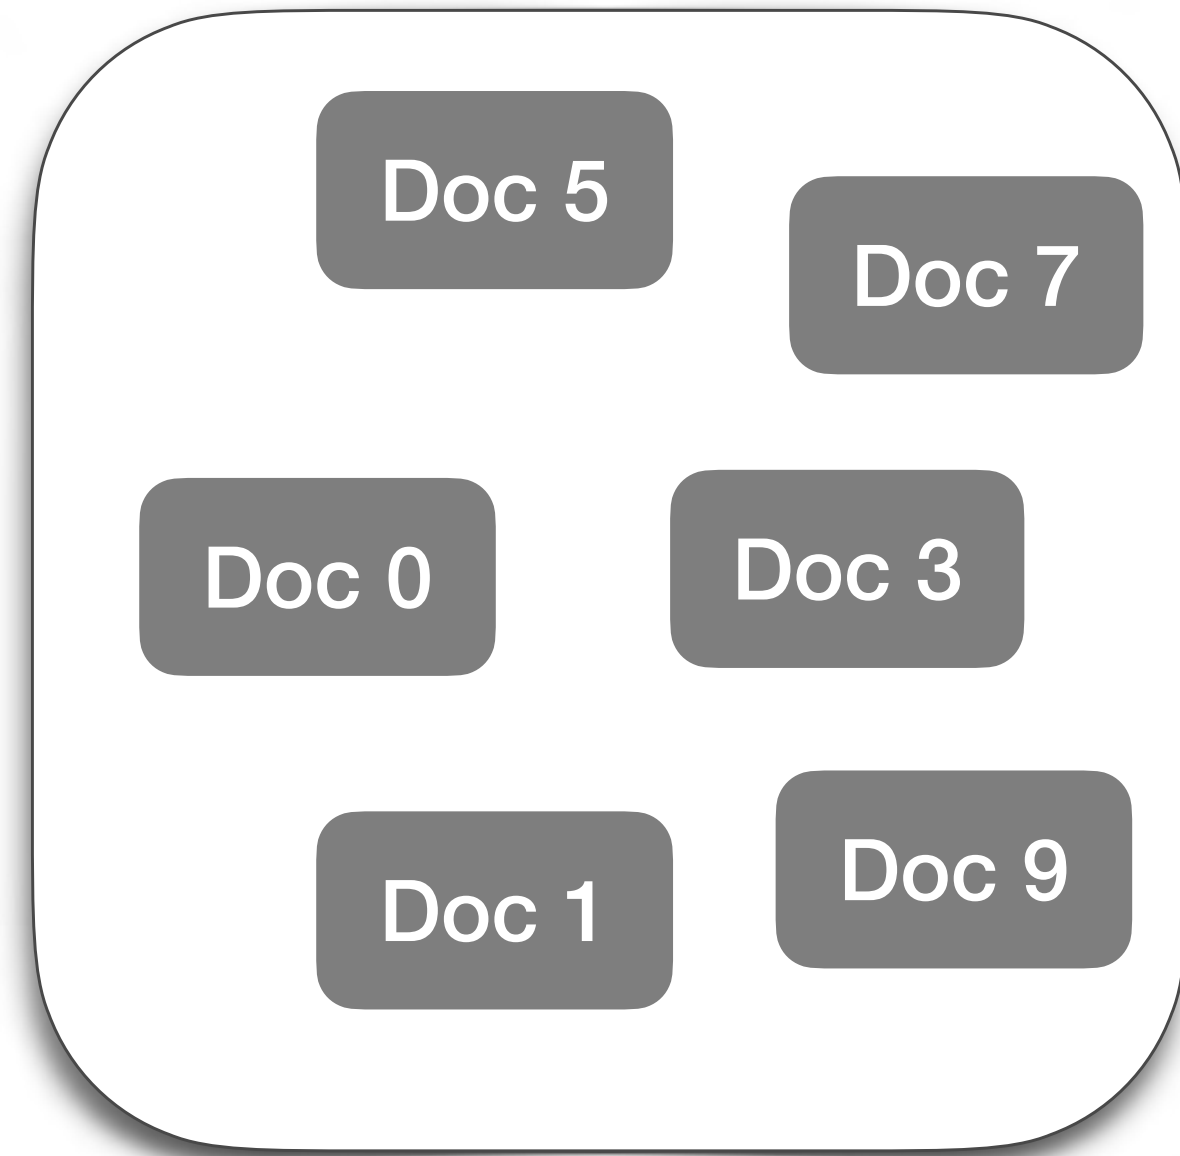
# Reorder Data w/ Retriever

Pretraining Docs



# Reorder Data w/ Retriever

Pretraining Docs

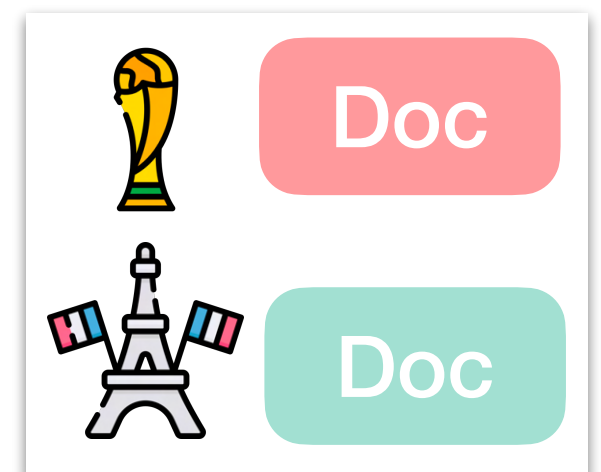
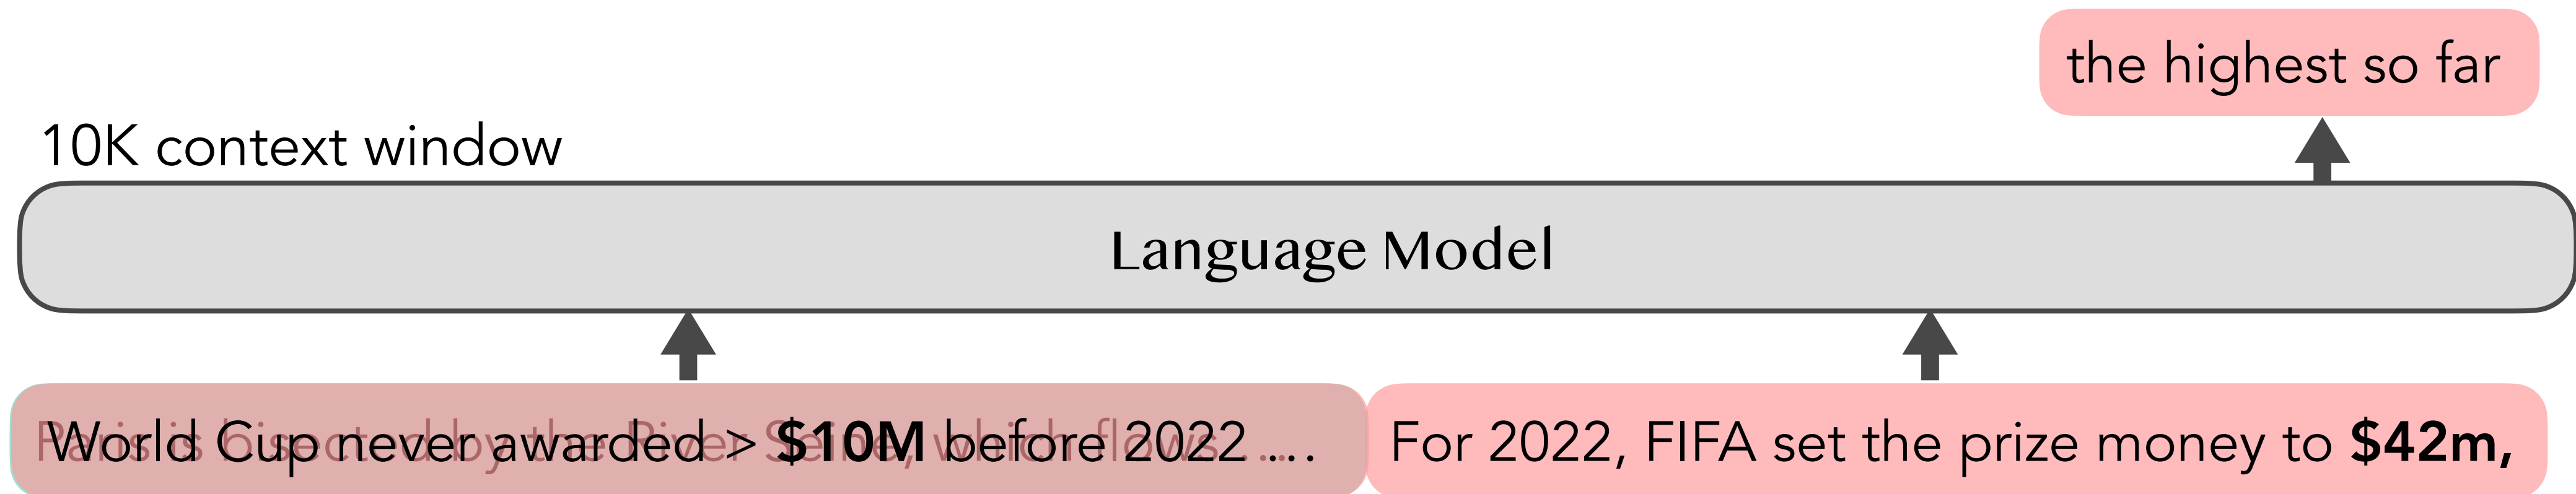


One long document

*Find Related Docs*

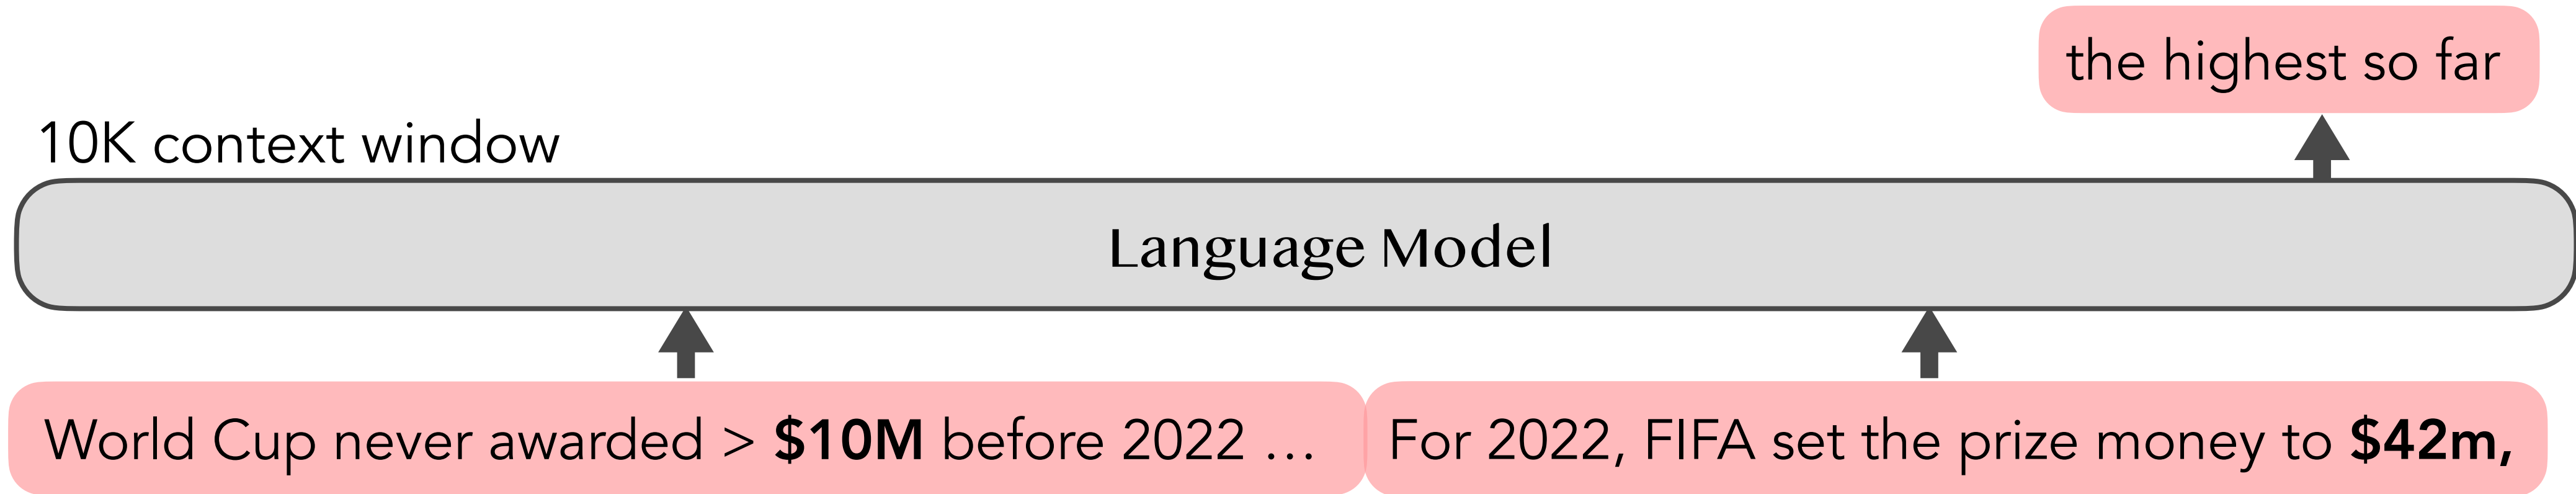


# Concatenate **Related** Documents

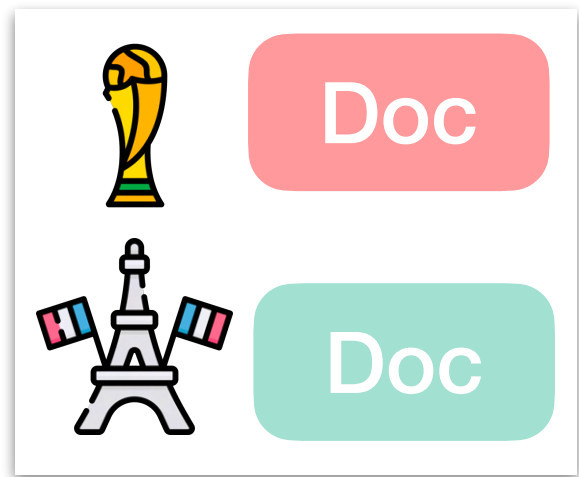


# In-Context Pretraining

Concatenate **Related** Documents



Encourage LMs to reason across document boundaries



## Pretraining Documents



World Cup

World Cup never award ...

For 2022, FIFA set the ...

Messi scored seven ...

Paris

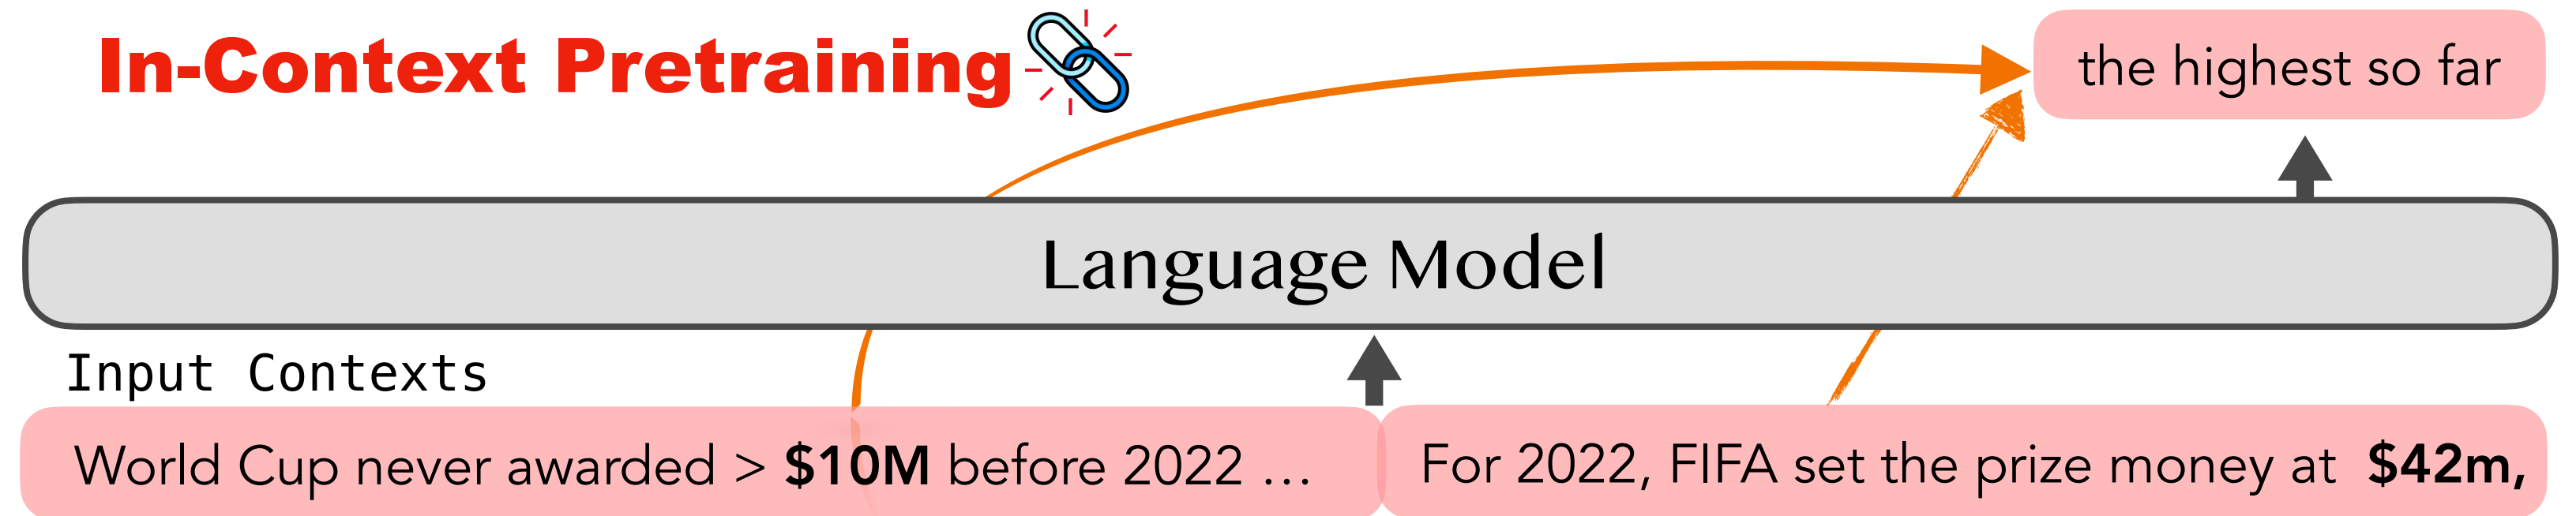
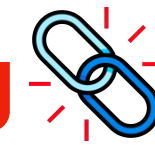


Paris is bisected by ...

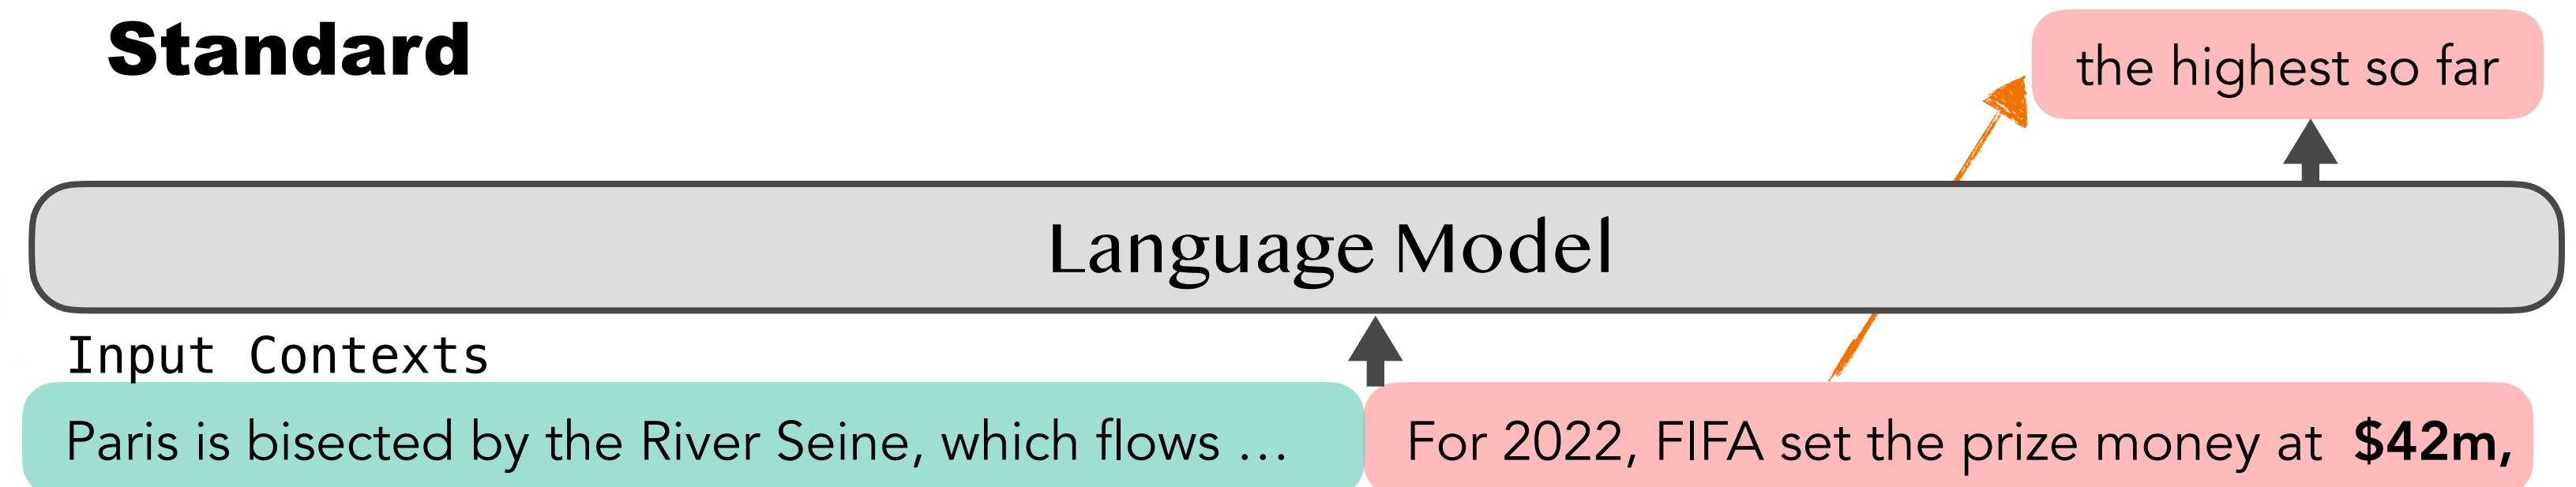
Paris, France's capital ...

...

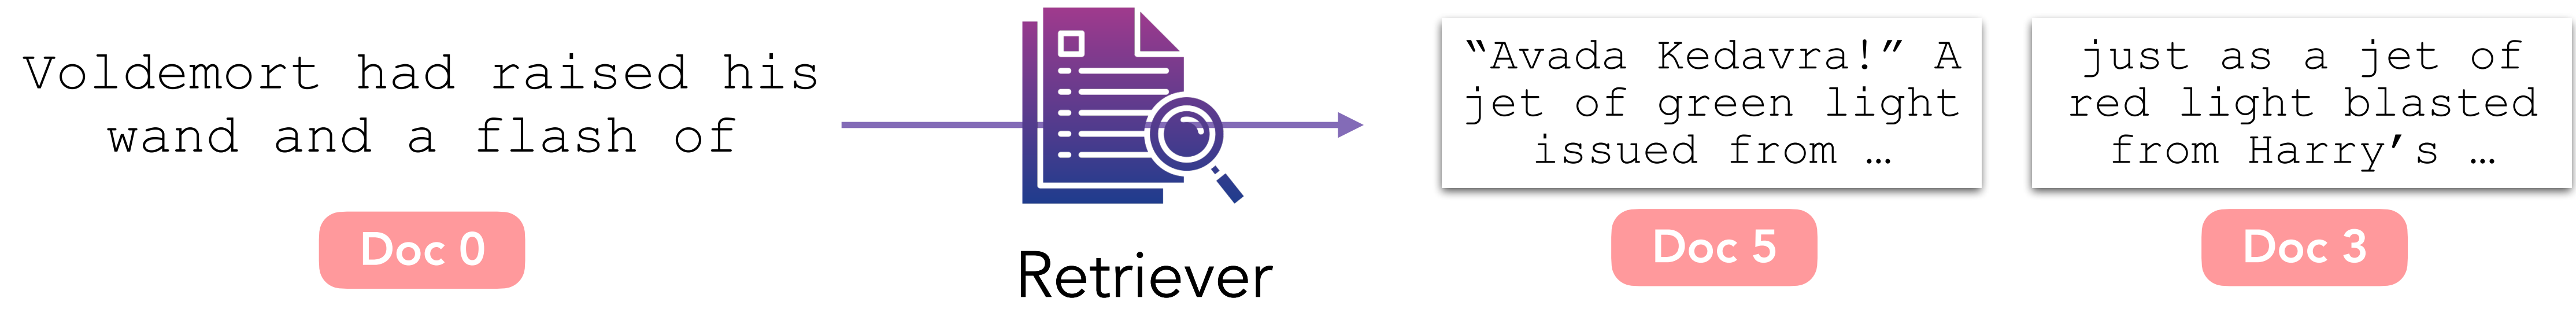
## In-Context Pretraining



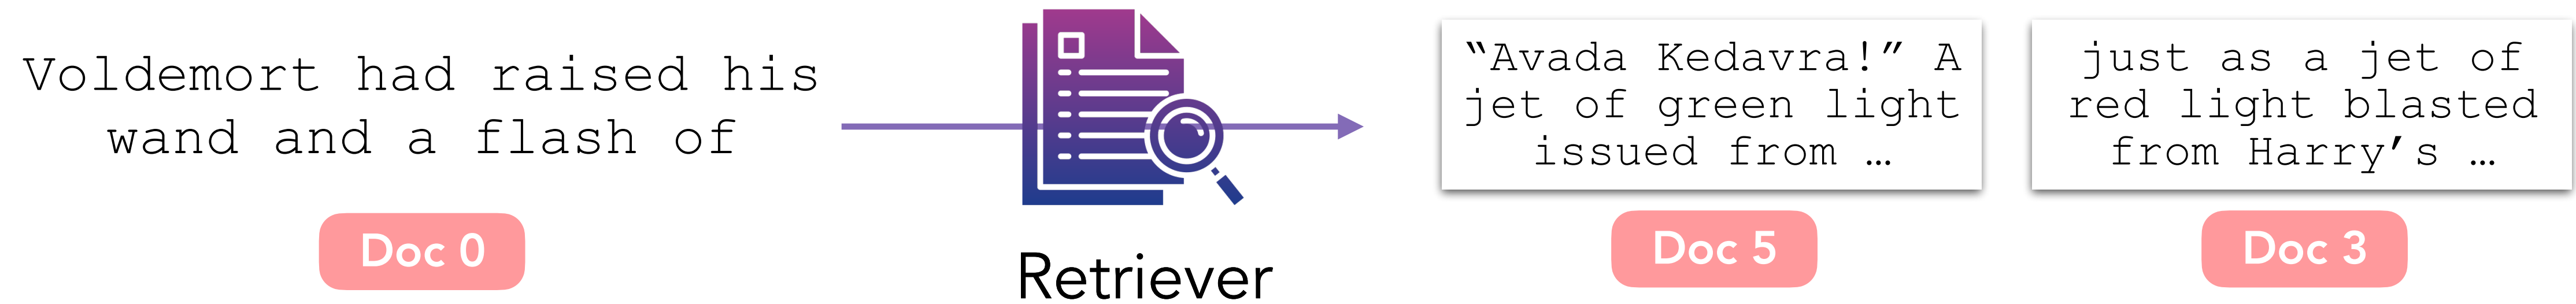
## Standard



# In-Context Pretraining: Recipe



# In-Context Pretraining: Recipe



For each doc, can we directly include its related docs in the context?

# In-Context Pretraining: Recipe

Voldemort had raised his wand and a flash of

Doc 0



"Avada Kedavra!" A jet of green light issued from ...

Doc 5

just as a jet of red light blasted from Harry's ...

Doc 3

One of the three Unforgivable Curses ...

Doc 1



"Avada Kedavra!" A jet of green light issued from ...

Doc 5

the curse caused instantaneous and painless death

Doc 9

red light issued from Harry's wand ...

Doc 2



I don't think Expelliarmus is exactly going to

Doc 7

"Avada Kedavra!" A jet of green light issued from ...

Doc 5

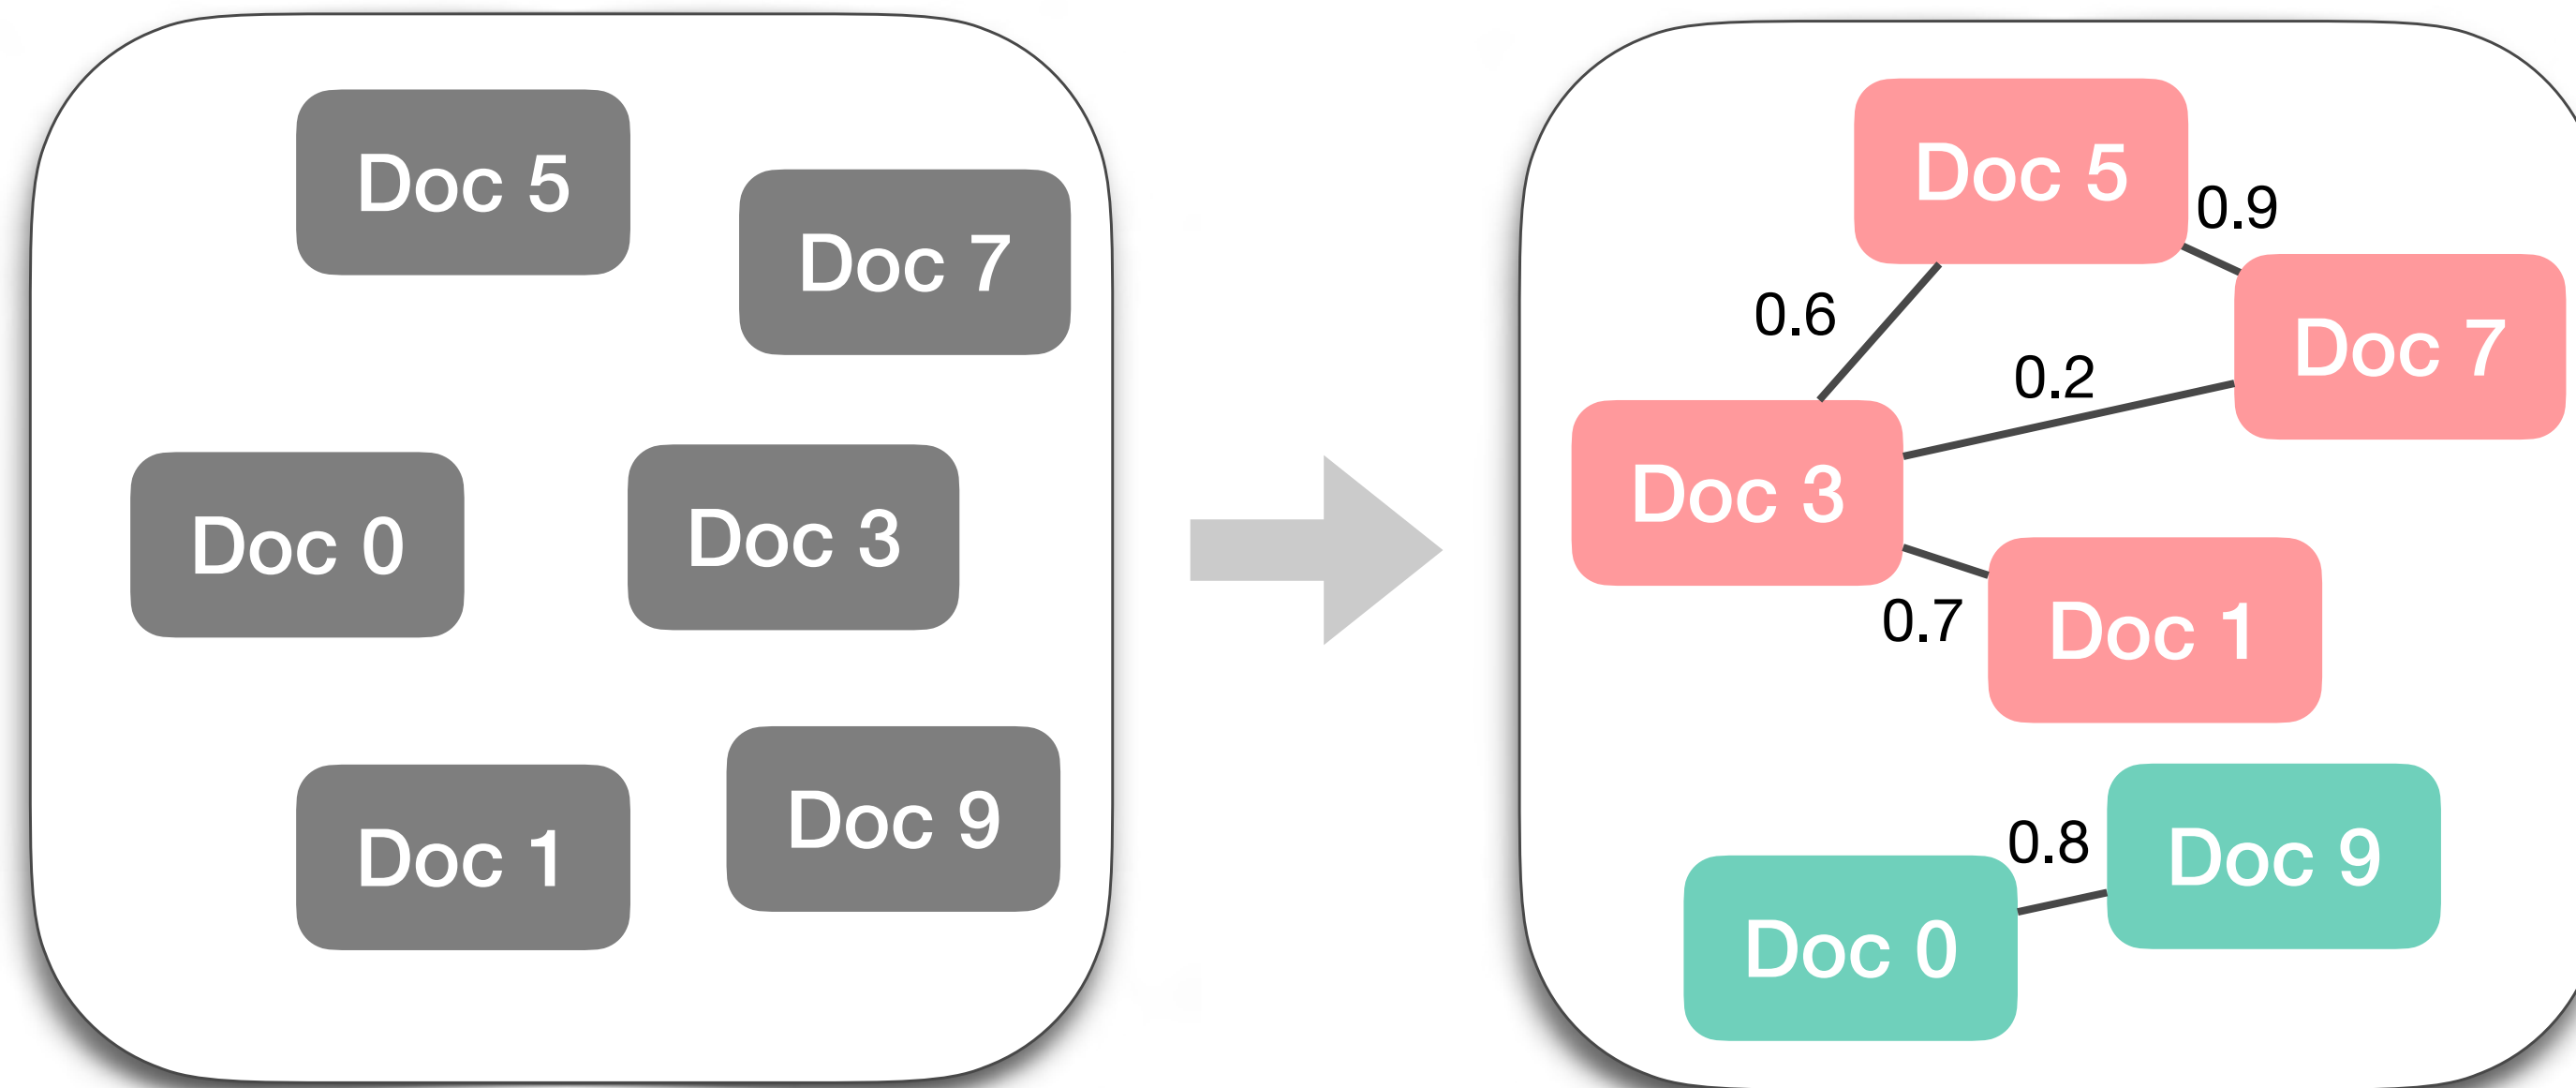
1) **Related** documents in the same context

2) Each document appears **exactly once**

# In-Context Pretraining: Recipe

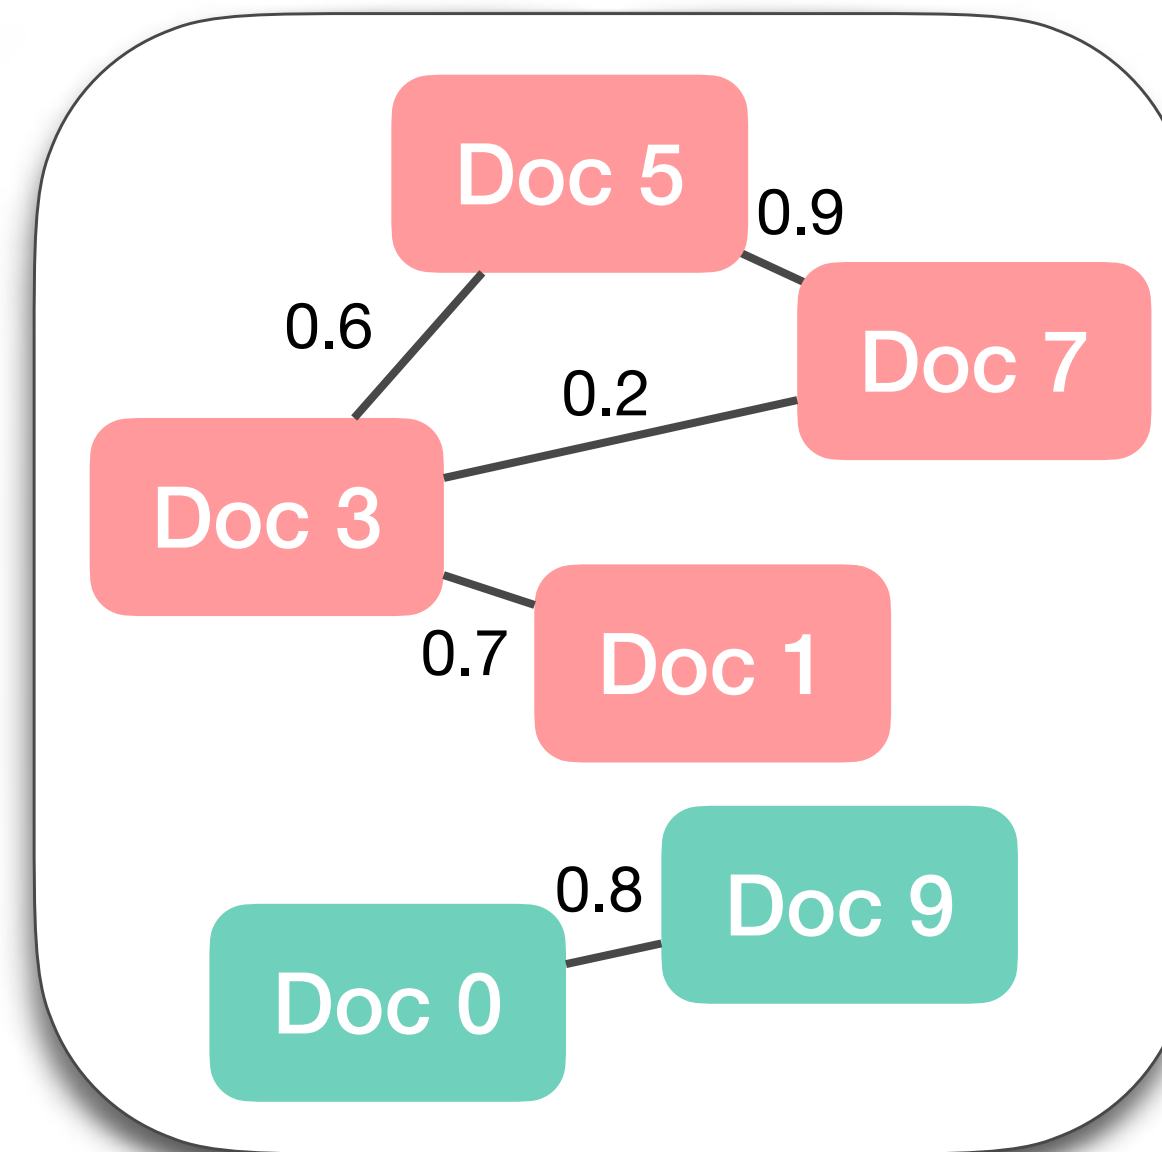
Document ordering problem

Pretraining Docs



# In-Context Pretraining: Recipe

Document ordering problem

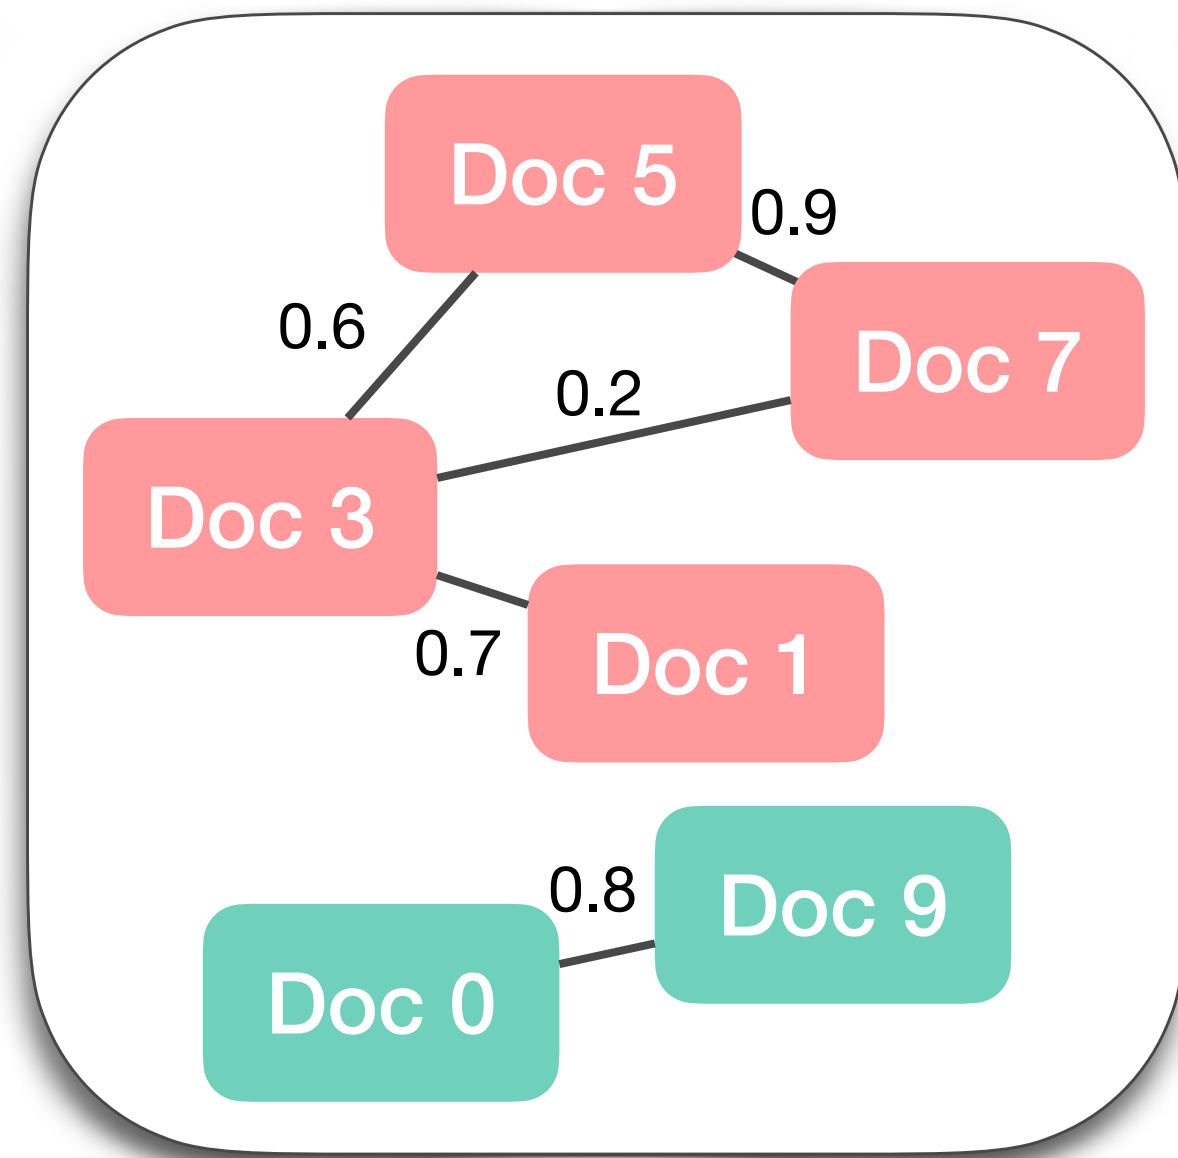




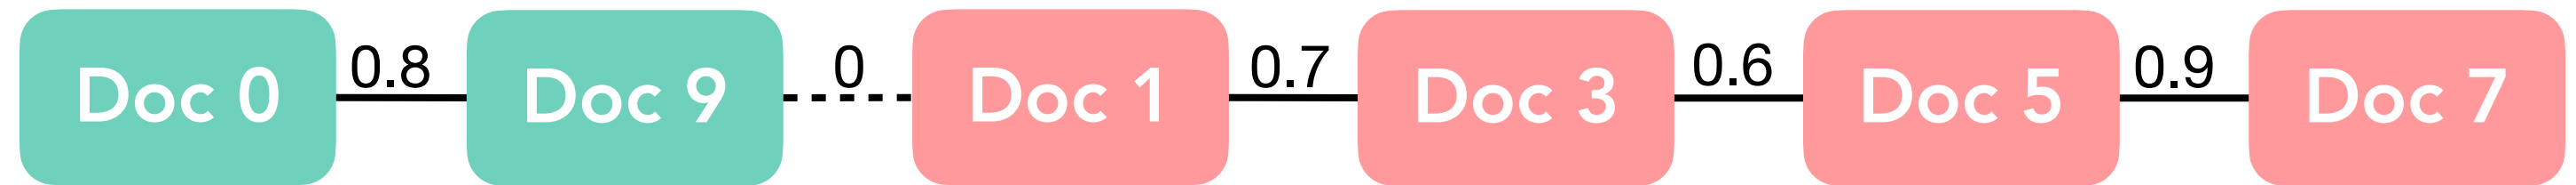
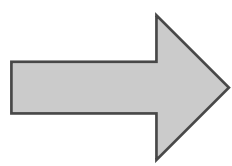
# In-Context Pretraining: Recipe

## Document ordering problem

*Find a path that visits each doc once, making related docs to be visited consecutively*

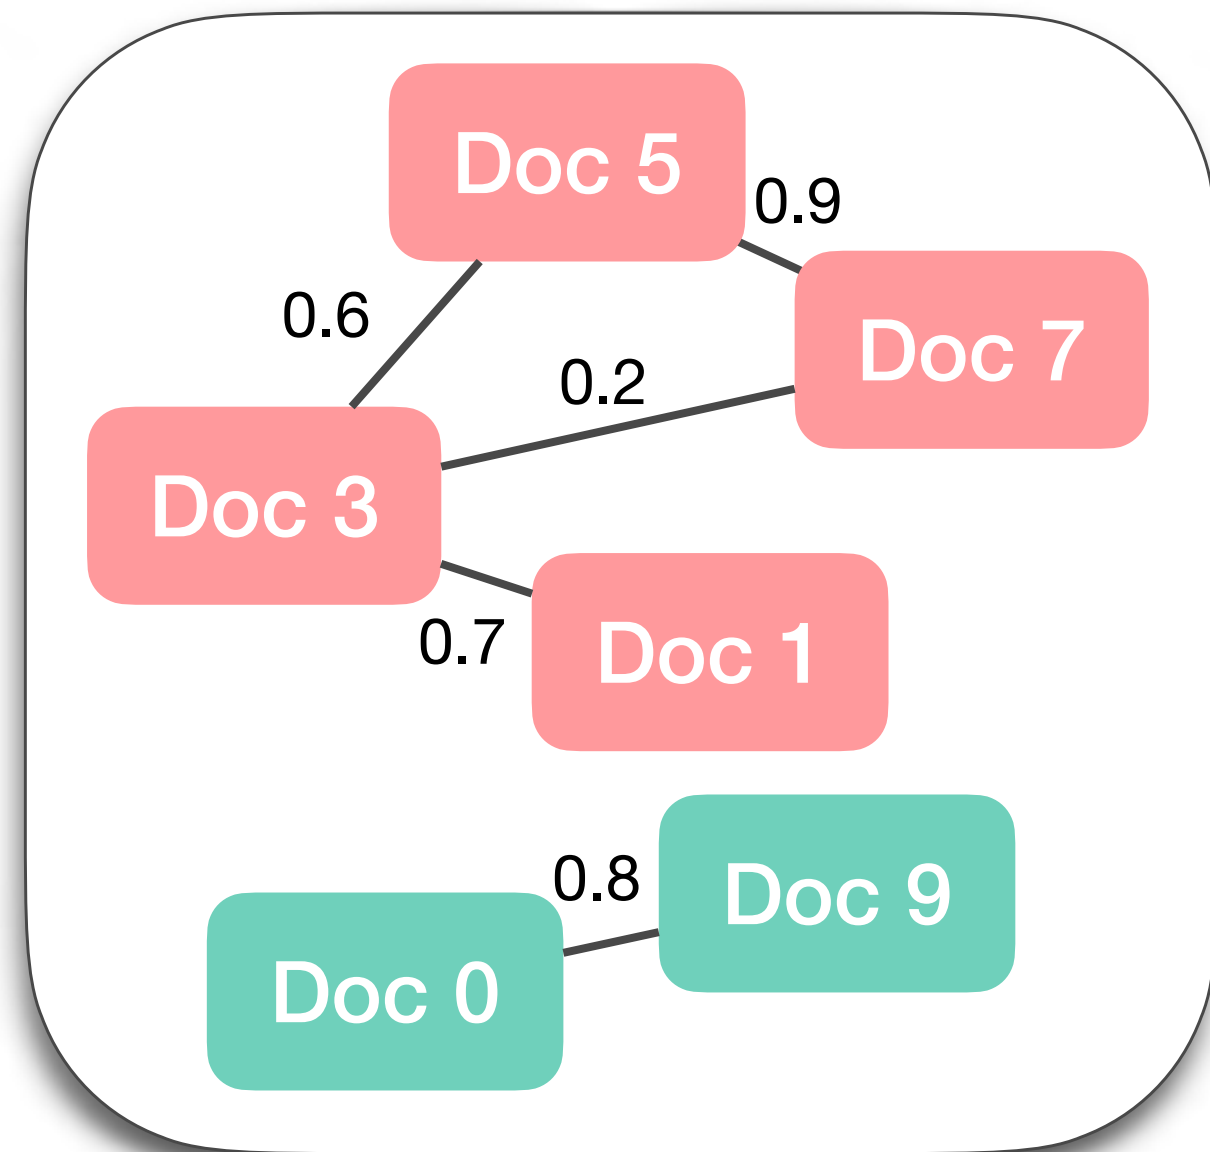


Maximum Traveling Salesman



# Document Ordering Problem

Input:



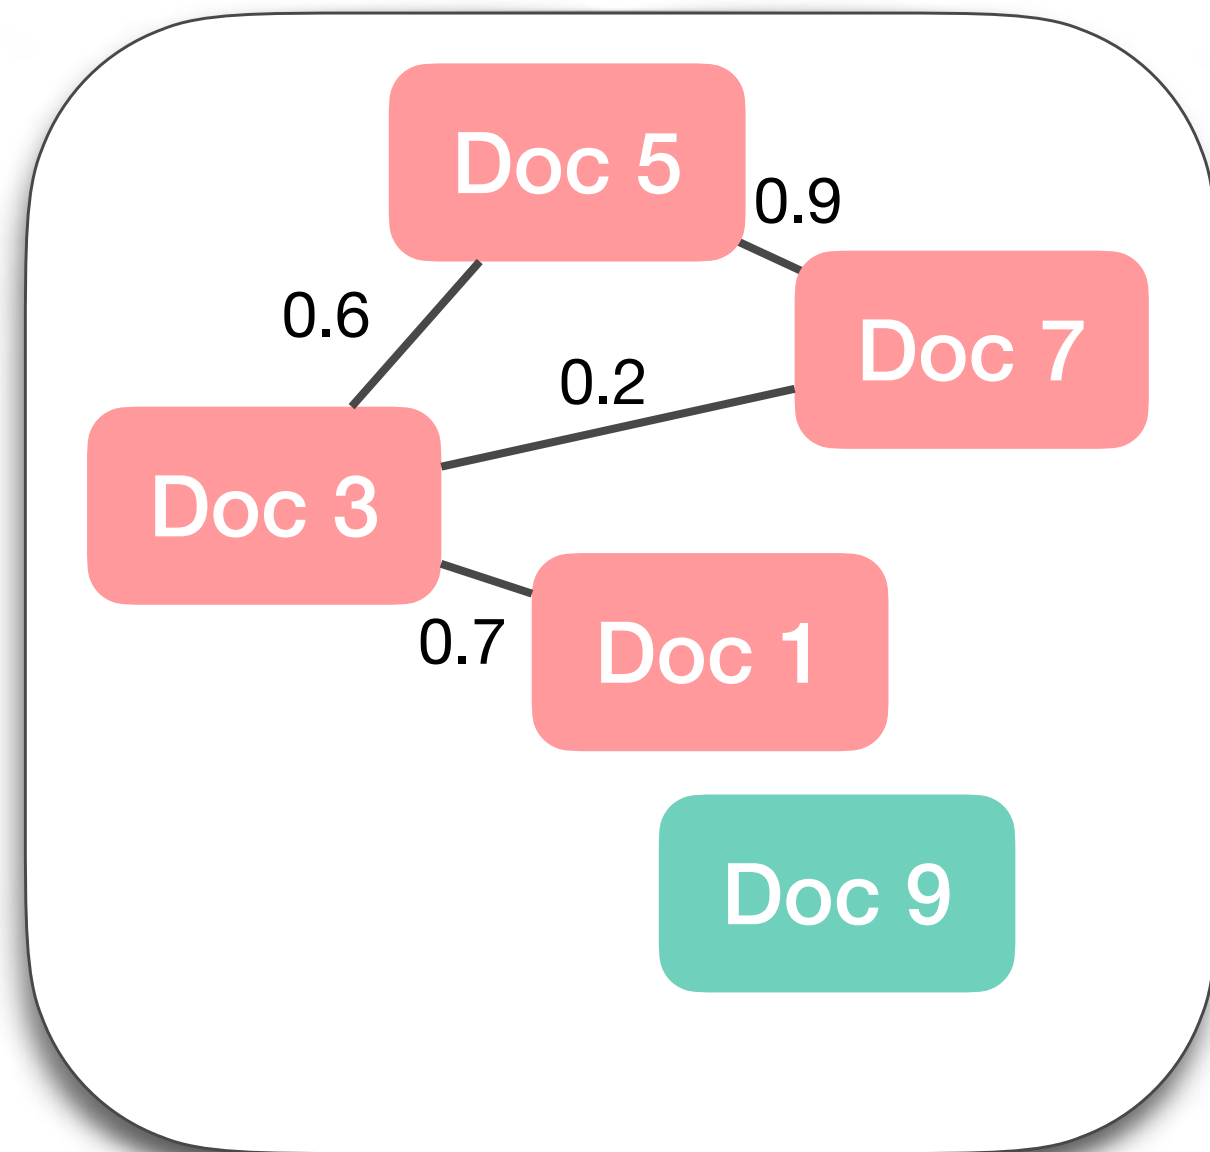
Output: path

**Procedure:**

select an unvisited doc with the min degree

# Document Ordering Problem

Input:



Output: path

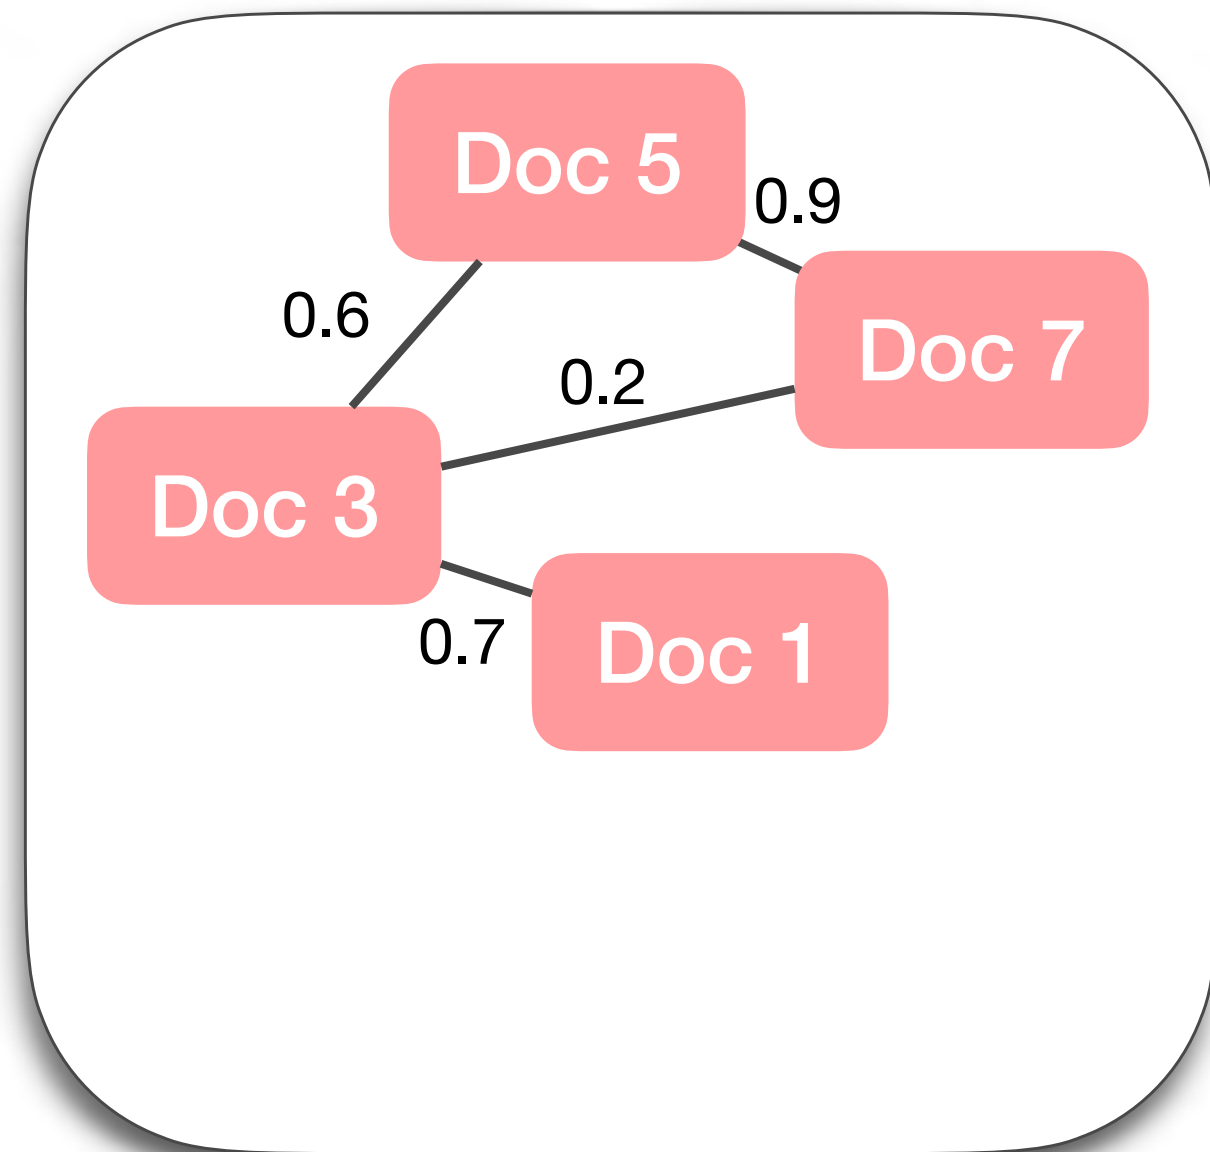


Procedure:

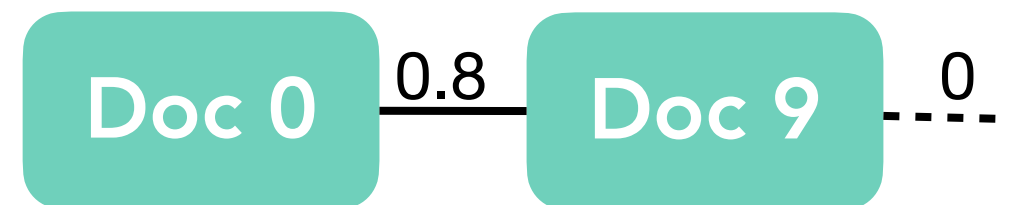
Move to the unvisited neighbor with max weight until all neighbors are visited

# Document Ordering Problem

Input:



Output: path

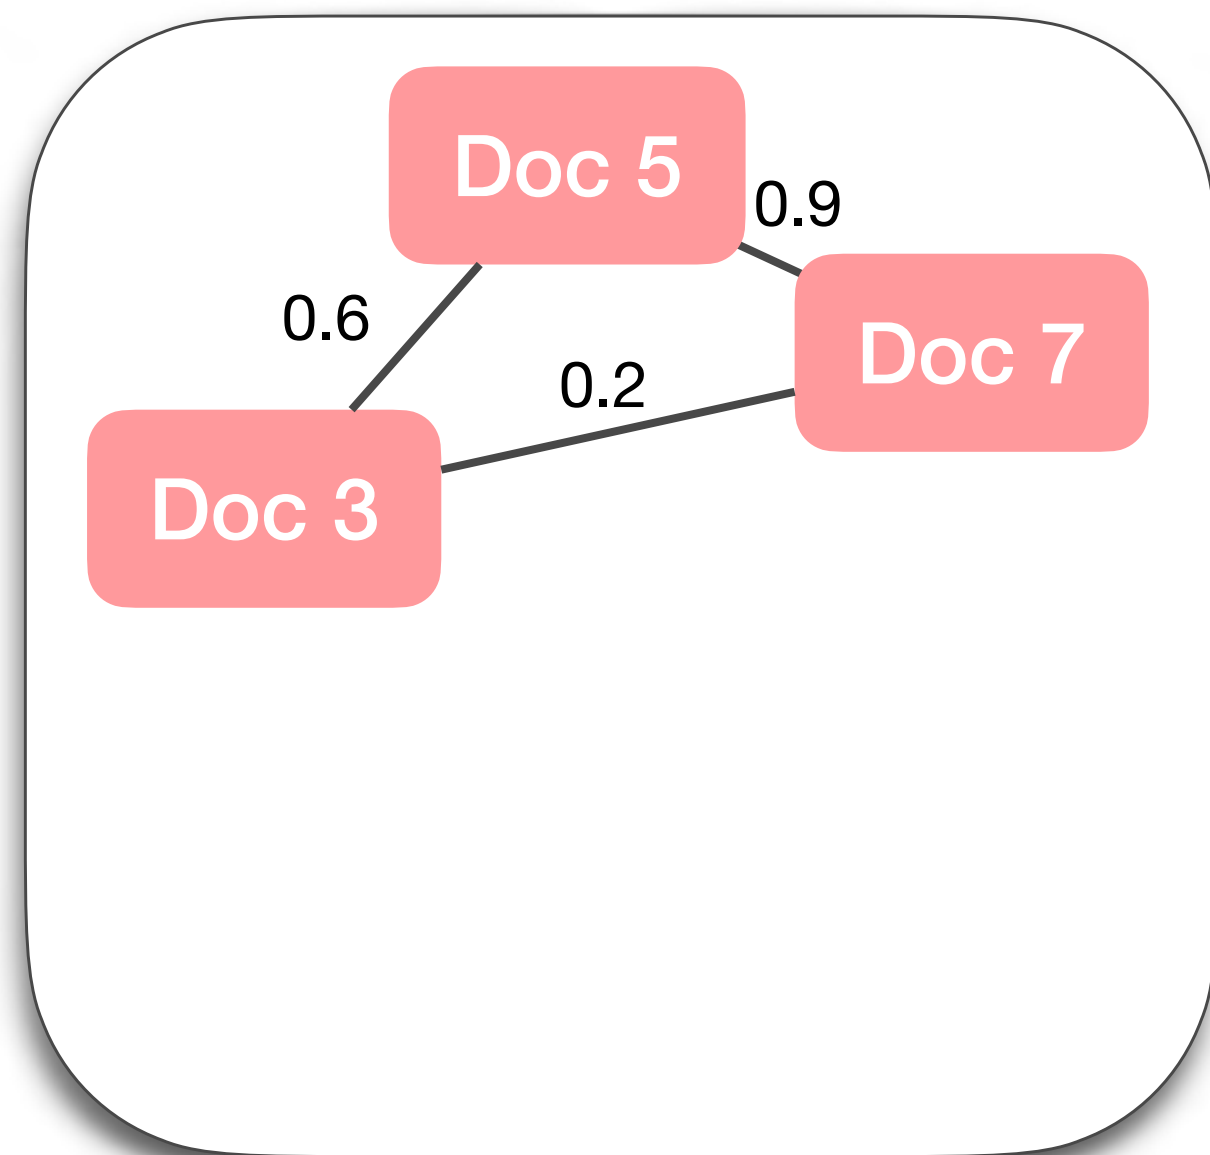


Procedure:

select an unvisited doc with the min degree

# Document Ordering Problem

Input:



Output: path

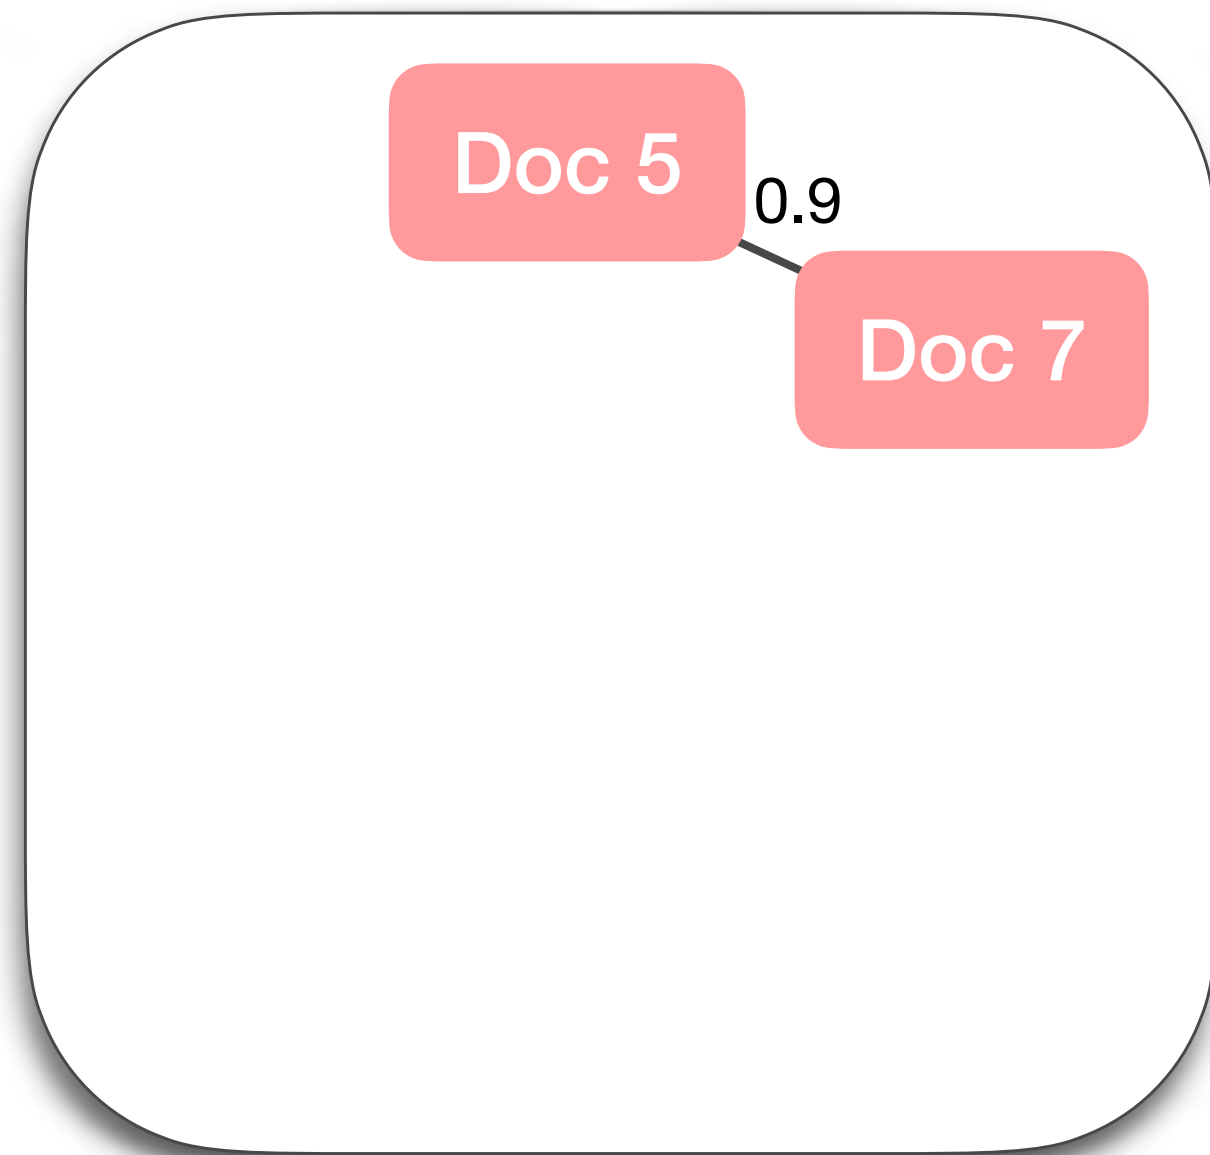


Procedure:

Move to the unvisited neighbor with max weight until all neighbors are visited

# Document Ordering Problem

Input:



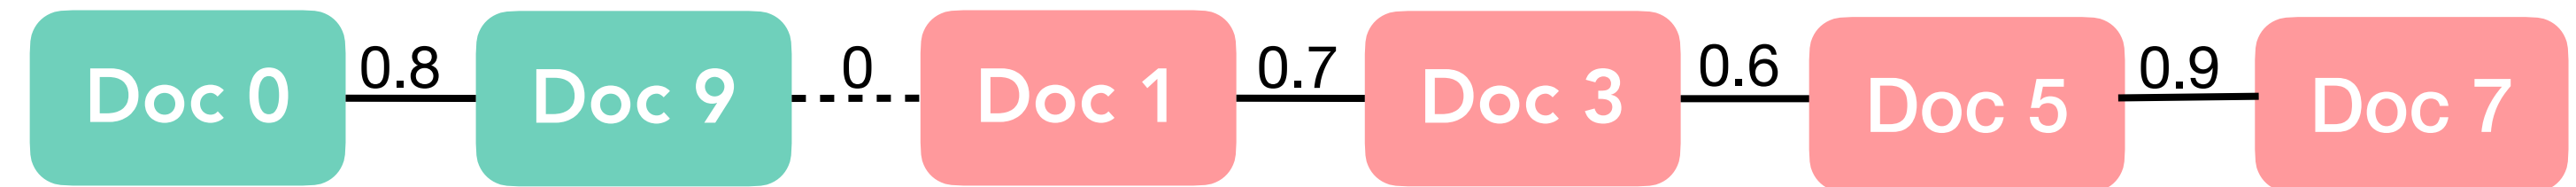
Output: path



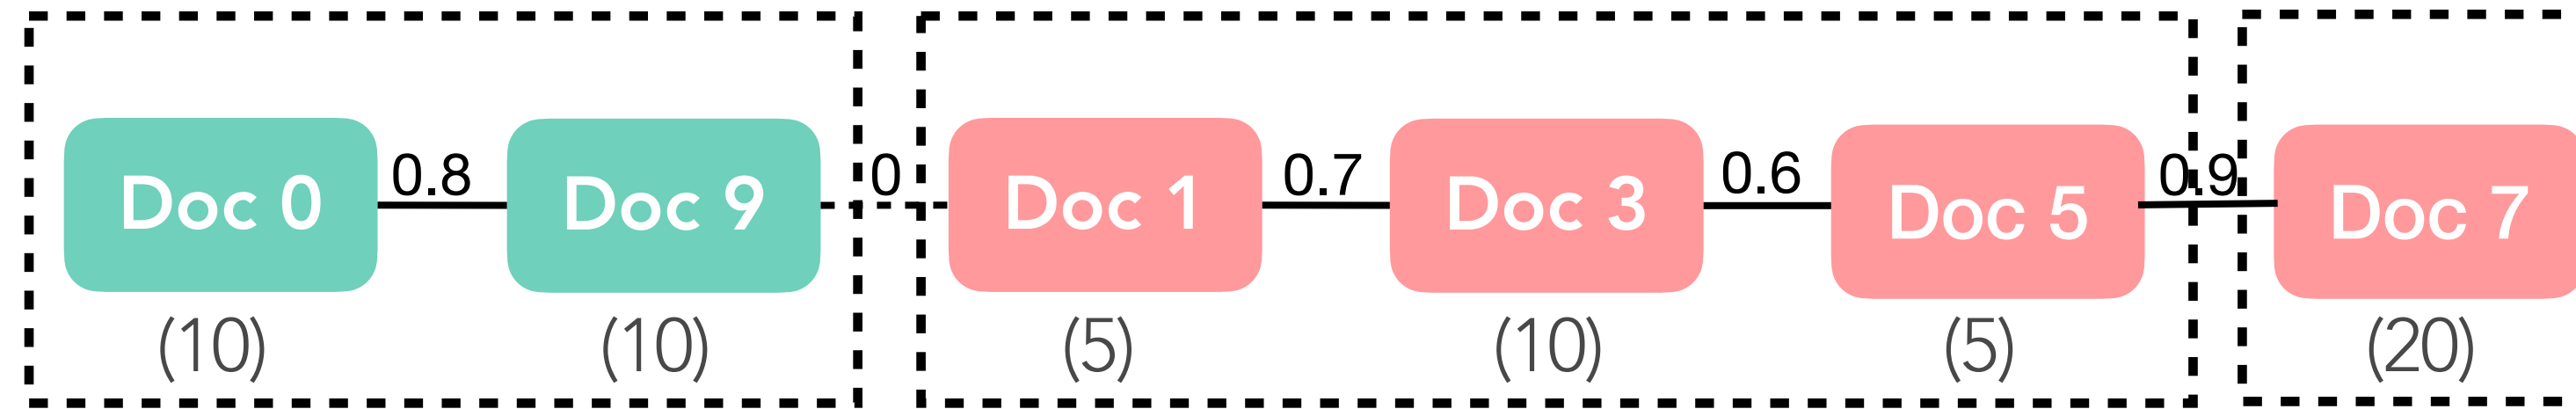
Procedure:

Move to the unvisited neighbor with max weight until all neighbors are visited

# Document Ordering Problem



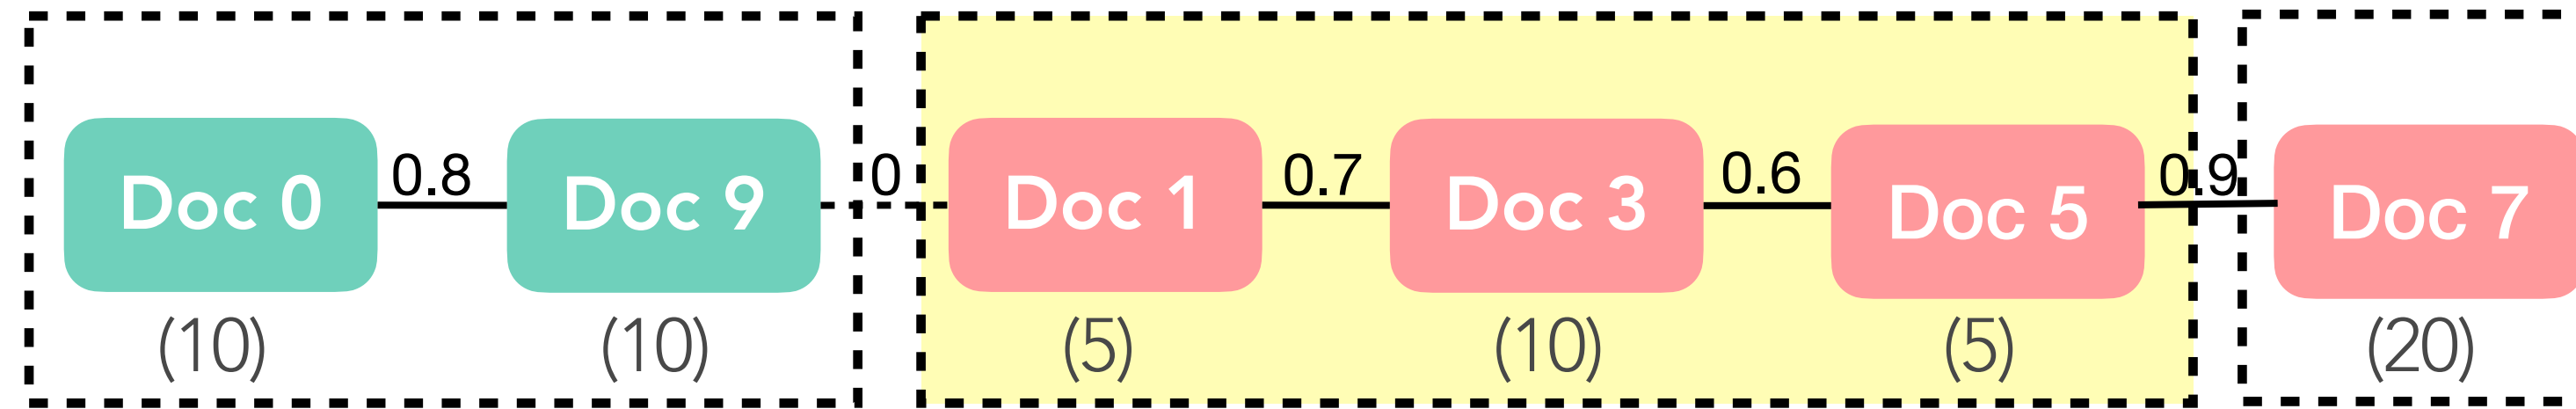
# Document Ordering Problem



If  $\text{max\_seq\_len} = 20$



# Document Ordering Problem



If  $\text{max\_seq\_len} = 20$

... Kadavra!" green light ... a jet of red light ... wand and a flash of green

Language Model

"Avada Kadavra!" ... green ... ... as a jet of red light ... ... his wand and a flash of

Doc 1

Doc 3



Doc 5

1) **Related** documents in the same context

2) Each document appears **exactly once**

Simple!  
(Training code remains same)

# Training Details

- **Architecture:** LLaMA 
- **Model:** 0.3, 0.7, 1.5, and **7B** model with sequence length of 8192 from scratch (**128 A100s for 9 days**)
- **Data:** 306B tokens from  (235M docs)
- **Retriever:** Contriever

# Baselines

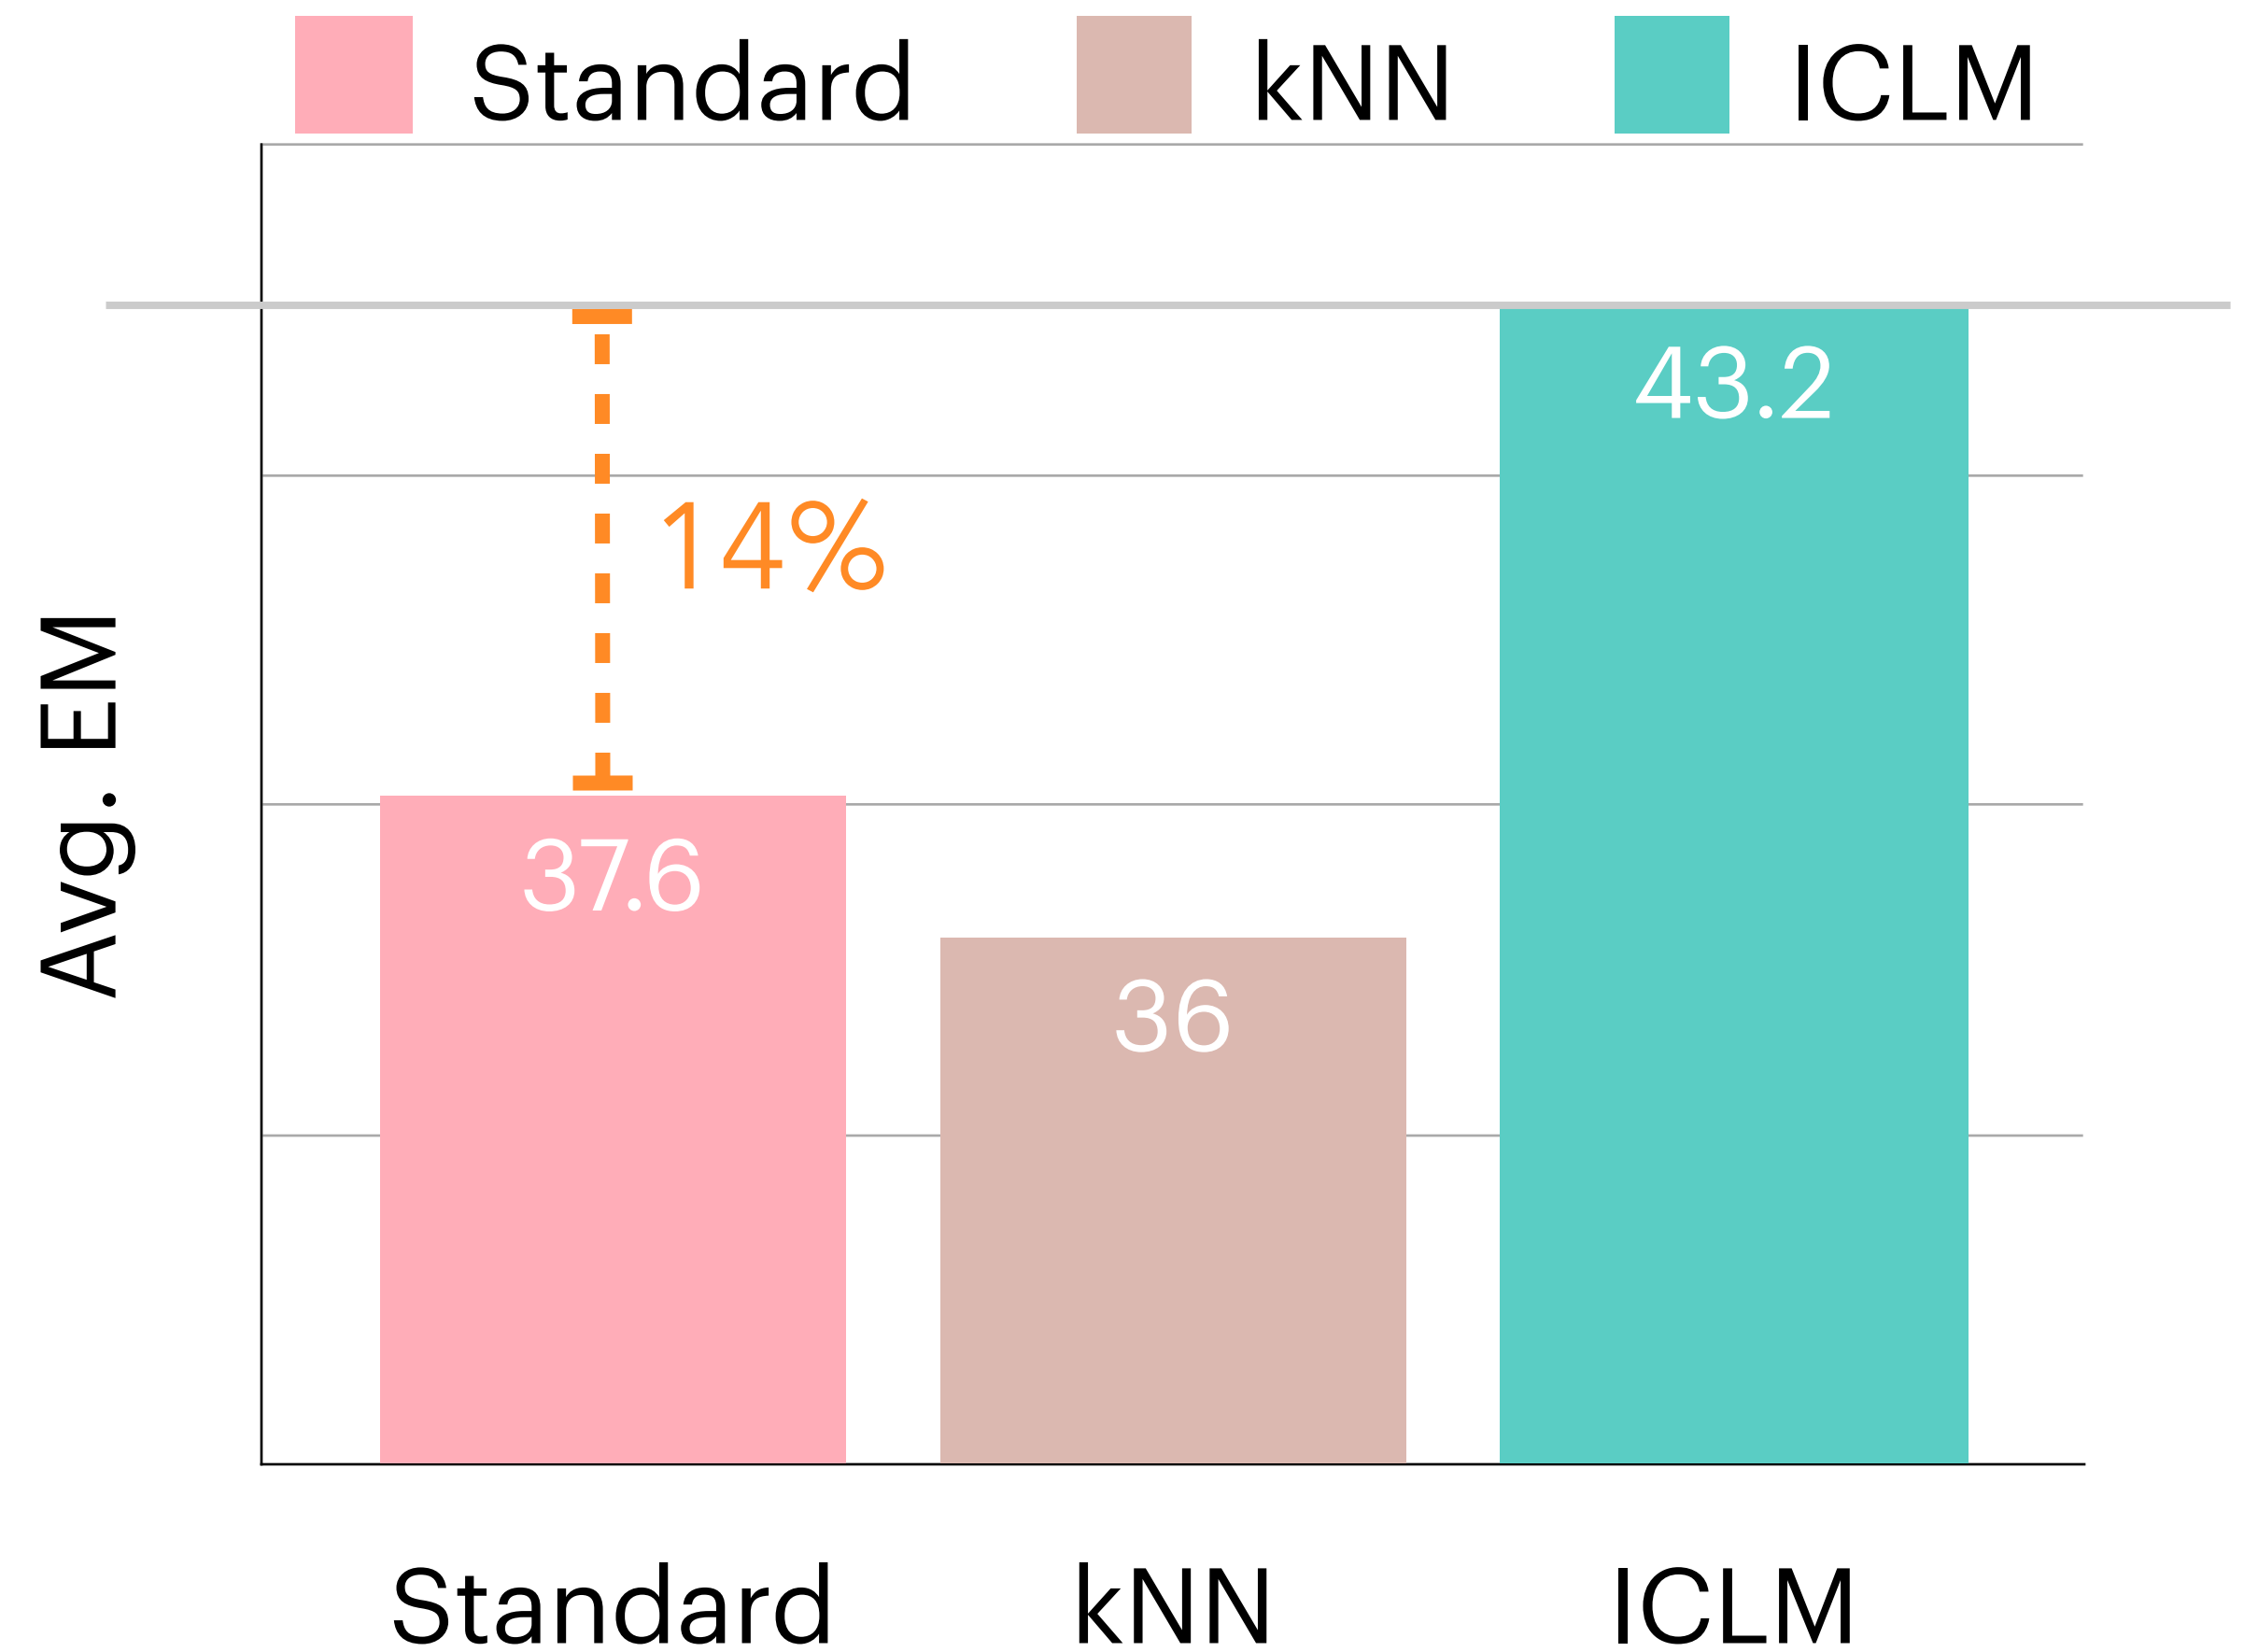
- **Standard:** places random docs in the input contexts
- **kNN:** places each doc and its retrieved top-k docs in the input

*Given the same number of training steps, kNN exposes LMs to a less diverse set of documents, since documents can **repeat***

# Results: Reading Comprehension

## Tasks:

- 1. Single document:** race-high, race-middle, boolq, squad
- 2. Multi document:** hotpotQA, drop



# Results: Open-Domain QA

## Tasks: NQ, TQA

### With retrieved docs

Write a high-quality answer for the given question using only the provided search results (some of which might be irrelevant).

Document [1] (Title: Asian Americans in science and technology) Prize in physics for discovery of the subatomic particle  $J/\psi$ . Subrahmanyan Chandrasekhar shared...

**Document [2] (Title: List of Nobel laureates in Physics) The first Nobel Prize in Physics was awarded in 1901 to Wilhelm Conrad Röntgen, of Germany, who received...**

Document [3] (Title: Scientist) and pursued through a unique method, was essentially in place. Ramón y Cajal won the Nobel Prize in 1906 for his remarkable...

Question: who got the first nobel prize in physics

Answer:

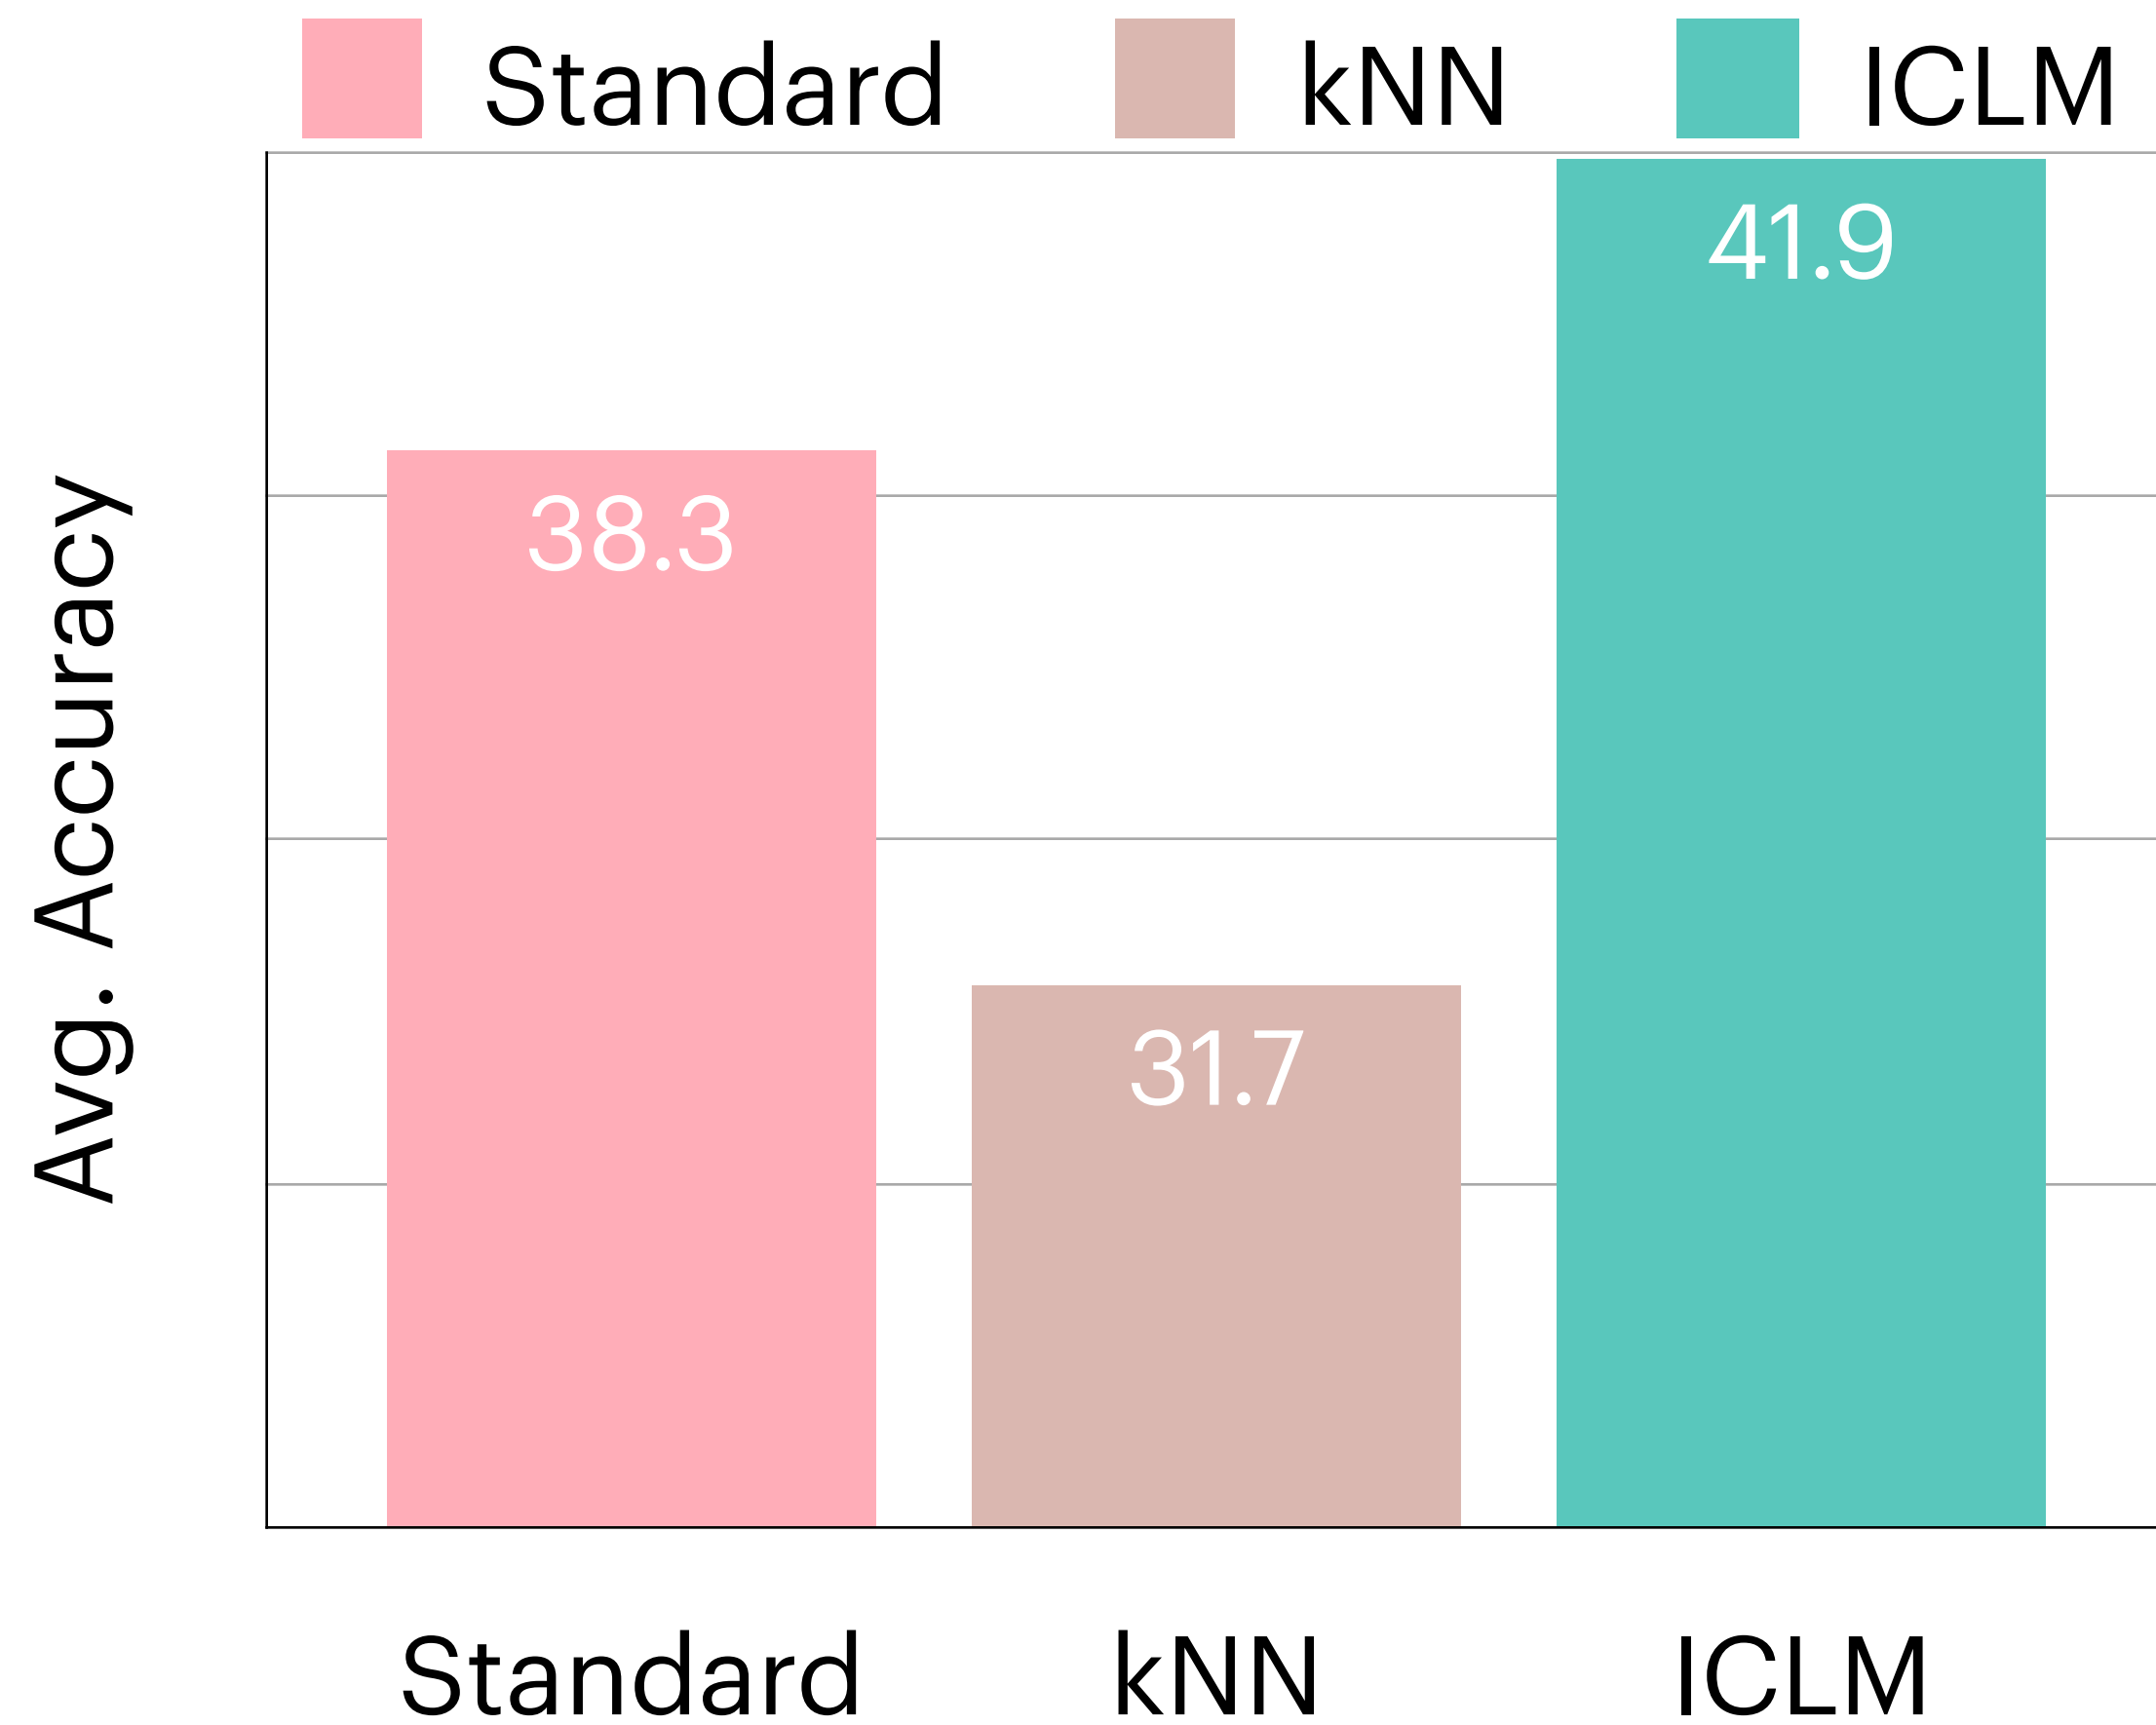
### W/o retrieved docs

Question: who got the first nobel prize in physics

Answer:

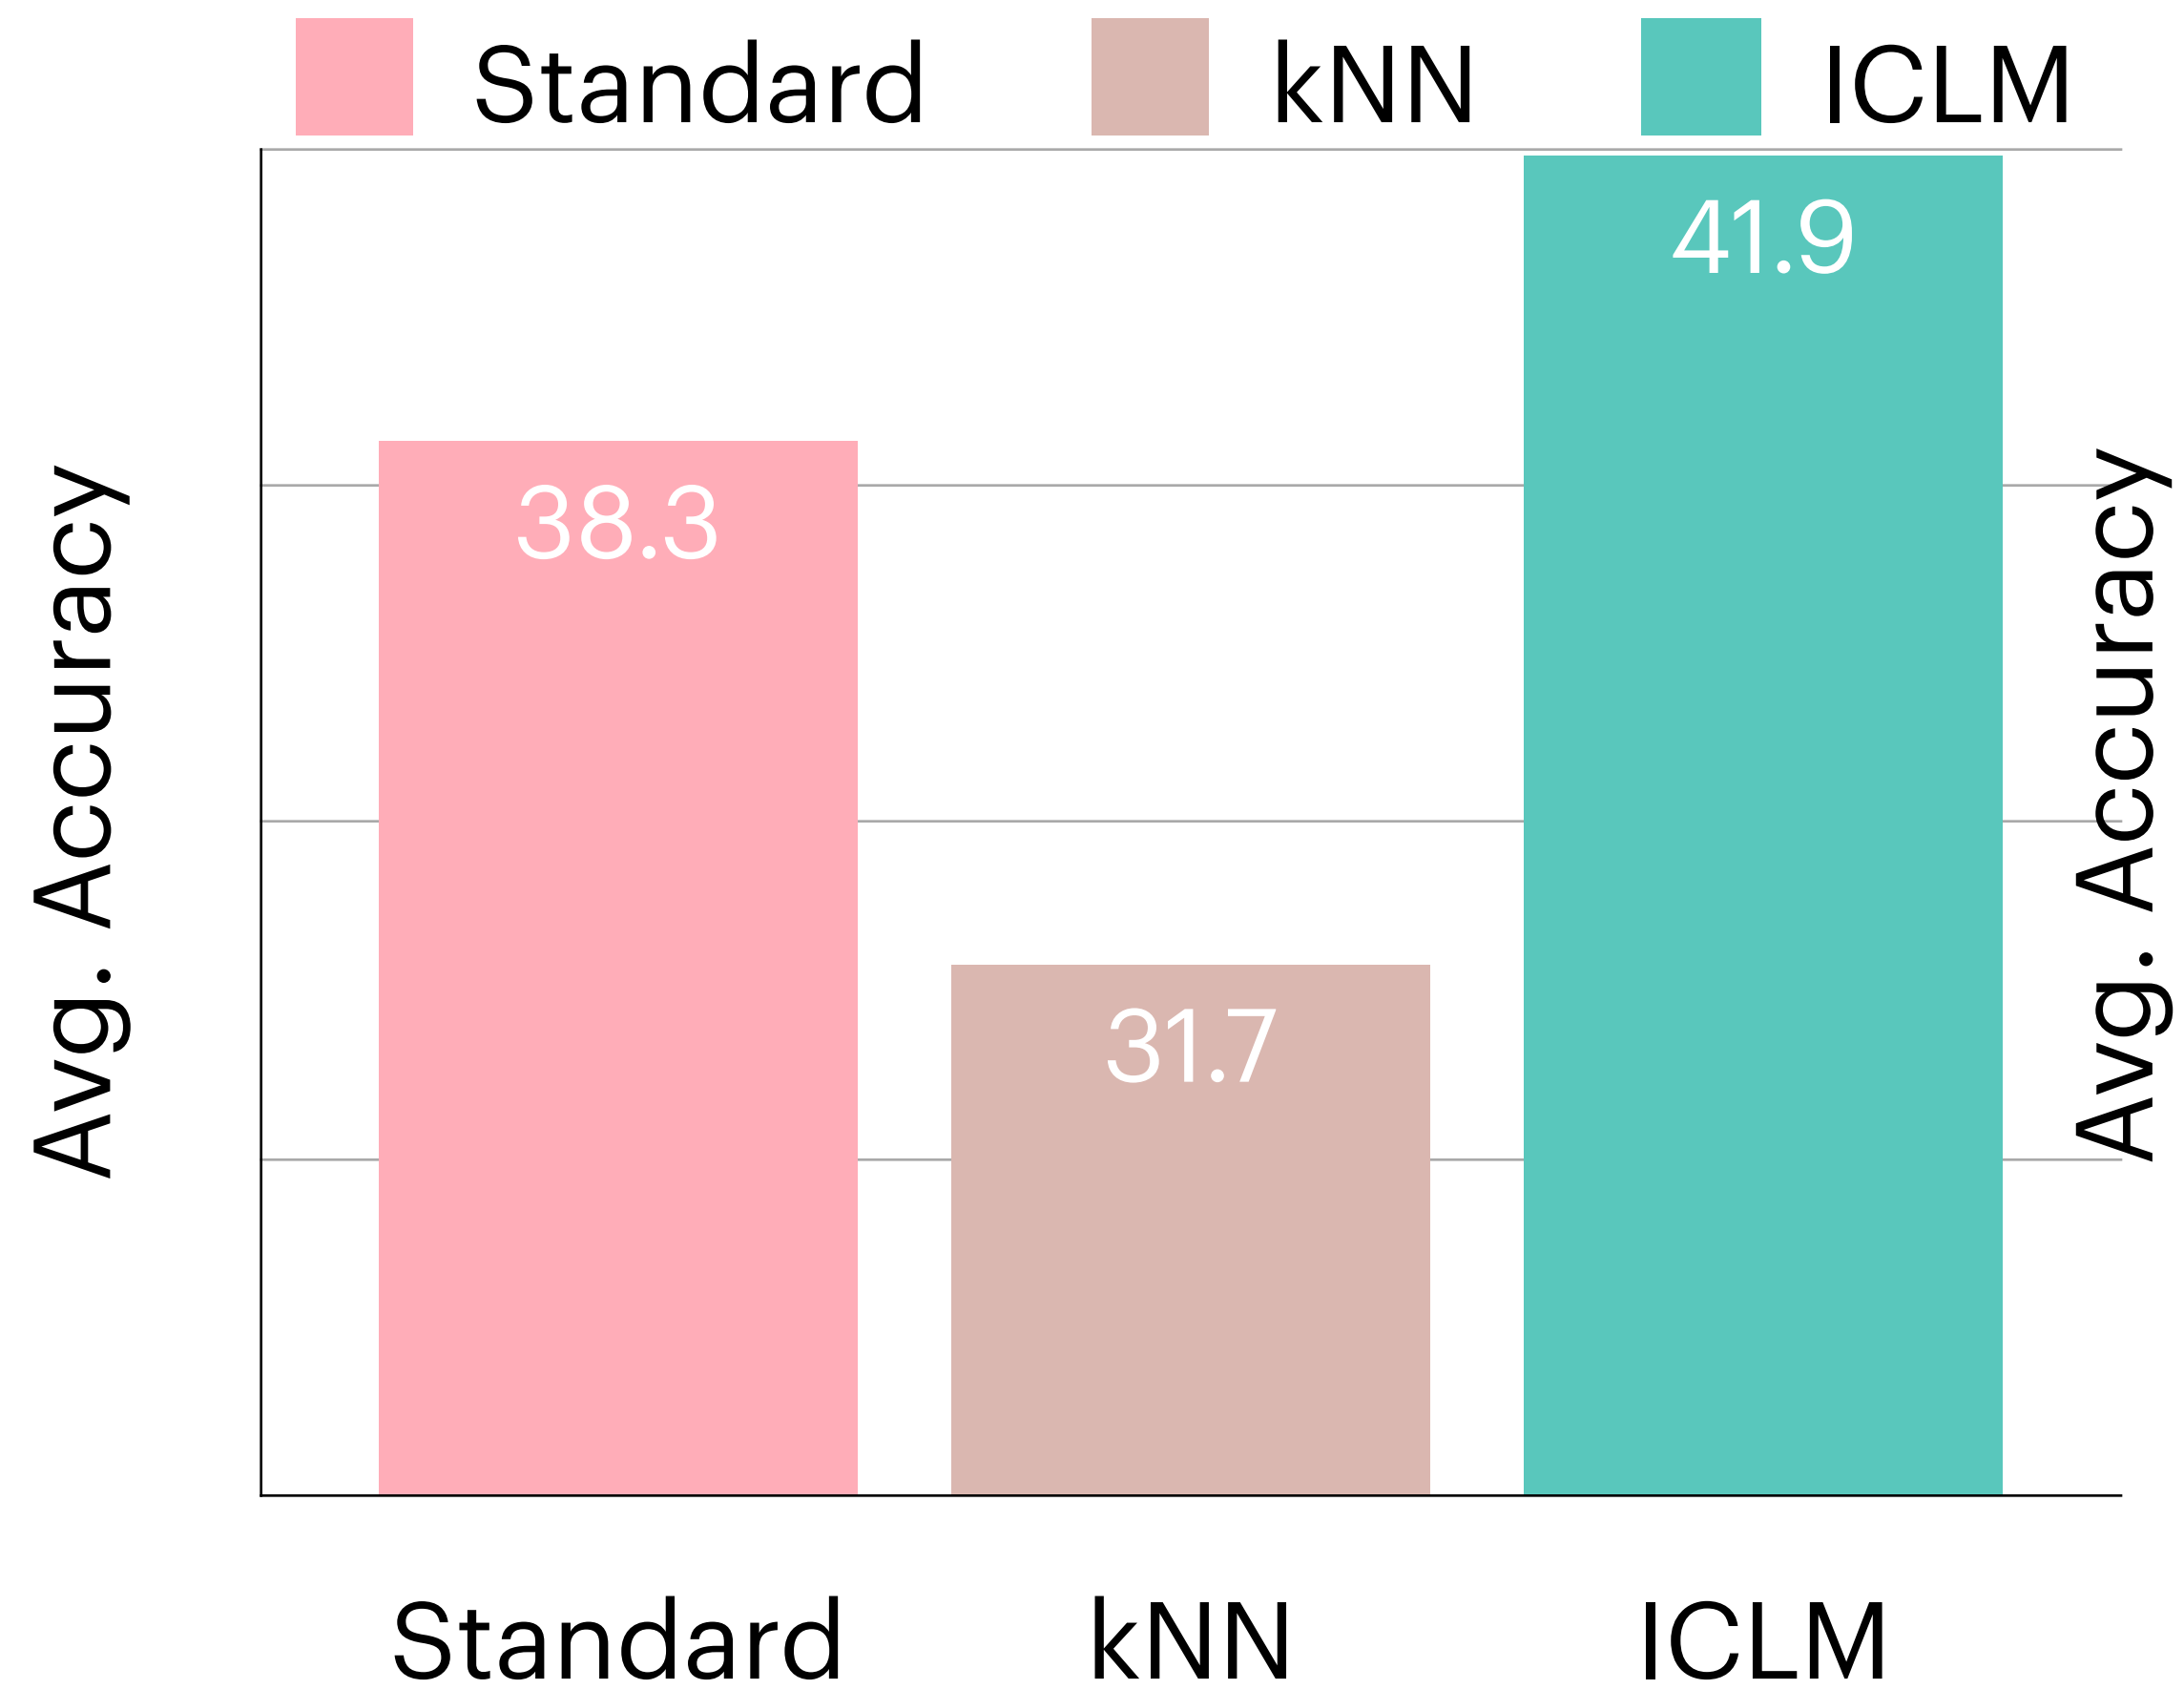
# Results: Open-Domain QA

**With** retrieved docs

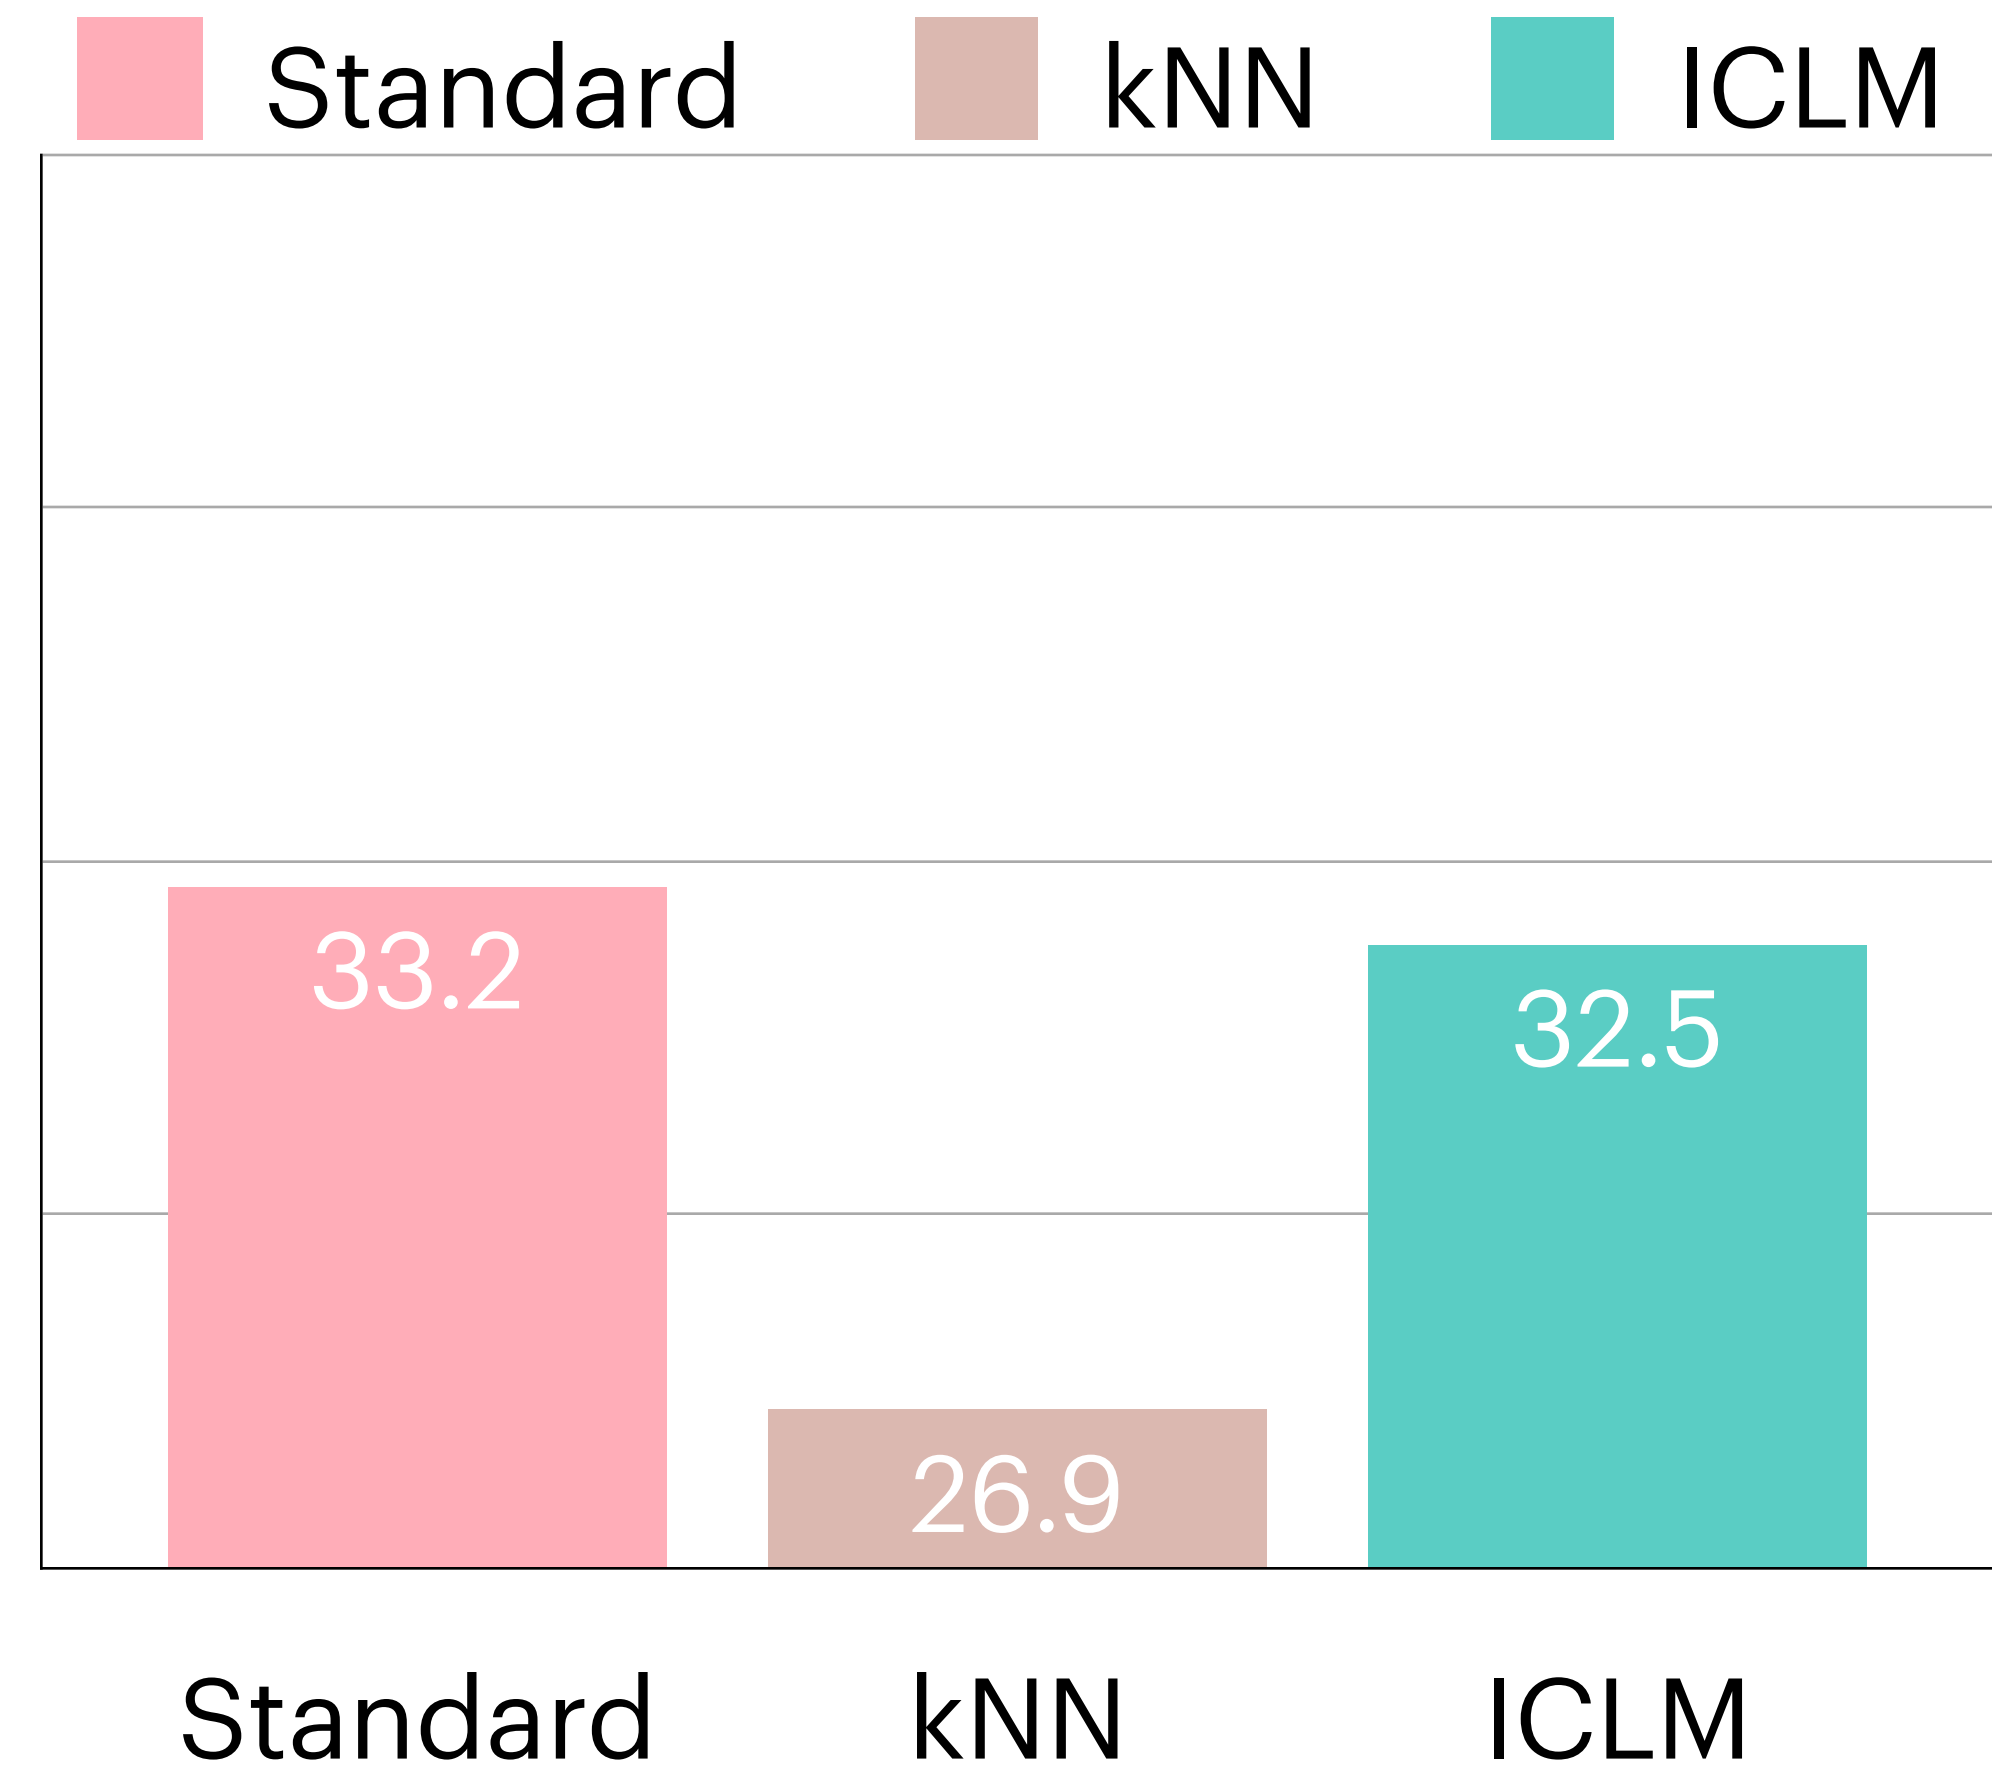


# Results: Open-Domain QA

**With** retrieved docs



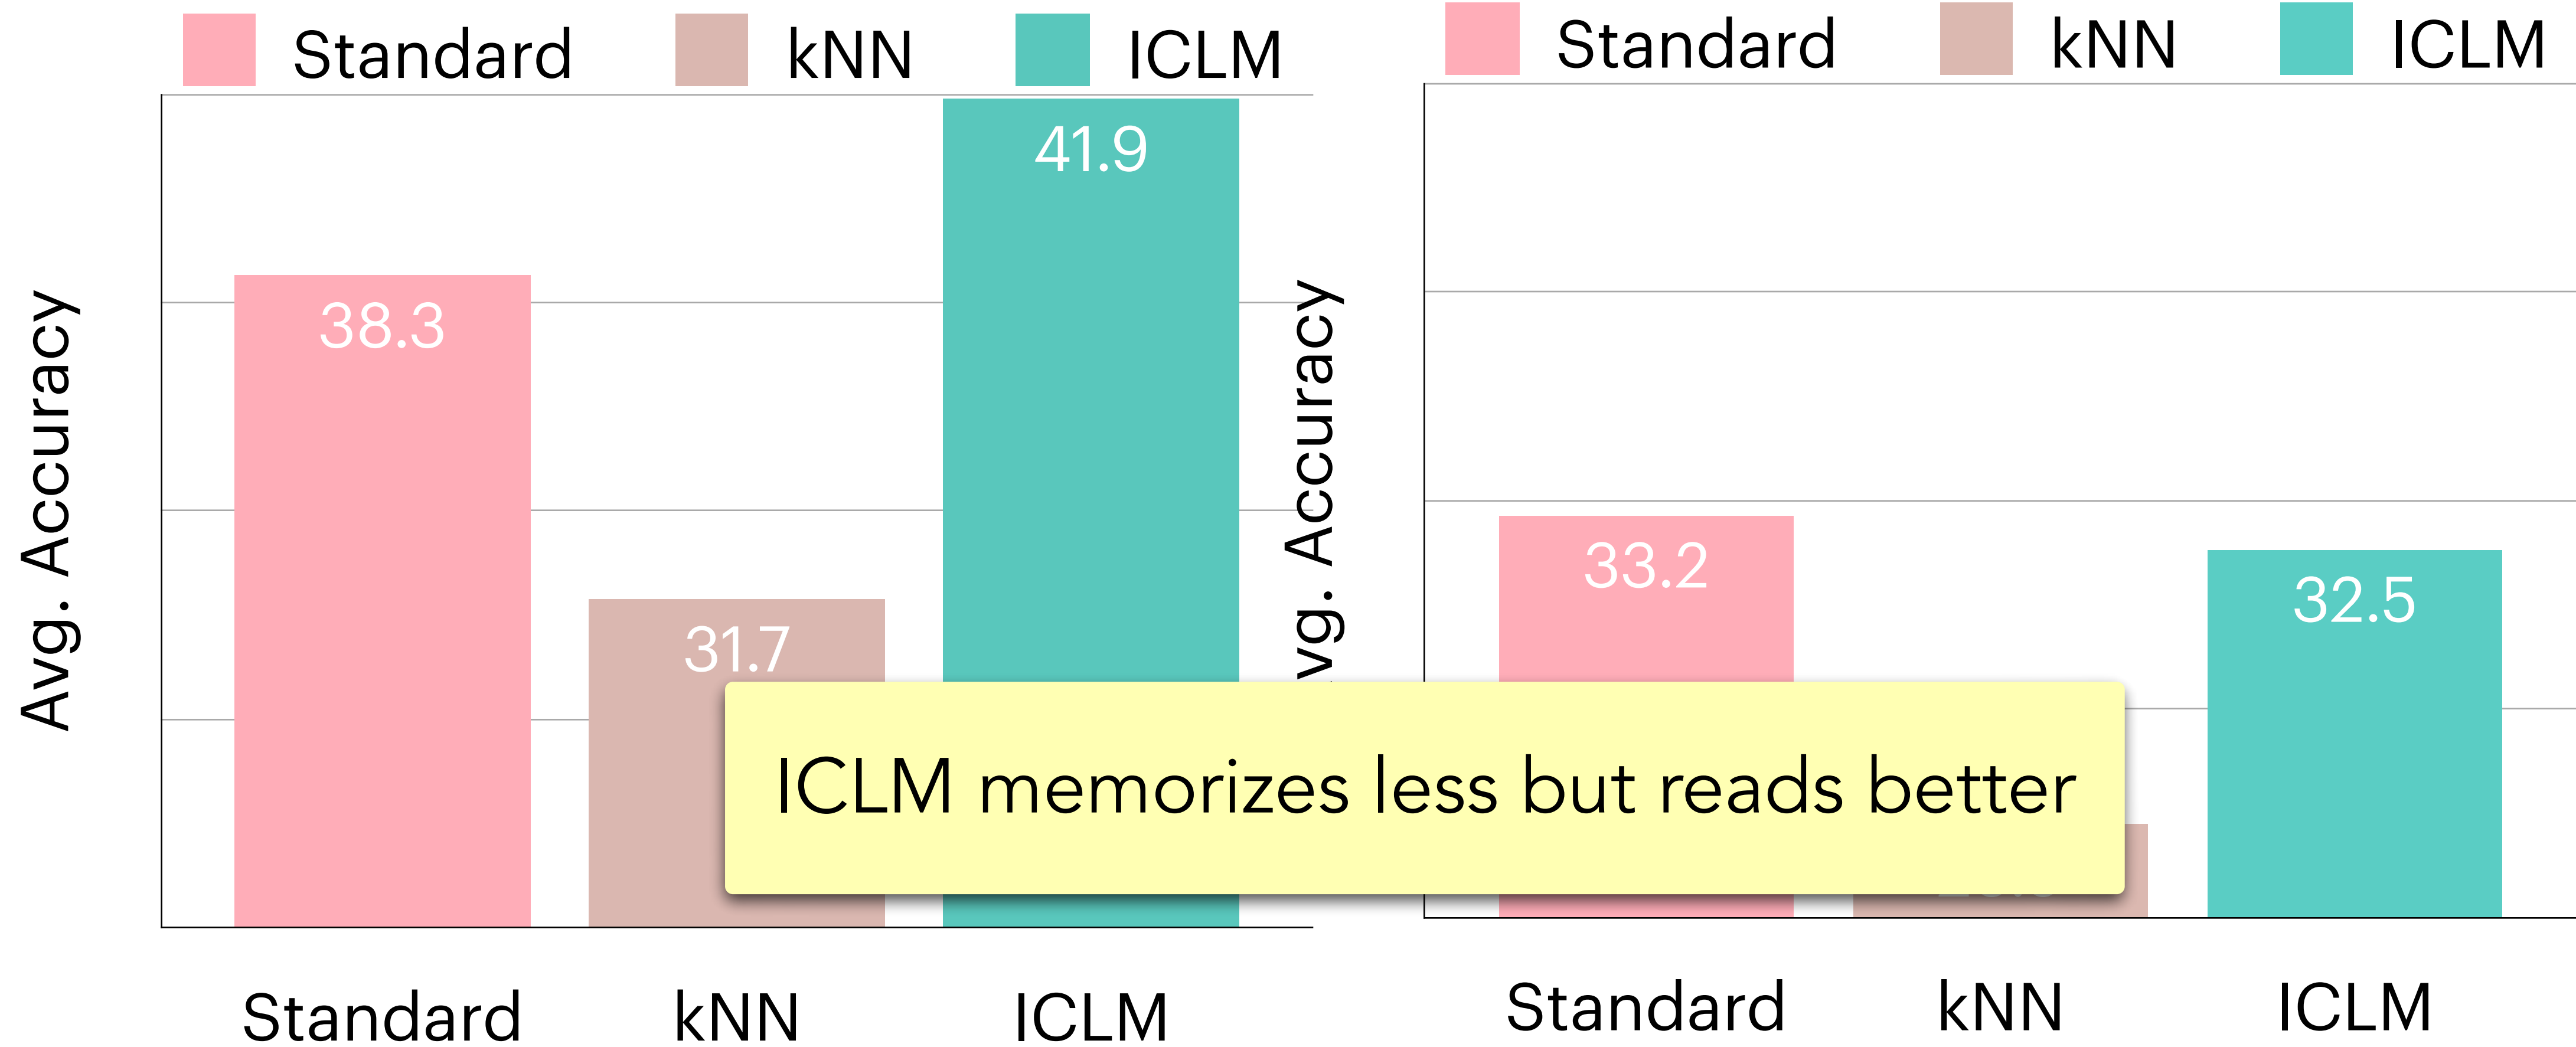
**W/o** retrieved docs



# Results: Open-Domain QA

**With** retrieved docs



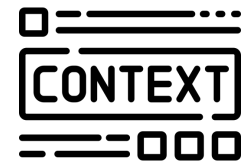

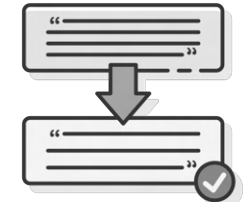

**W/o** retrieved docs





# Results

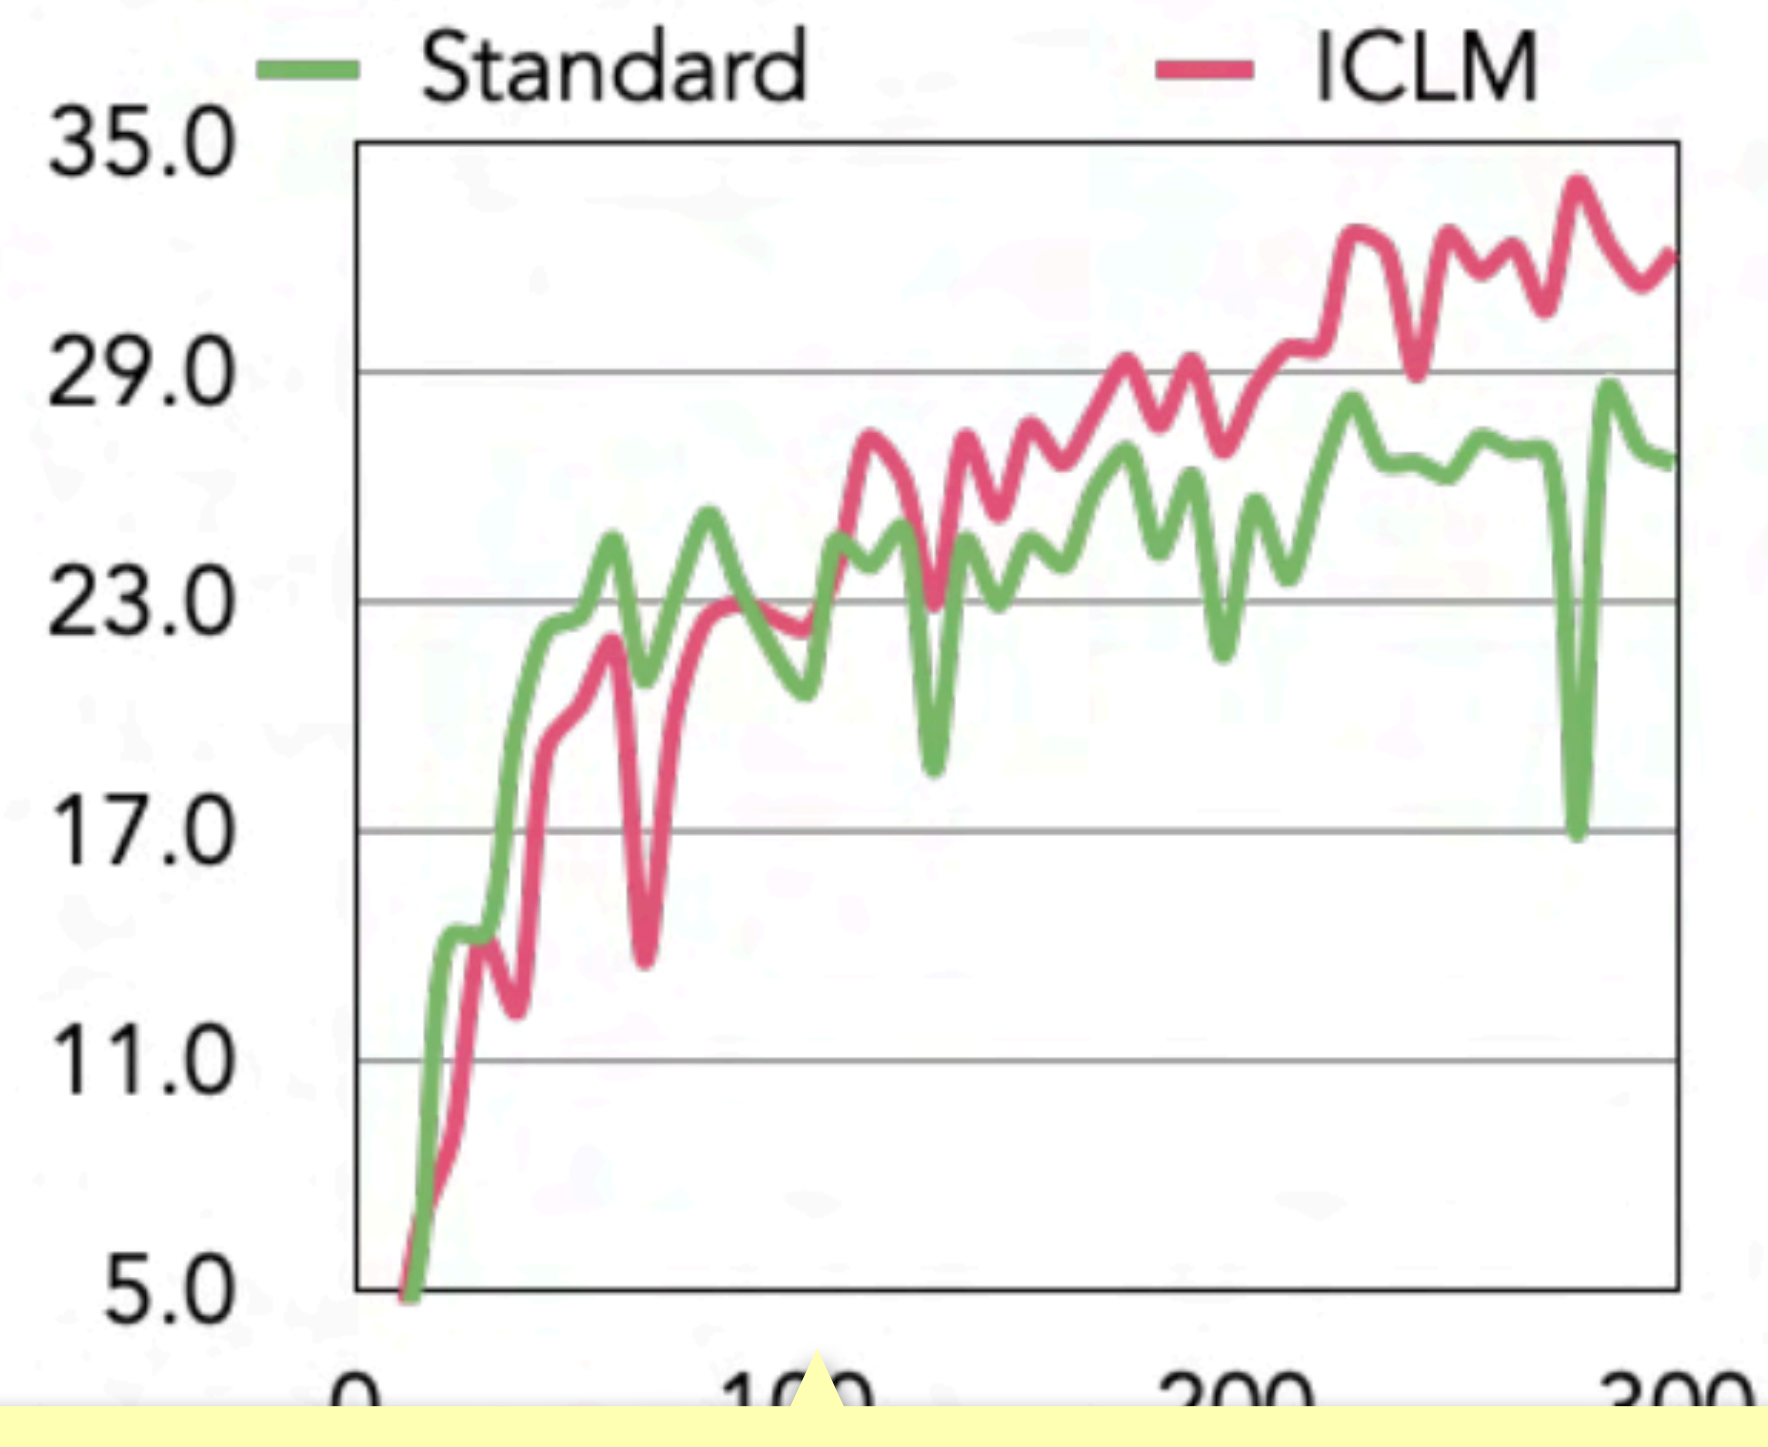
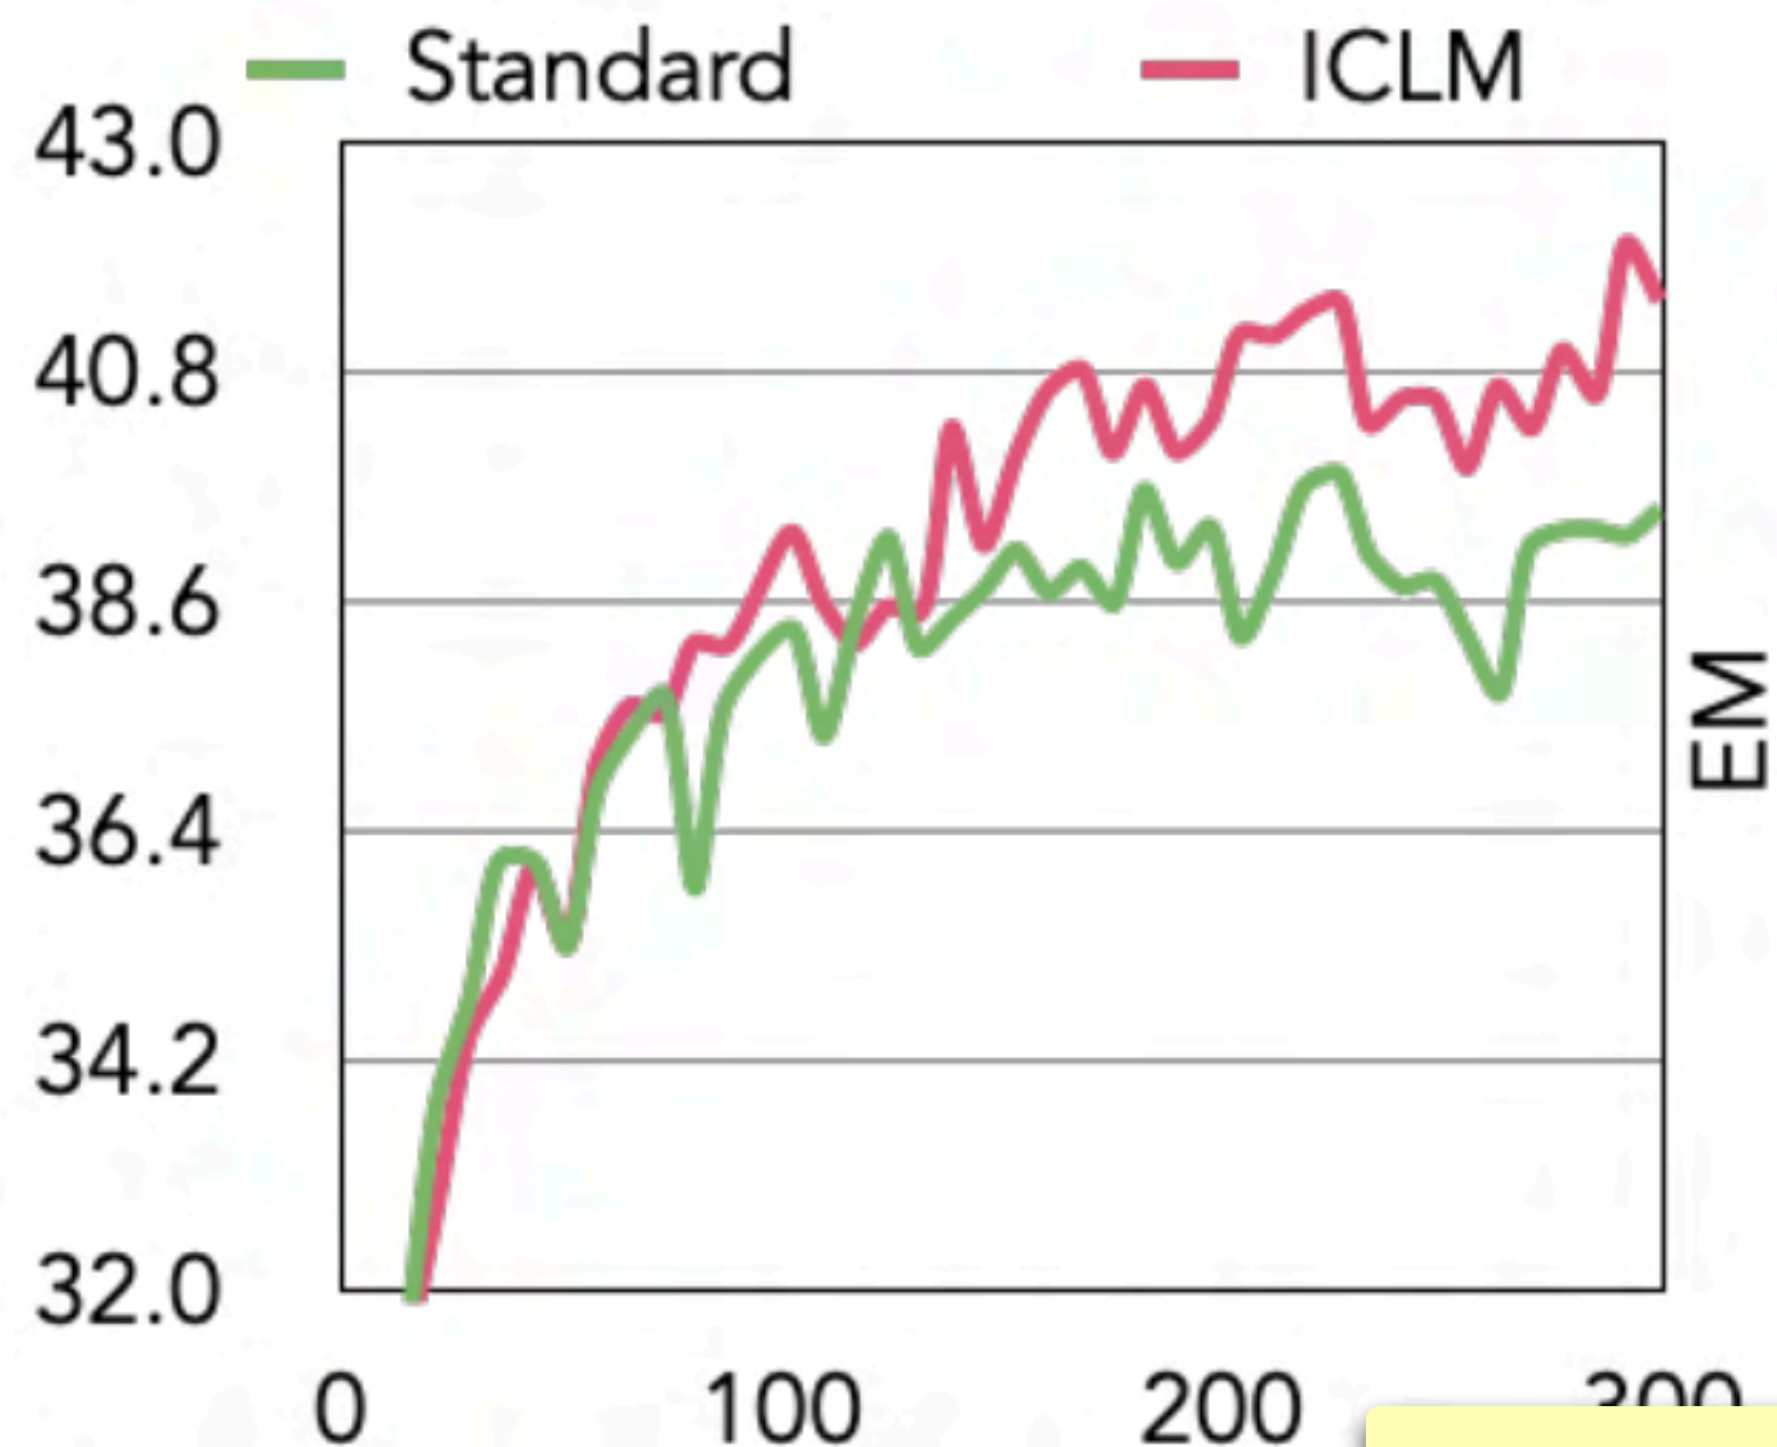
23 benchmarks in total

		Accuracy	
		Standard	In-Context Pretraining
	Open-Domain QA (w/ retrieval)	38%	<b>10.5%</b> → 42%
	In-Context Learning	66%	<b>7.5%</b> → 71%
	Reading Comprehension	37%	<b>14.0%</b> → 43%
	Factuality	44%	<b>15.9%</b> → 51%
	Long Document Reasoning	32%	<b>7.5%</b> → 34%

# Evolution of Performance

Reading comprehension

Open-domain QA (w/ retrieval)



(b) Race-High

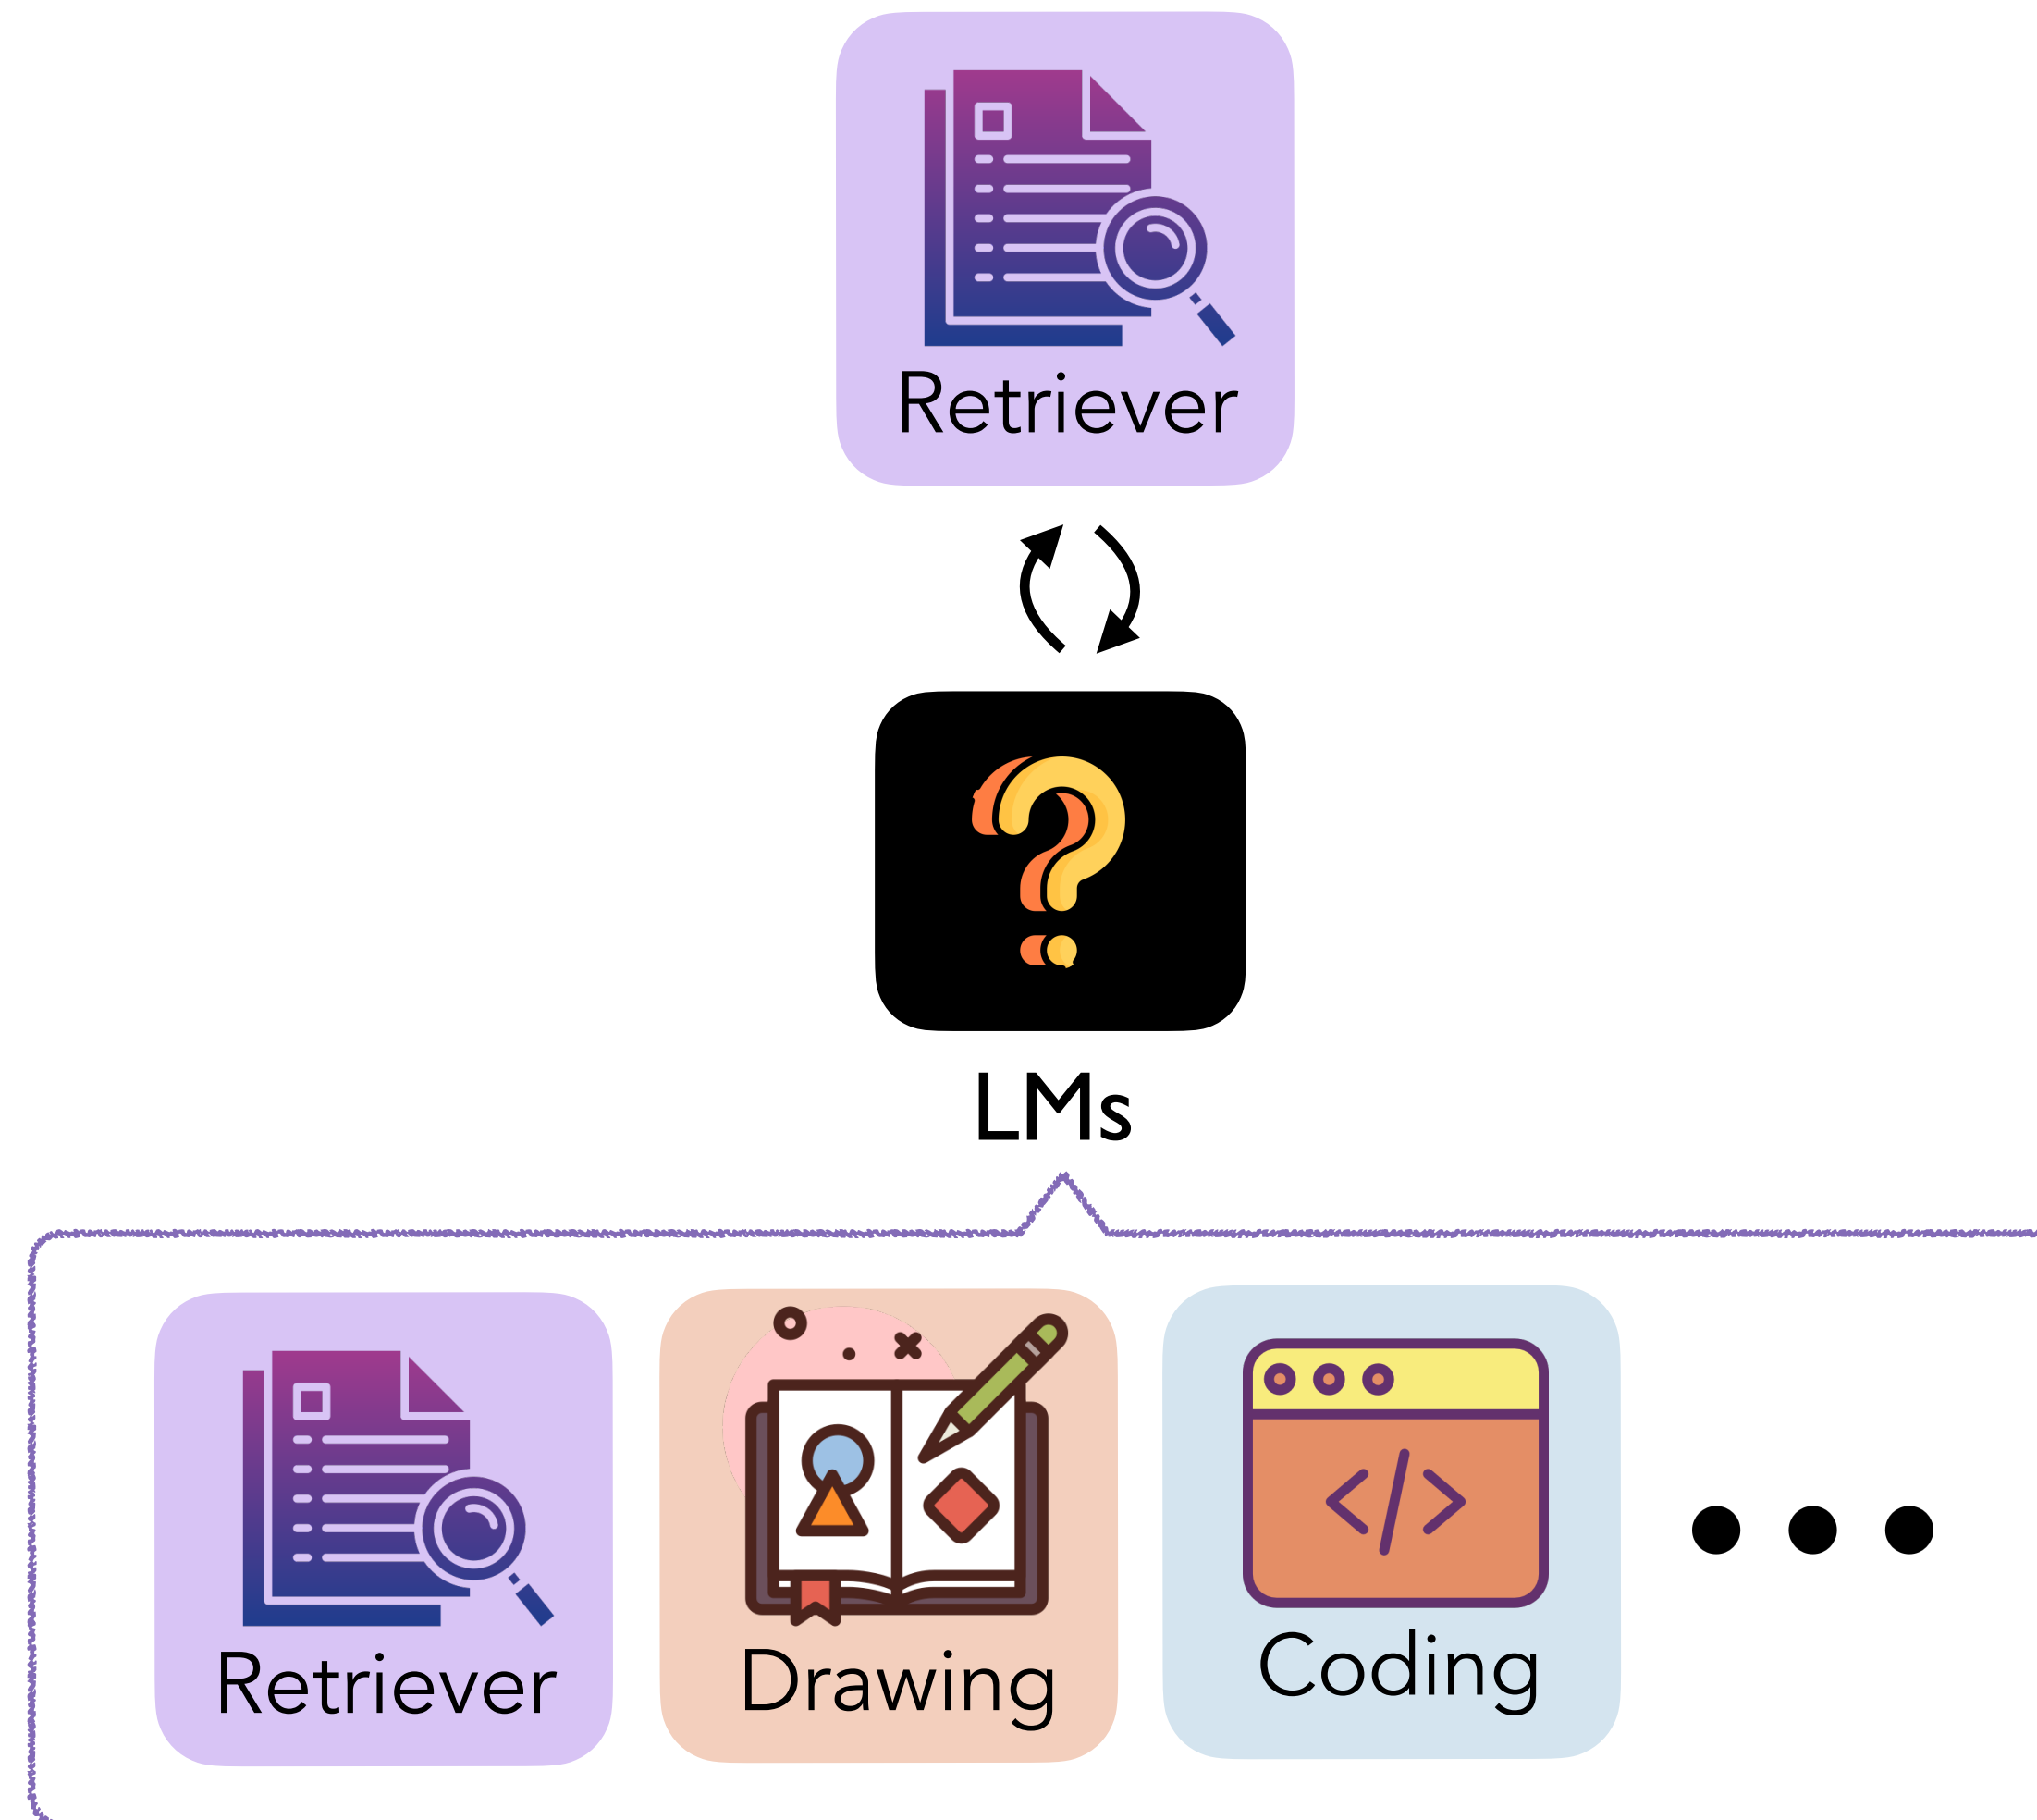
Consistent performance improvement

(c) HQ (Open)

# What's Next?

- How similar should the documents be in the same context?
- Does In-Context Pretraining work for continual pretraining?

# Augmented Models



Pretraining

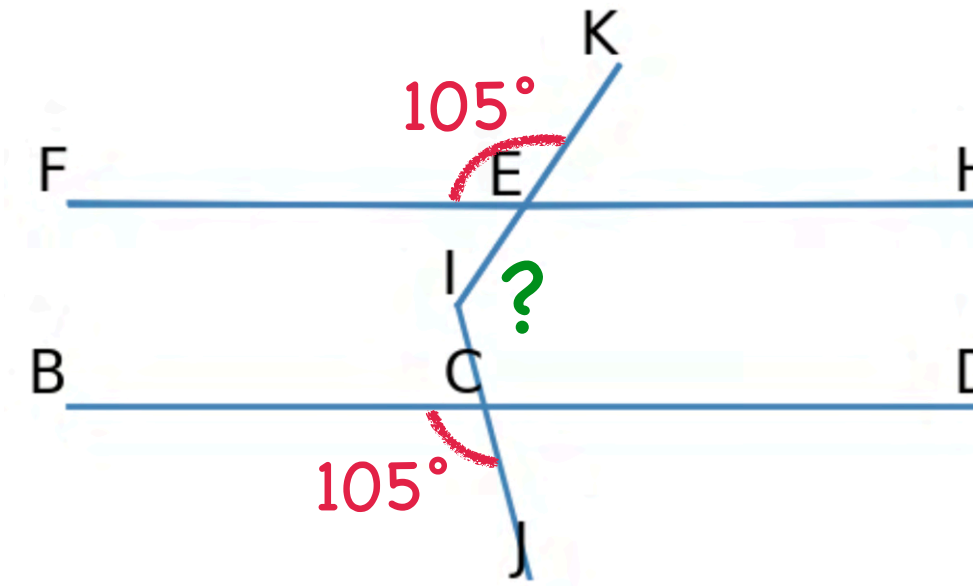
Shi, et al. ICLR 2024 Spotlight

Inference

Hu\*, Shi\* et al., NeurIPS 2024



Given  $\angle BCJ = 105^\circ$ ,  
 $\angle KEF = 105^\circ$ . Find  $\angle EIC$



From the given image, we are provided with two angles:

1.  $\angle BCJ = 105^\circ$
2.  $\angle KEF = 105^\circ$

We need to find the measure of  $\angle EIC$ .

Looking at the diagram, it seems that  $\angle EIC$  is the angle formed by lines CE and CI.  
Let's break this down:

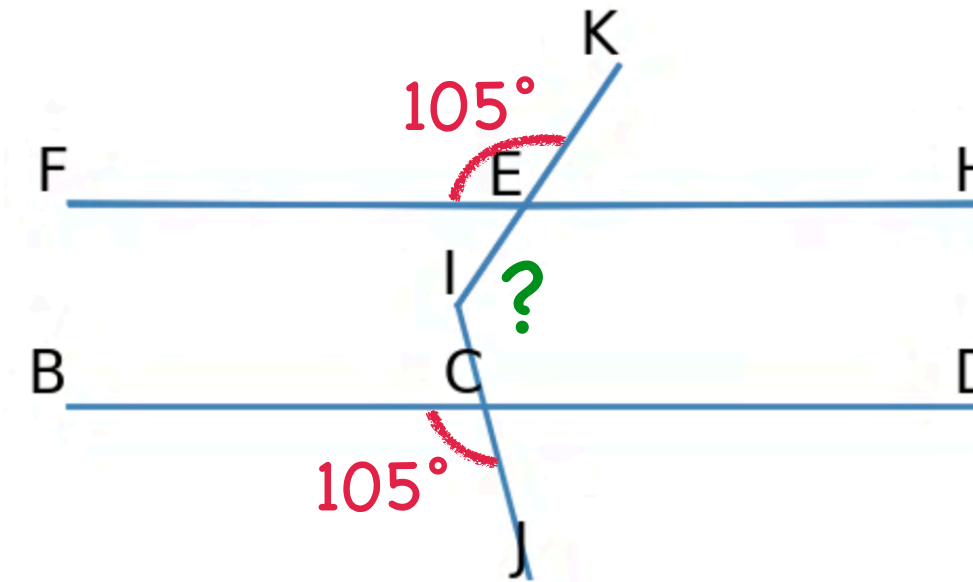
1. **Angle relationships:** The lines BC and CD are parallel, and similarly, the lines FE and HD are parallel. The transversal KE cuts both sets of parallel lines.
  - Since  $\angle BCJ$  and  $\angle KEF$  are both  $105^\circ$  and form corresponding angles with the transversal, they imply the parallelism of lines.
2. **Vertically opposite angles:** Since  $\angle BCJ = 105^\circ$ , the vertically opposite angle  $\angle EIC$  must also be  $105^\circ$  because vertically opposite angles are always equal.

Therefore, the measure of  $\angle EIC$  is  **$105^\circ$** .





Given  $\angle BCJ = 105^\circ$ ,  
 $\angle KEF = 105^\circ$ . Find  $\angle EIC$



From the given image, we are provided with two angles:

1.  $\angle BCJ = 105^\circ$
2.  $\angle KEF = 105^\circ$

We need to find the measure of  $\angle EIC$ .

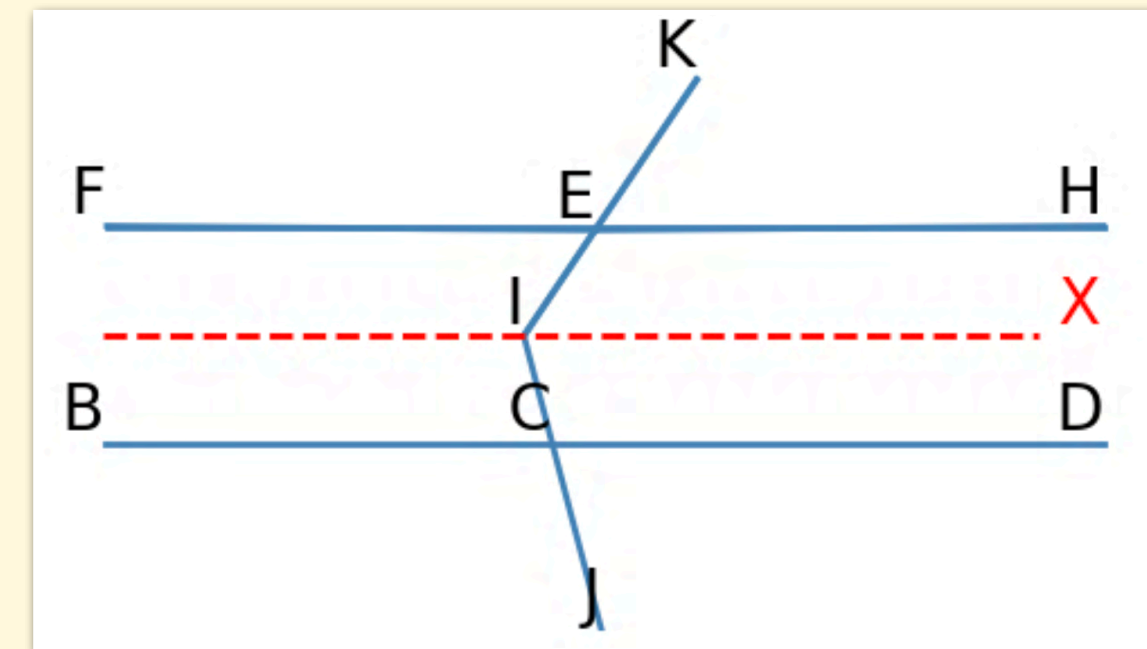
Looking at the diagram, it seems that  $\angle EIC$  is the angle formed by lines CE and CI.  
Let's break this down:

1. **Angle relationships:** The lines BC and CD are parallel, and similarly, the lines FE and HD are parallel. The transversal KE cuts both sets of parallel lines.
  - Since  $\angle BCJ$  and  $\angle KEF$  are both  $105^\circ$  and form corresponding angles with the transversal, they imply the parallelism of lines.
2. **Vertically opposite angles:** Since  $\angle BCJ = 105^\circ$ , the vertically opposite angle  $\angle EIC$  must also be  $105^\circ$  because vertically opposite angles are always equal.

Therefore, the measure of  $\angle EIC$  is  **$105^\circ$** .



**Draw line IX parallel to EH**



Given  $\angle BCJ = \angle KEF = 105^\circ$ , their supplementary angles,  $\angle KEH = \angle DCJ = 180^\circ - 105^\circ = 75^\circ$ .

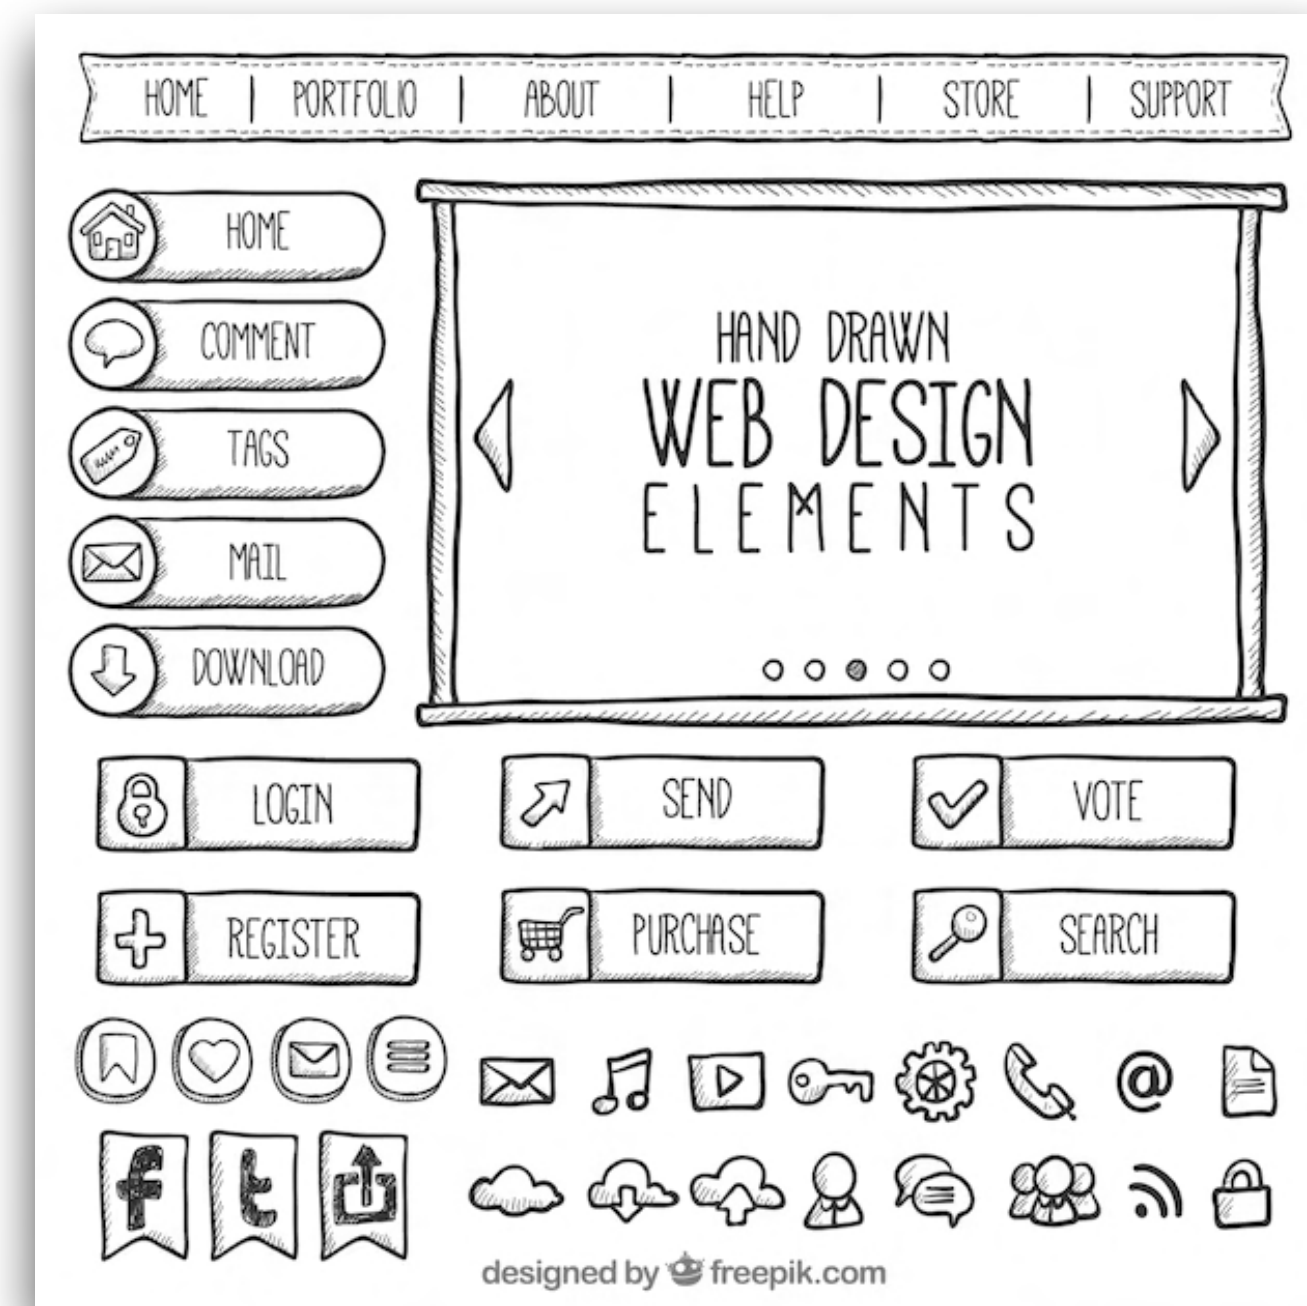
Given BD parallel to IX, it follows that  $\angle EIX = \angle KEH$ , and  $\angle XIC = \angle DCJ$ , both measuring  $75^\circ$ .

Therefore,  $\angle EIC = \angle EIX + \angle XIC = 150^\circ$



# Sketching 🖋️

Sketching is a fundamental human activity, serving as a versatile tool for **communication**, **ideation**, and **problem-solving**



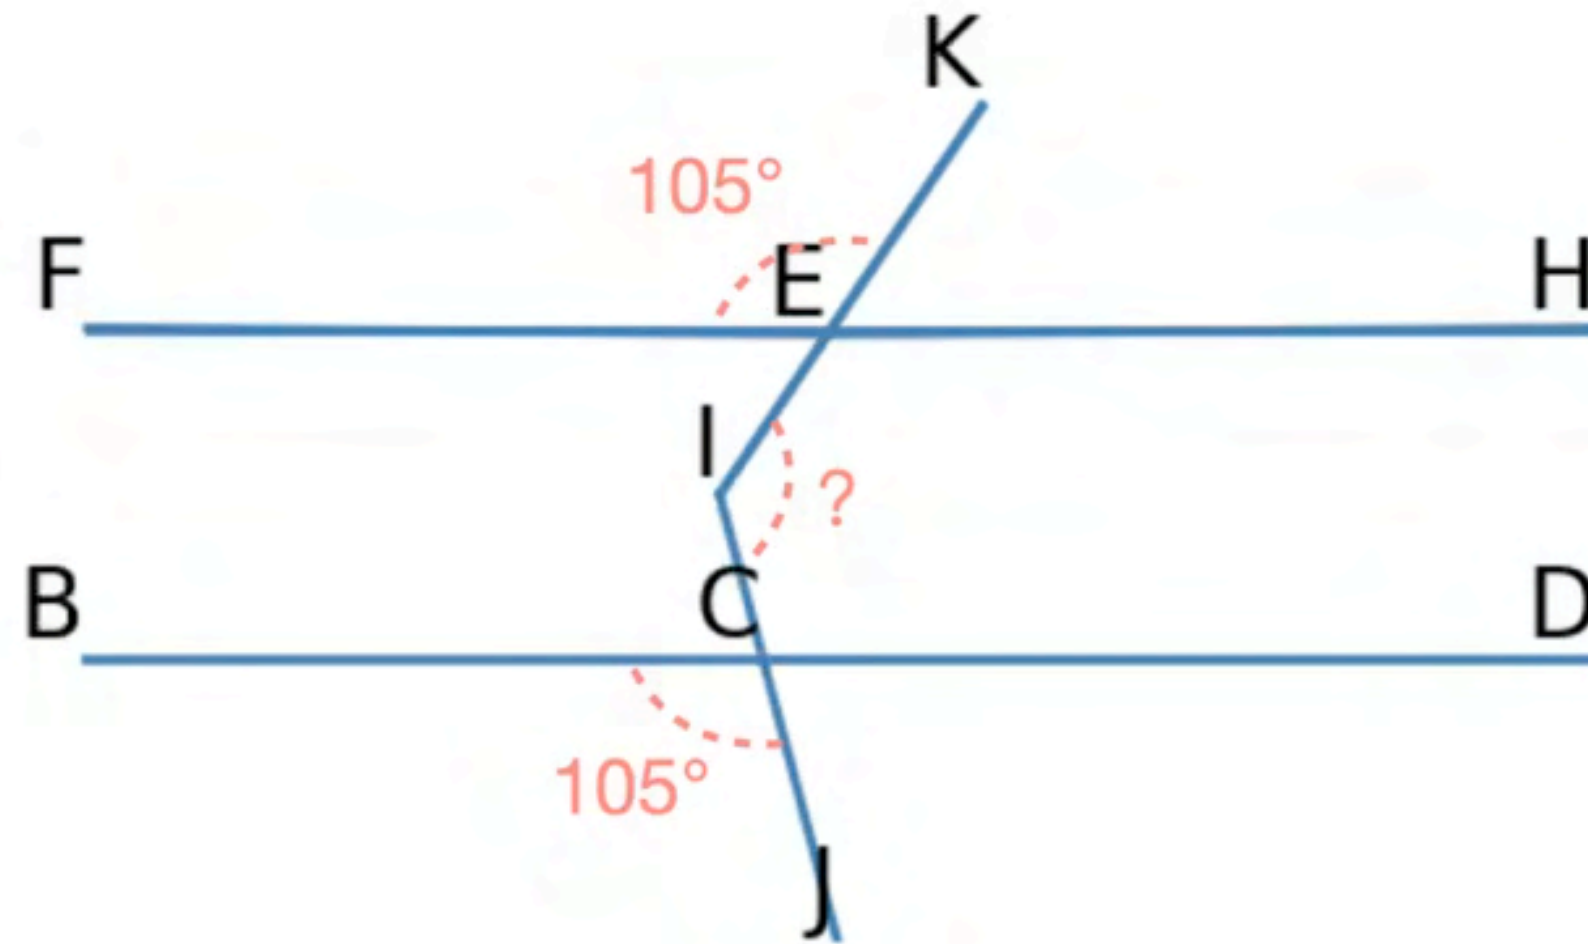
Why not let multimodal LMs do so?

# Visual Reasoning w/ Sketchpad

Augmenting multimodal LMs with a visual canvas and drawing tools 🖋️

**Query:** Given  $\angle BCJ = 105^\circ$ ,  $\angle KEF = 105^\circ$ . Find  $\angle EIC$

**Input Image:**





# Visual Sketchpad: Recipe

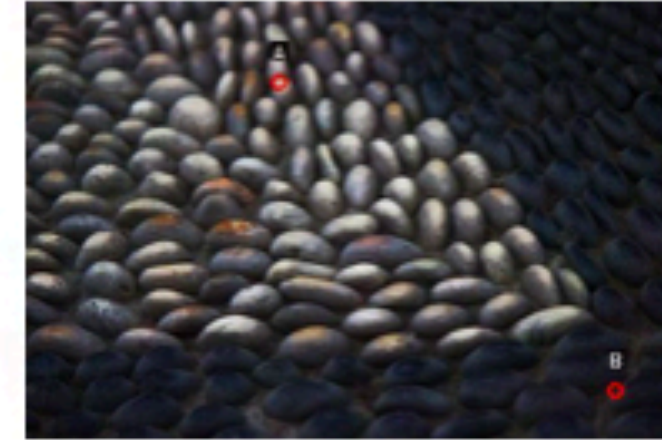
## Sketching via Code Generation

# Visual Sketchpad: Recipe

## Sketching via Code Generation



Which point is closer to the camera, A or B?



A

B

# Visual Sketchpad: Recipe

## Sketching via Code Generation



Which point is closer to the camera, A or B?



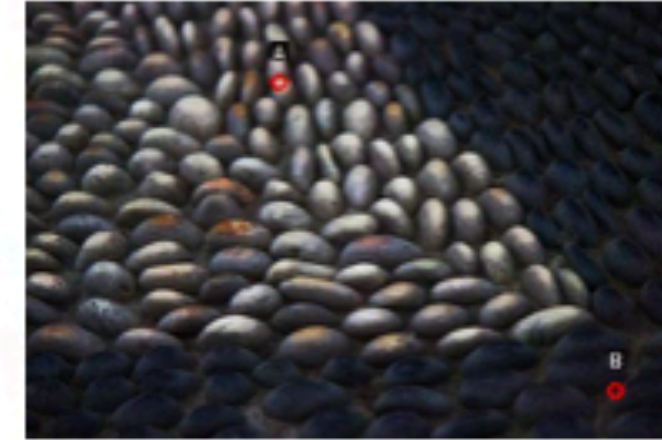
**THOUGHT 1:** I can use a depth estimation tool that provides a depth map with colors indicating distances.

# Visual Sketchpad: Recipe


## Sketching via Code Generation



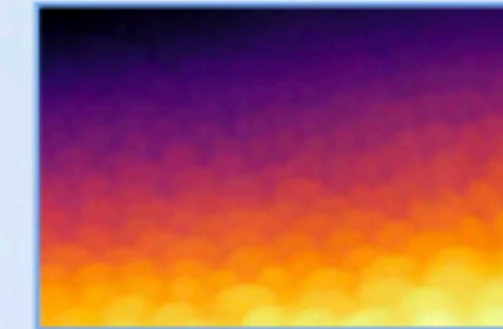
Which point is closer to the camera, A or B?



**THOUGHT 1:** I can use a depth estimation tool that provides a depth map with colors indicating distances.

**ACTION 1:**  Call depth estimation

```
depth_map = depth(image)  
display(depth_map)
```

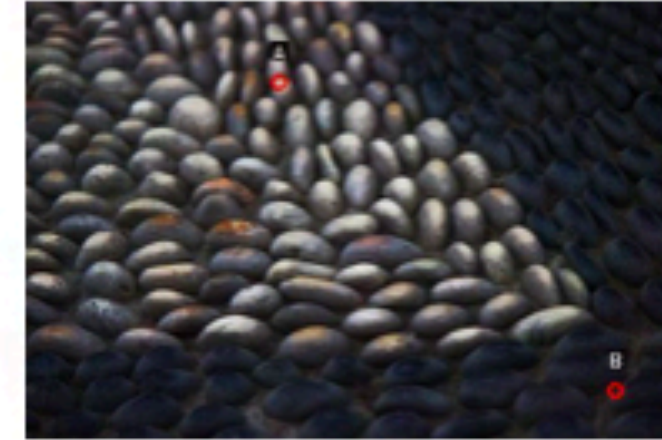


# Visual Sketchpad: Recipe


## Sketching via Code Generation



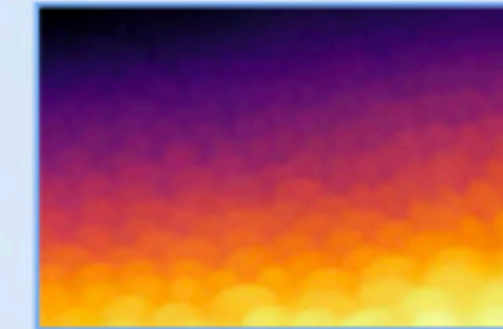
Which point is closer to the camera, A or B?



**THOUGHT 1:** I can use a depth estimation tool that provides a depth map with colors indicating distances.

**ACTION 1:**  Call depth estimation

```
depth_map = depth(image)
display(depth_map)
```



**THOUGHT 2:** The depth map uses colors from purple (farther) to yellow (closer). I need to overlay the original image with the depth map

# Visual Sketchpad: Recipe


## Sketching via Code Generation



Which point is closer to the camera, A or B?



**THOUGHT 1:** I can use a depth estimation tool that provides a depth map with colors indicating distances.

**ACTION 1:**  Call depth estimation

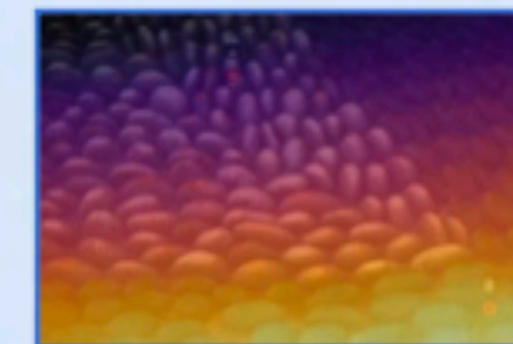
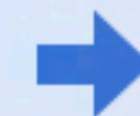
```
depth_map = depth(image)
display(depth_map)
```



**THOUGHT 2:** The depth map uses colors from purple (farther) to yellow (closer). I need to overlay the original image with the depth map

**ACTION 2:**  Call overlay

```
overlay = overlay_images(depth_map, image)
display(overlay)
```



# Visual Sketchpad: Recipe

## Sketching via Code Generation


visual reasoning chains



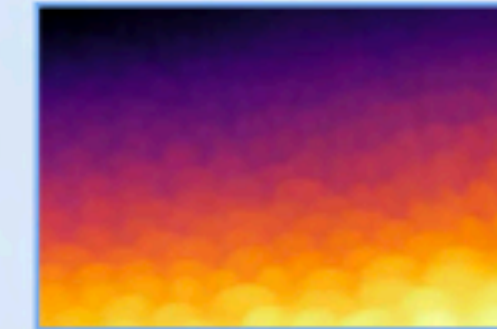
Which point is closer to the camera, A or B?



**THOUGHT 1:** I can use a depth estimation tool that provides a depth map with colors indicating distances.

**ACTION 1:**  Call depth estimation

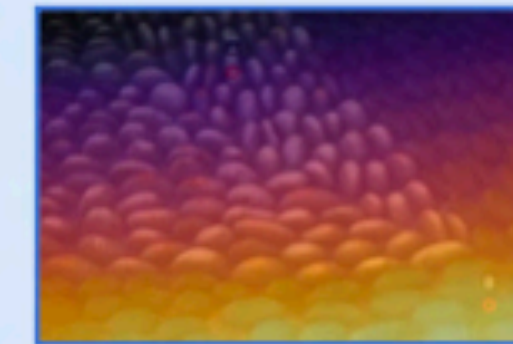
```
depth_map = depth(image)
display(depth_map)
```



**THOUGHT 2:** The depth map uses colors from purple (farther) to yellow (closer). I need to overlay the original image with the depth map

**ACTION 2:**  Call overlay

```
overlay = overlay_images(depth_map, image)
display(overlay)
```

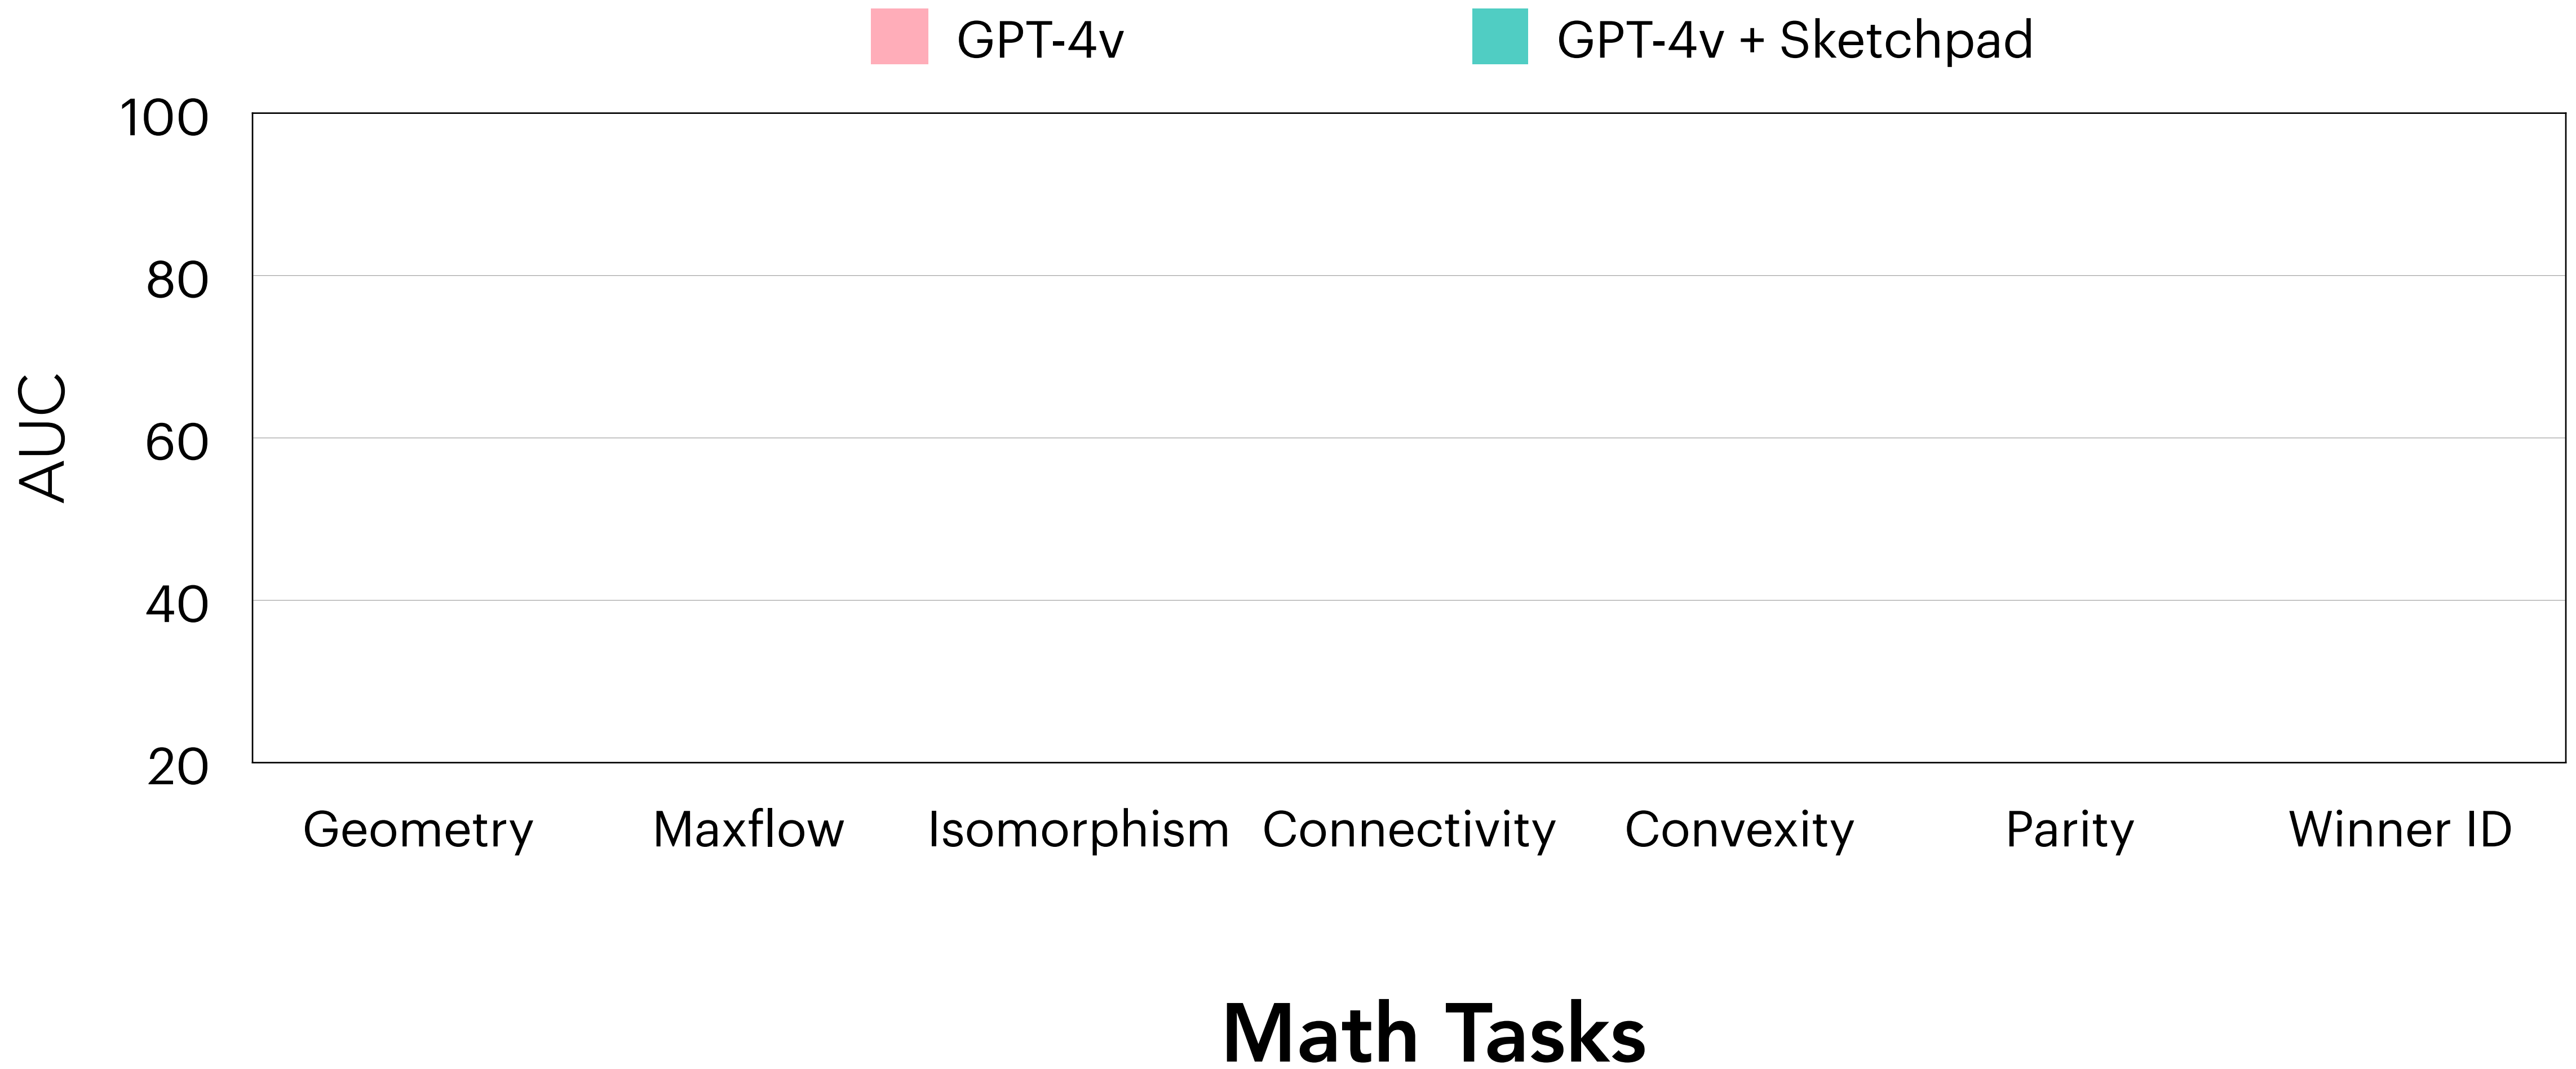


**THOUGHT 3:** It is evident that point B, in the yellow region, is closer to the camera compared to point A, which is in the purple region

# Results

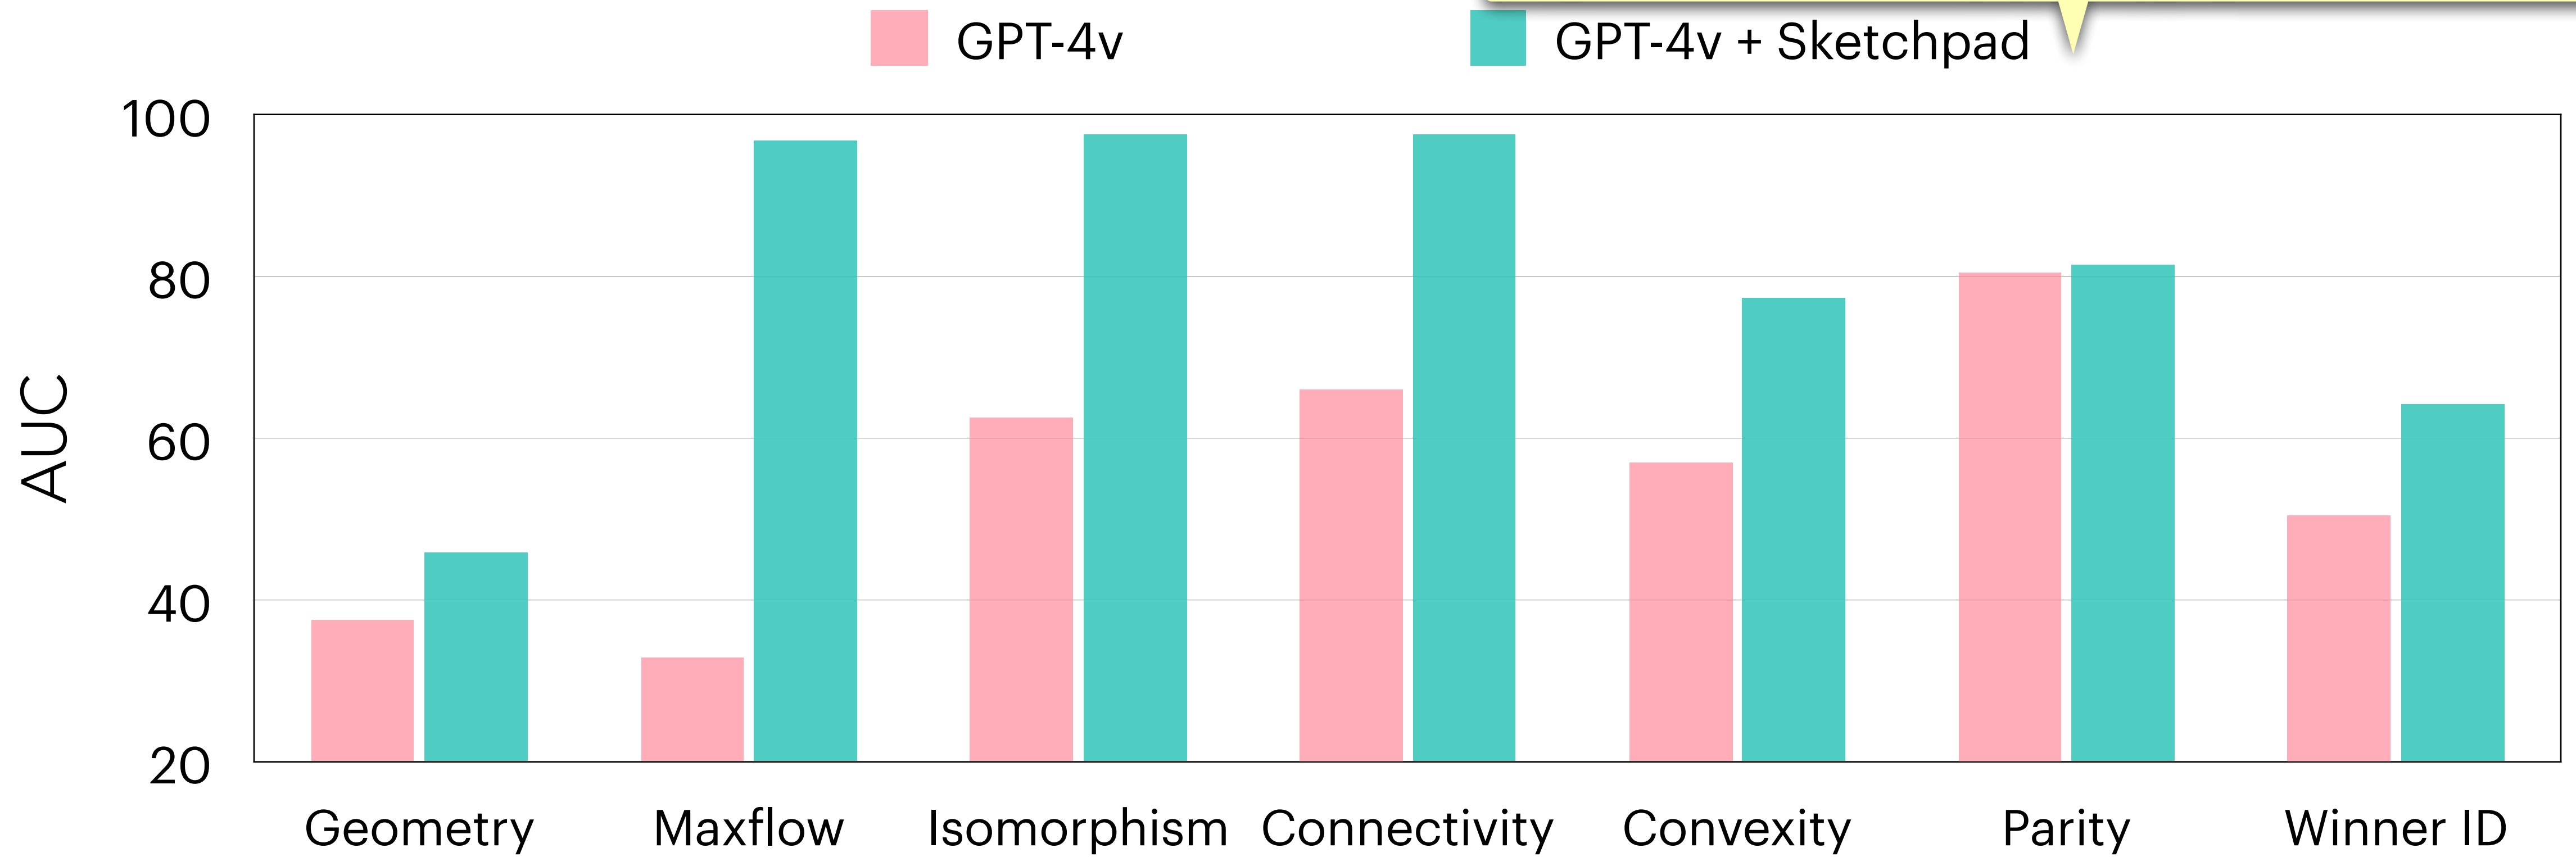


# Results



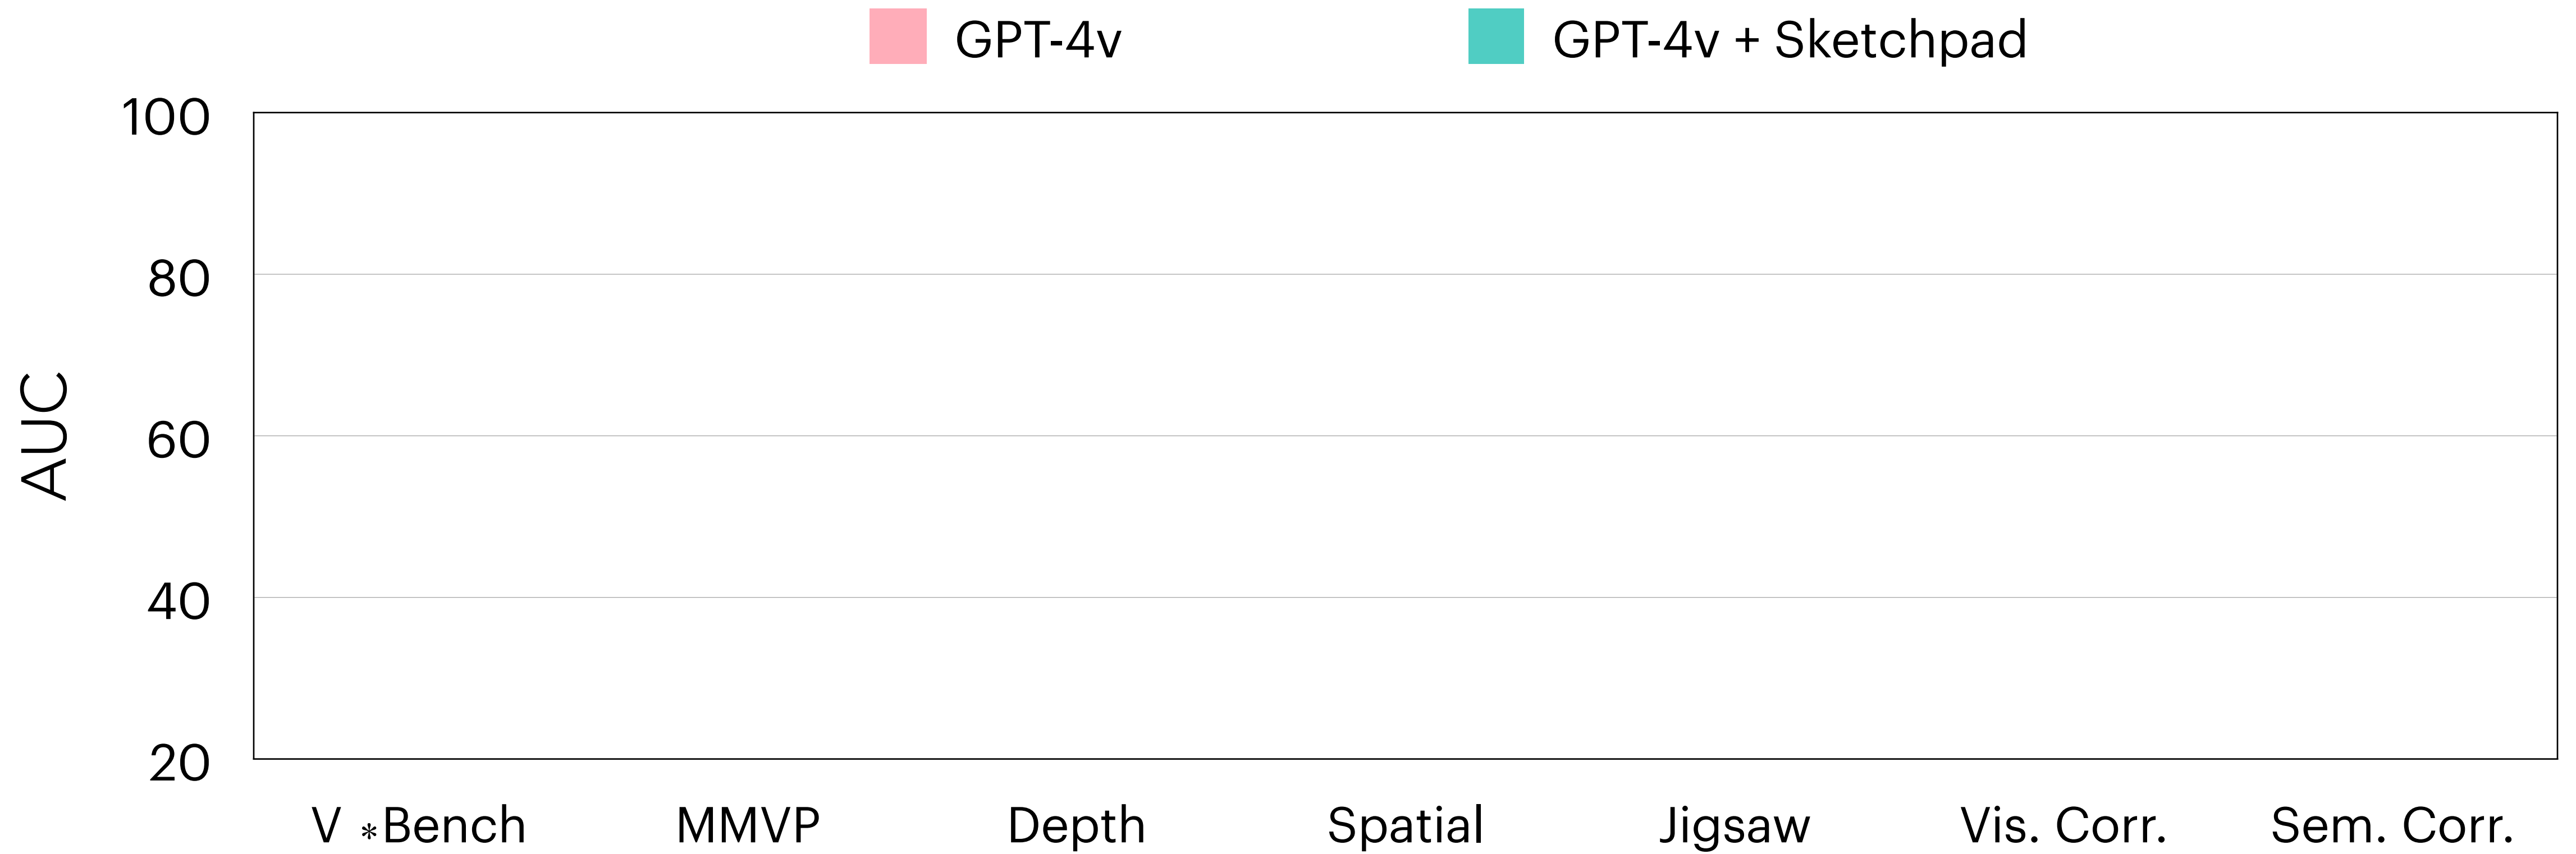
# Results

12.7% improvement on average



Math Tasks

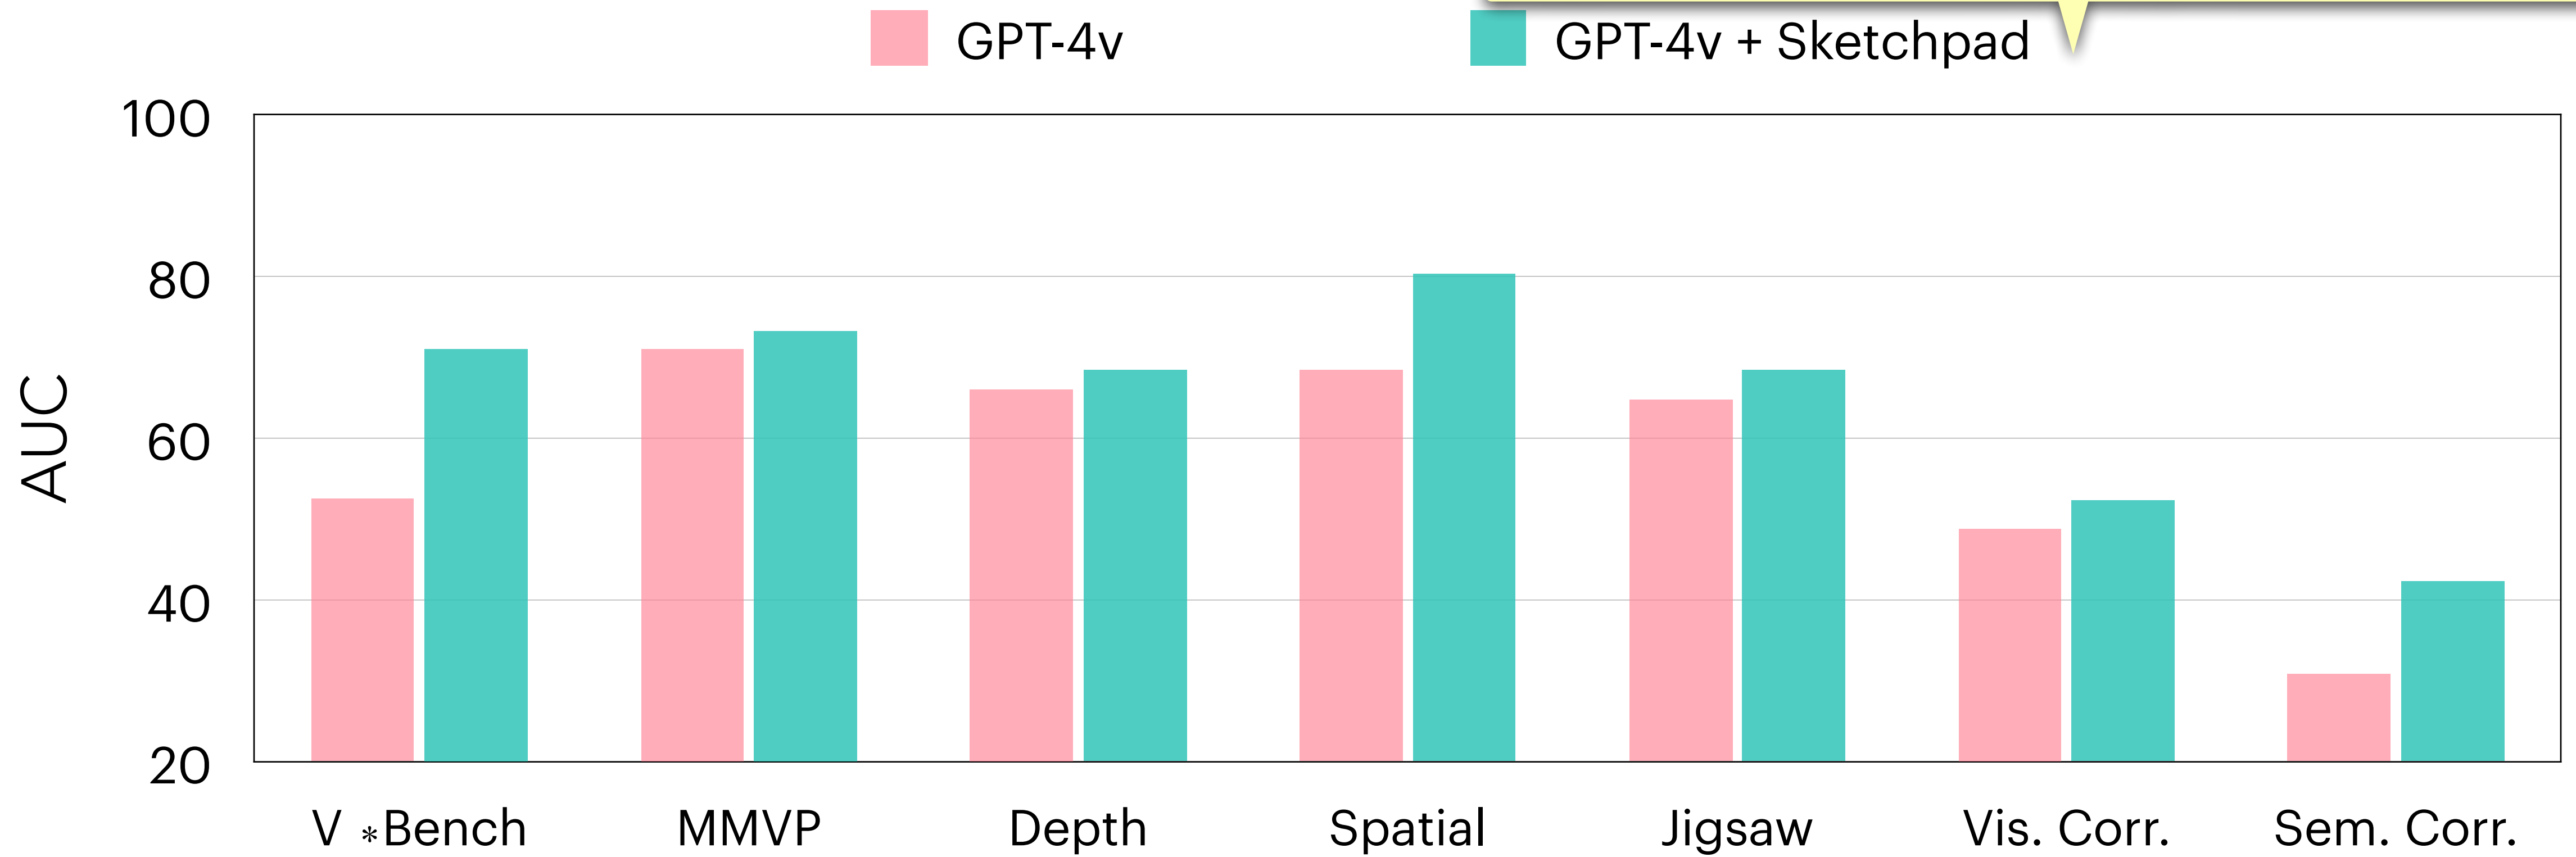
# Results



## Visual Reasoning Tasks

# Results

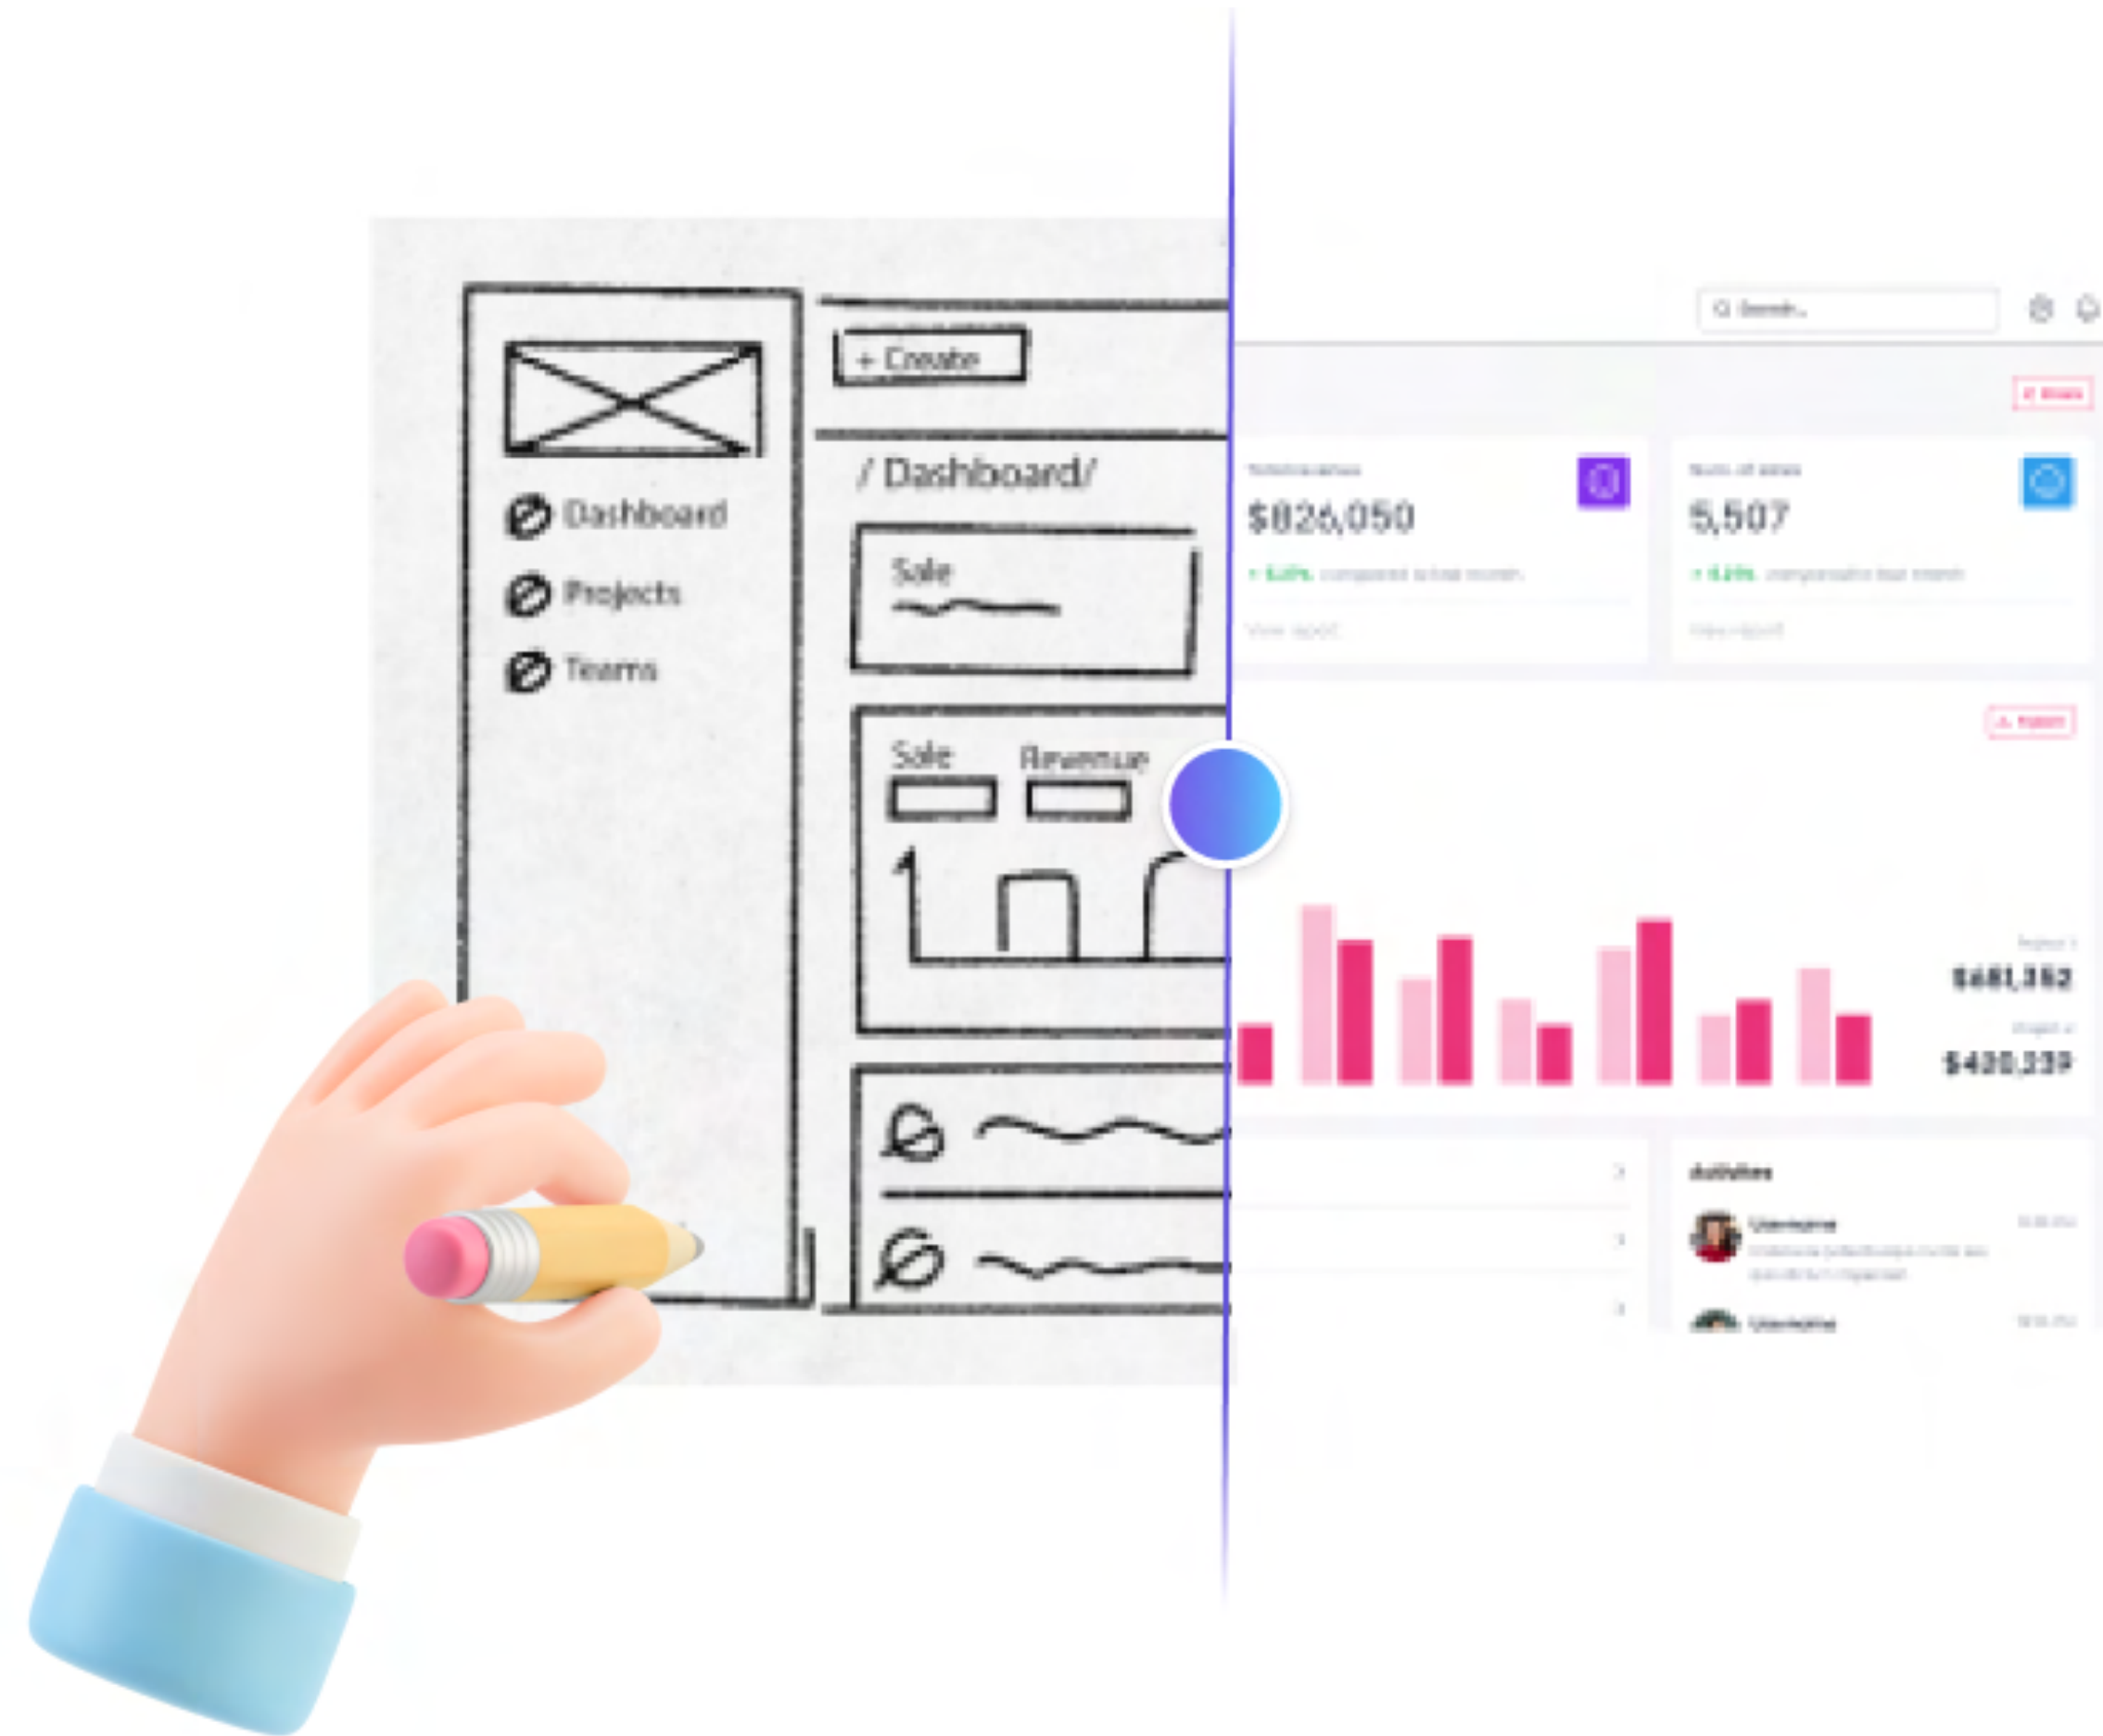
8.6% improvement on average



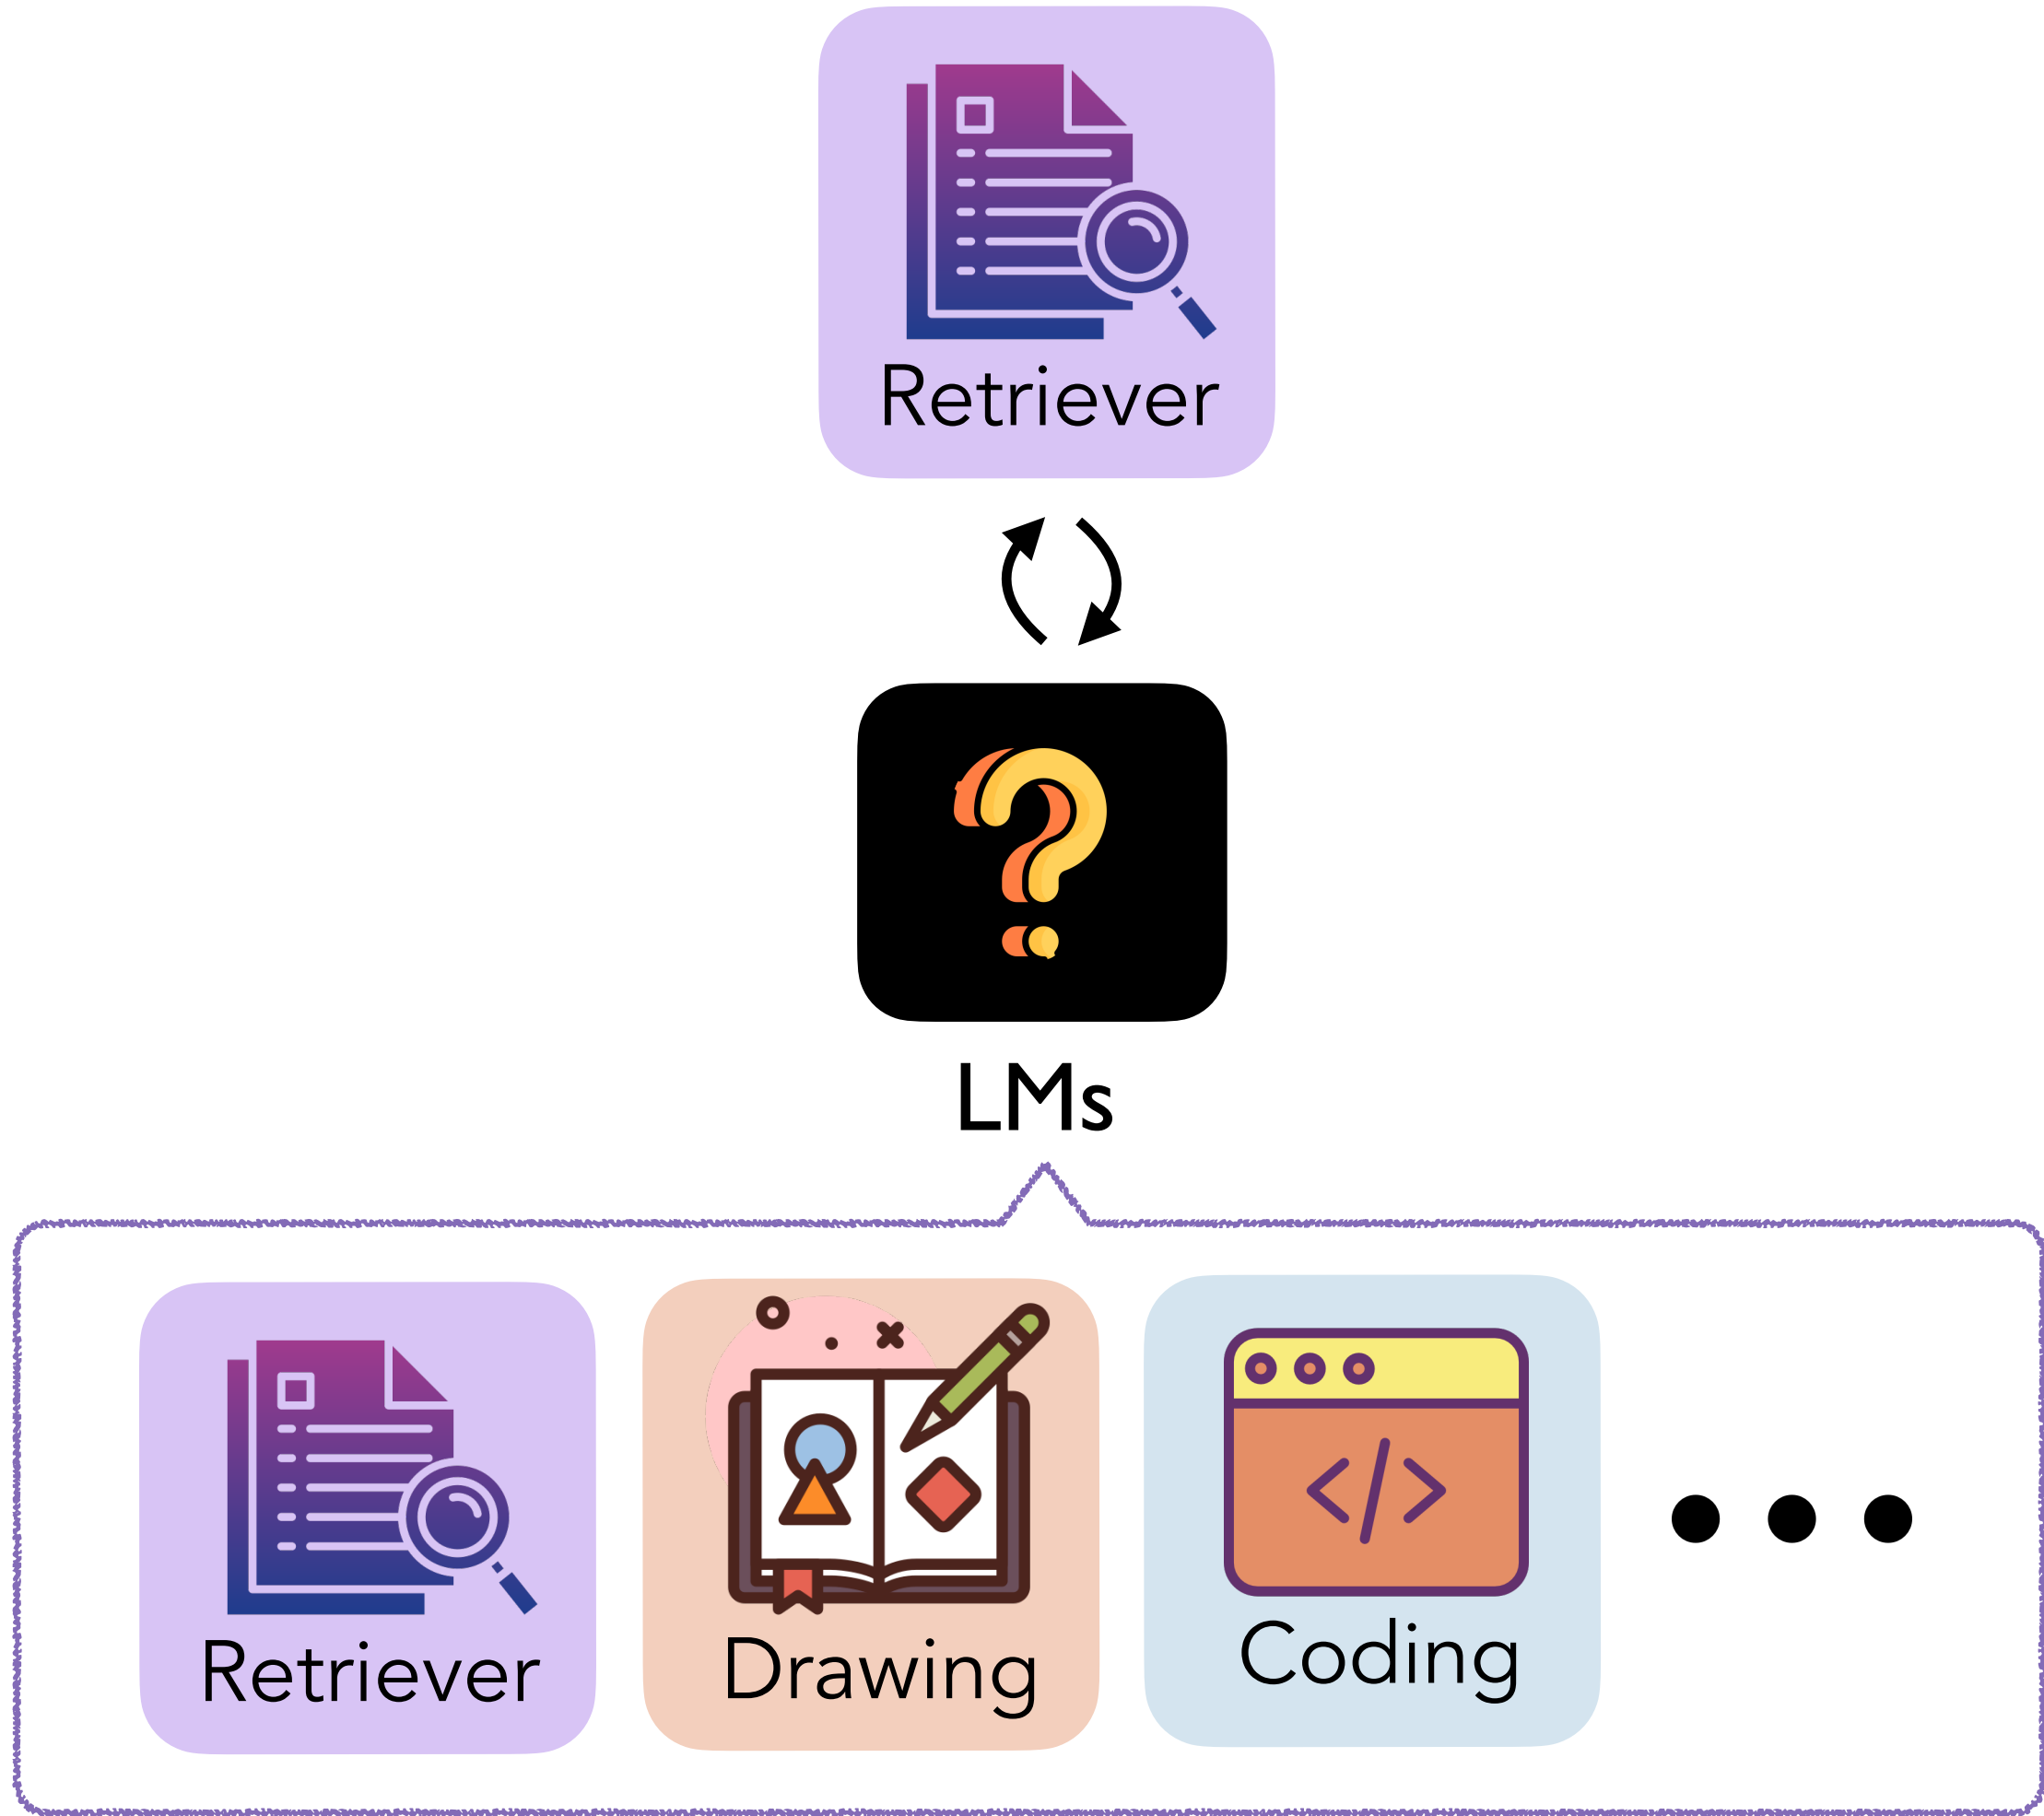
## Visual Reasoning Tasks

# What's Next?

Sketch to UI Design with Multimodal LMs



# Summary: Augmented Models



Pretraining

Shi, et al. ICLR 2024 Spotlight

Inference

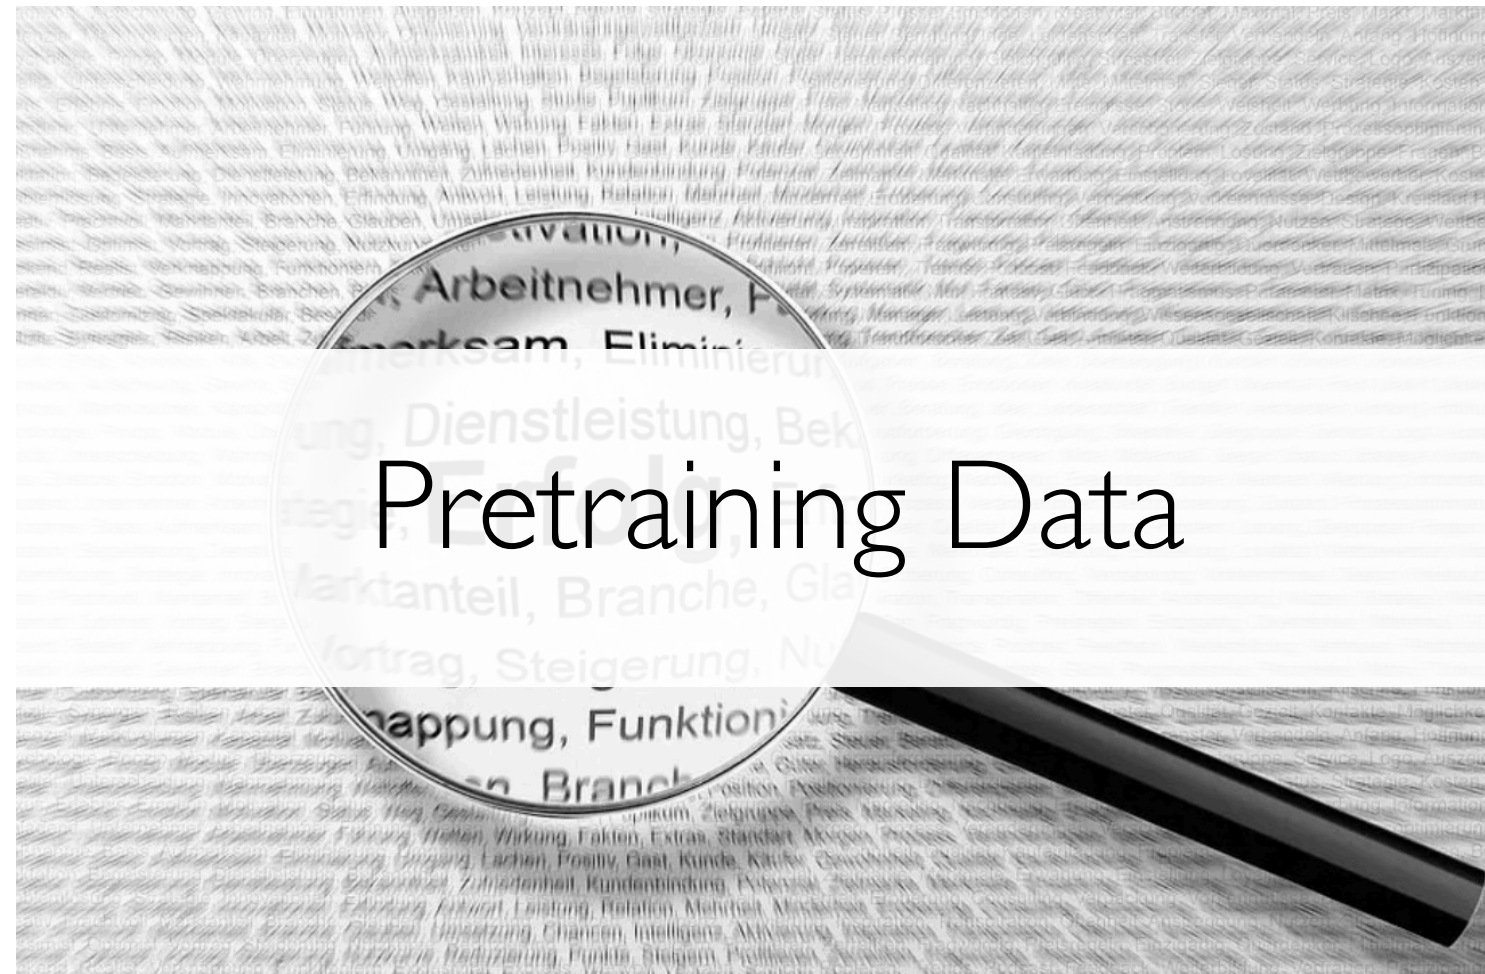
Hu\*, Shi\* et al., NeurIPS 2024

# Beyond Monolithic Language Models

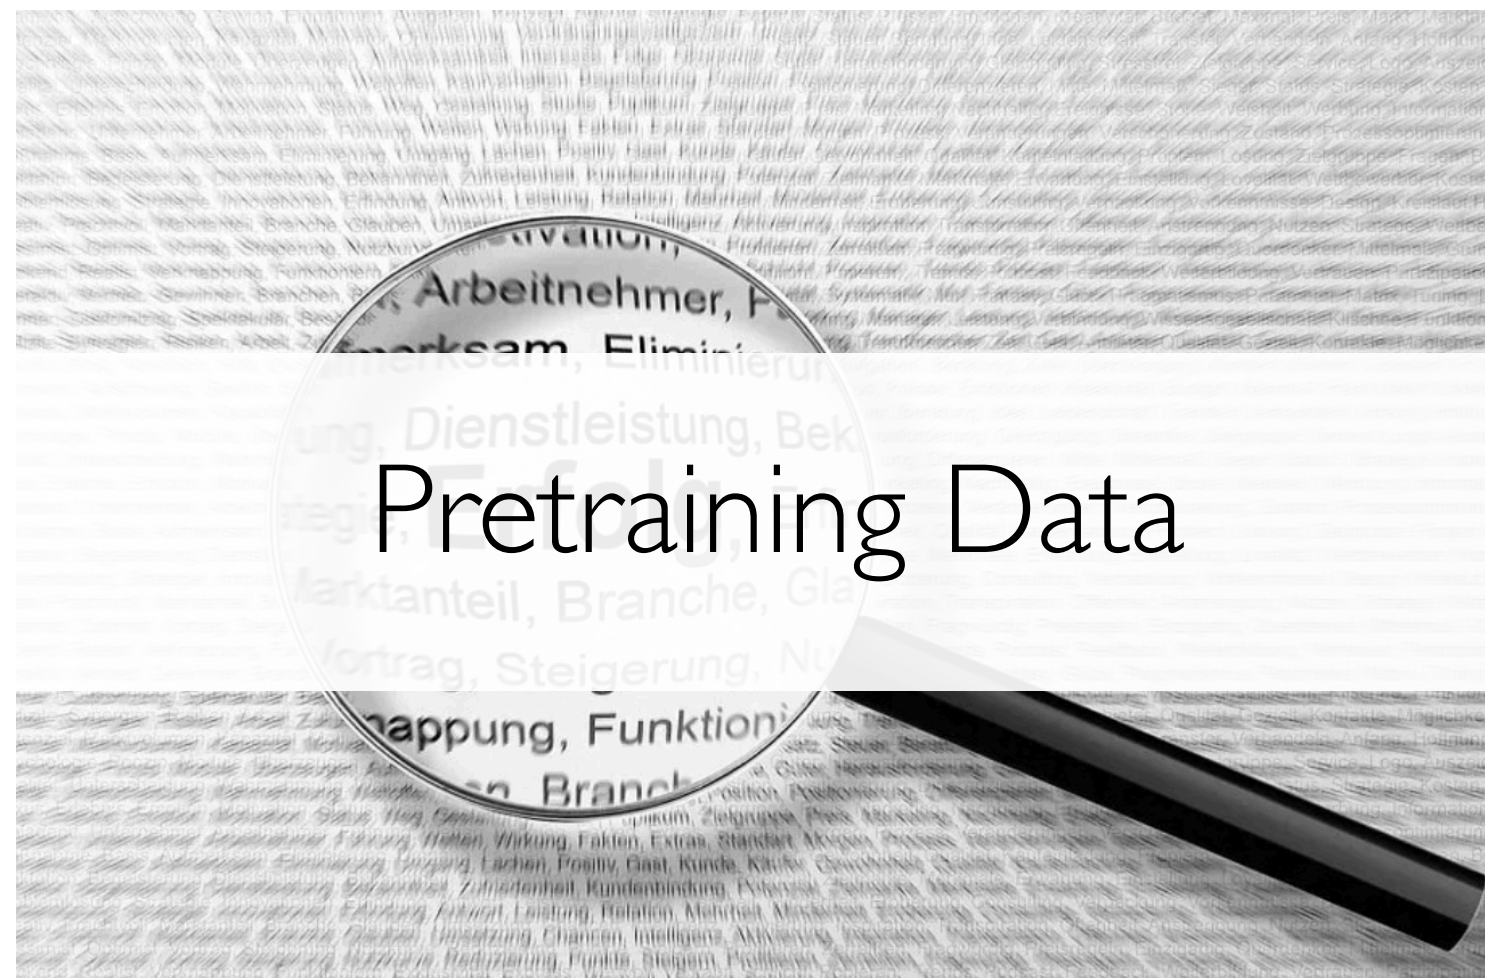
*Augmented Models* 

*Data Modularity*









# Copyright Risks in LMs

## The Times Sues OpenAI and Microsoft Over A.I. Use of Copyrighted Work

Millions of articles from The New York Times were used to train chatbots that now compete with it, the lawsuit said.

Dec. 27, 2023

 [World](#) [Business](#) [Markets](#) [Sustainability](#) [Legal](#) [Breakingviews](#) [Technology](#) [Investigations](#)

[Litigation](#) | [Copyright](#) | [Litigation](#) | [Technology](#) | [Intellectual Property](#)

## Music publishers ask court to halt AI company Anthropic's use of lyrics

By Dawn Chmielewski

November 17, 2023 11:39 AM EST · Updated 7 months ago

# Not Just in LMs...

He\*, Huang\*, **Shi\***, et al. Under Review, 2024

# Not Just in LMs...



Videogame plumber

 DALL-E



He\*, Huang\*, **Shi\***, et al. Under Review, 2024

# Not Just in LMs...



Videogame plumber

 DALL-E



Superhero Gotham

 DALL-E



He\*, Huang\*, **Shi\***, et al. Under Review, 2024

*(Slides adapted from Yangsibo's talk:  
Open Technical Questions in GenAI Copyright)*

# Not Just in LMs...



Videogame plumber



Superhero Gotham

For 20 out of 50 copyrighted characters, we can generate them using <5 keywords (w/o character names)

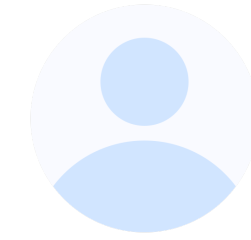


He\*, Huang\*, Shi\*, et al. Under Review, 2024

# Not Just in LMs...



Videogame plumber



Superhero Gotham

---

## Fantastic Copyrighted Beasts and How (Not) to Generate Them

---

**Luxi He<sup>\*1</sup>   Yangsibo Huang<sup>\*1</sup>   Weijia Shi<sup>\*2</sup>**  
**Tinghao Xie<sup>1</sup>   Haotian Liu<sup>3</sup>   Yue Wang<sup>4</sup>   Luke Zettlemoyer<sup>2</sup>**  
**Chiyuan Zhang   Danqi Chen<sup>1</sup>   Peter Henderson<sup>1</sup>**

<sup>1</sup>Princeton University   <sup>2</sup>University of Washington  
<sup>3</sup>University of Wisconsin-Madison   <sup>4</sup>University of Southern California

<https://copycat-eval.github.io/>

He\*, Huang\*, Shi\*, et al. Under Review, 2024

How can we *mitigate* copyright risks?



# Copyright Takedown in Search Engine

Takedown request









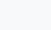
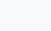
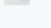
## Removing Content From Google

This page will help you get to the right place to report content that you would like removed from Google's services under applicable laws. Providing us with complete information will help us investigate your inquiry.

If you have non-legal issues that concern Google's [Terms of Service](#) or Product Policies, please visit <http://support.google.com>

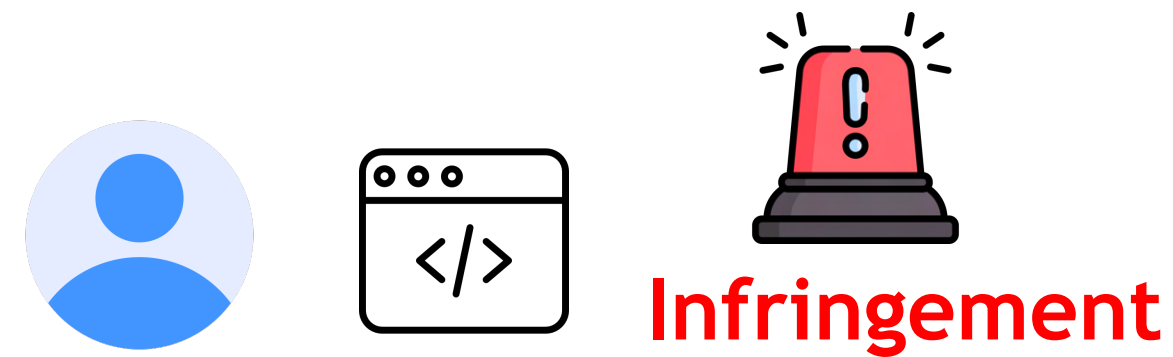
We ask that you submit a separate notice for each Google service where the content appears.

What Google product does your request relate to?

-  Google Search
-  Blogger/Blogspot
-  Google Maps and related products
-  Google Play: Apps
-  YouTube
-  Google Images
-  A Google Ad
-  Drive and Docs
-  Google Photos and Picasa Web Albums
-  Google Shopping
-  Google Play: Music
- See more products



Wei,\* Shi\*, et al. NeurIPS 2024



Remove the website

Google removes content in 30 days

# Copyright Takedown in Search Engine



[about](#) [issues & policy](#) [copyright law](#) [resources](#) [programs](#) [news & events](#) [education](#) [get involved](#)

[videos](#)

[faqs](#)

[copyright law explained](#)

[ai and copyright](#)

[ccb explained](#)

[copyright law by industry](#)

[copyright courses](#)

## The DMCA Notice and Takedown Process

In passing the notice and takedown provisions in the Digital Millennium Copyright Act (DMCA), Congress intended to encourage copyright owners and service providers to work together to combat existing and future forms of online copyright infringement.

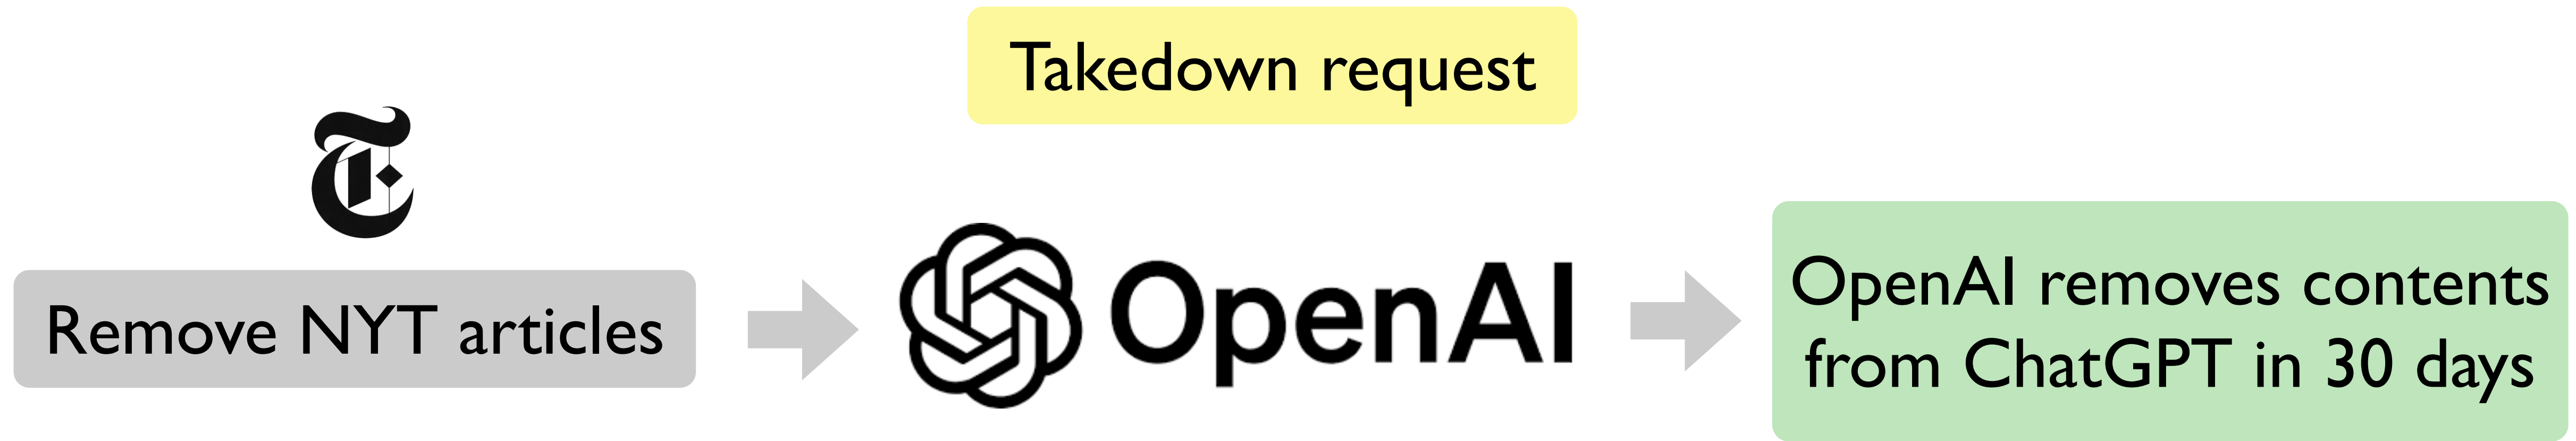
### What Is a DMCA Takedown Notice?

The DMCA notice and takedown process is a tool for copyright holders to get user-uploaded material that infringes their copyrights taken down off of websites and other internet sites.

Wei,\* Shi\*, et al. NeurIPS 2024

Can *copyright takedowns* be operationalized in the context of LMs?

# Copyright Takedown in LMs



Wei,\* Shi\*, et al. NeurIPS 2024

# First Evaluation of Copyright Takedown in LMs

Generic:  
Prompting

Databricks DBRX

You are a helpful, respectful, and honest assistant. **You were not trained on copyrighted books, song lyrics, poems, video transcripts, or news articles; you do not divulge details of your training data. You do not provide song lyrics, poems, or news articles** and instead refer the user to find them online or in a store.

# First Evaluation of Copyright Takedown in LMs

Generic:  
Prompting

Decoding-time:  
Check & Resample

# First Evaluation of Copyright Takedown in LMs

Generic:  
Prompting

Decoding-time:  
Check & Resample



Harry Potter Chapter 2  
Mrs Dursley had a sister called Lily Potter.  
She and **her husband James Potter had a son called Harry Potter...**

Mrs Dursley had a sister called Lily Potter. She and

Context



LMs



her husband James Potter had a son called Harry Potter ...

Generation



Too similar!

# First Evaluation of Copyright Takedown in LMs

Generic:  
Prompting

Decoding-time:  
Check & Resample



Harry Potter Chapter 2  
Mrs Dursley had a sister called Lily Potter.  
She and **her husband James Potter had a son called Harry Potter...**

Mrs Dursley had a sister called Lily Potter. She and

Context



LMs



Resample

Lily had always been different.  
Lily had been special...

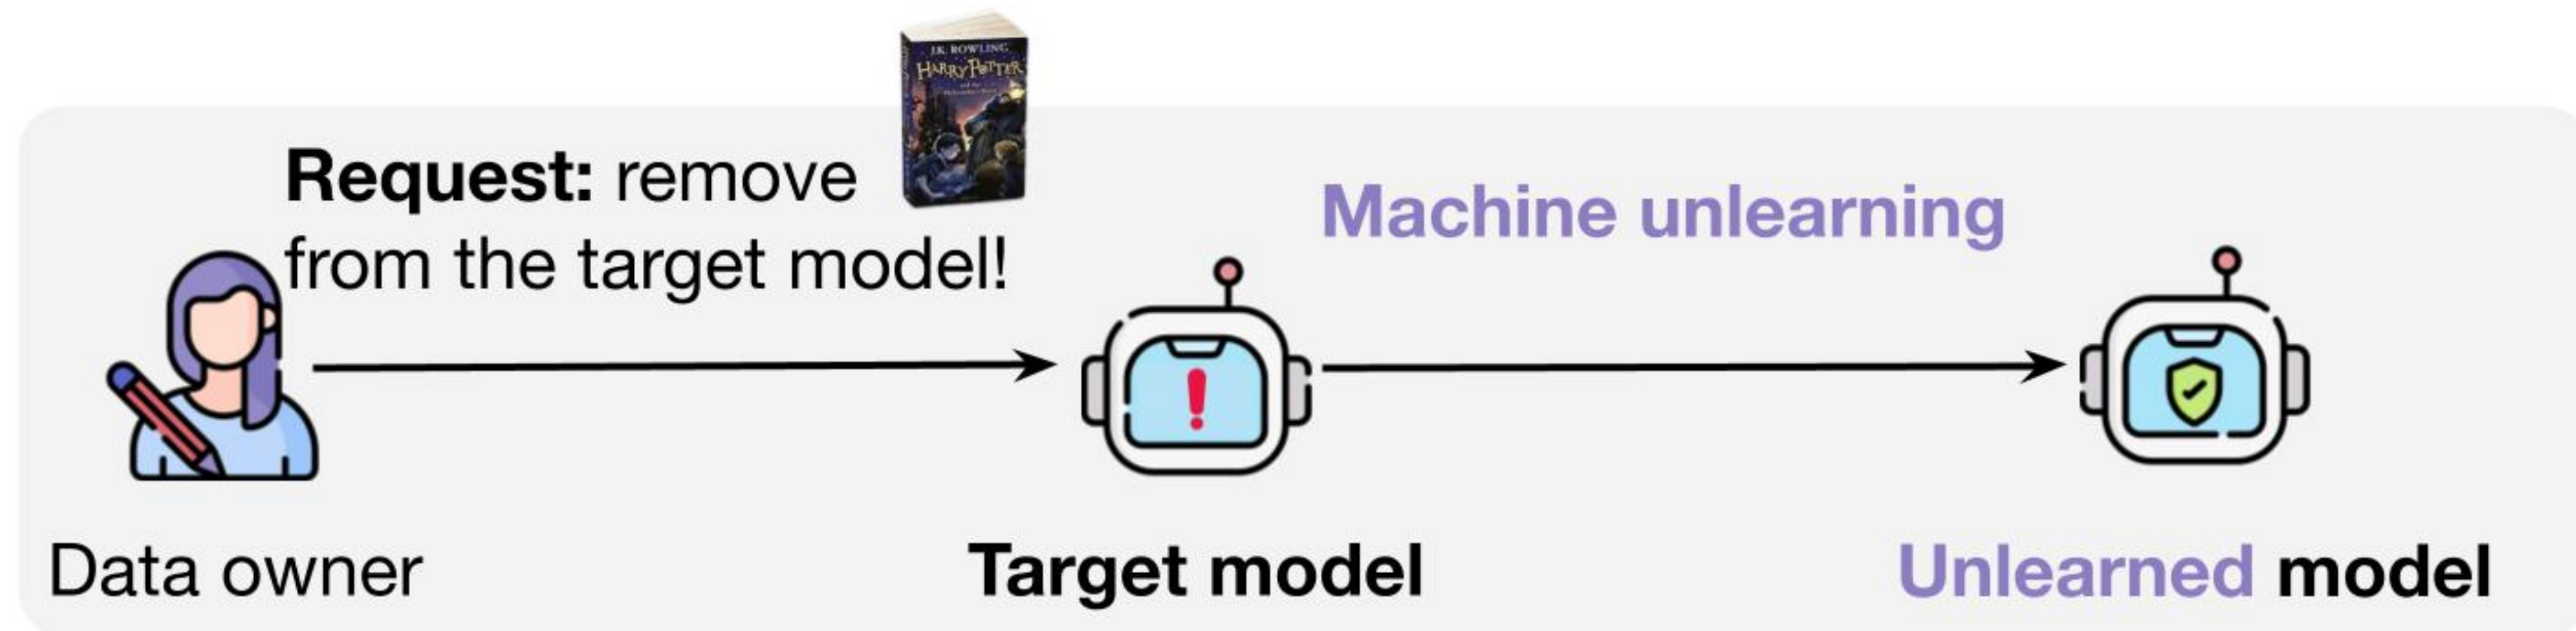


# First Evaluation of Copyright Takedown in LMs

Generic:  
Prompting

Decoding-time:  
Check & Resample

Training-based:  
Unlearning



Wei,\* Shi\*, et al. NeurIPS 2024

# First Evaluation of Copyright Takedown in LMs

Generic:  
Prompting

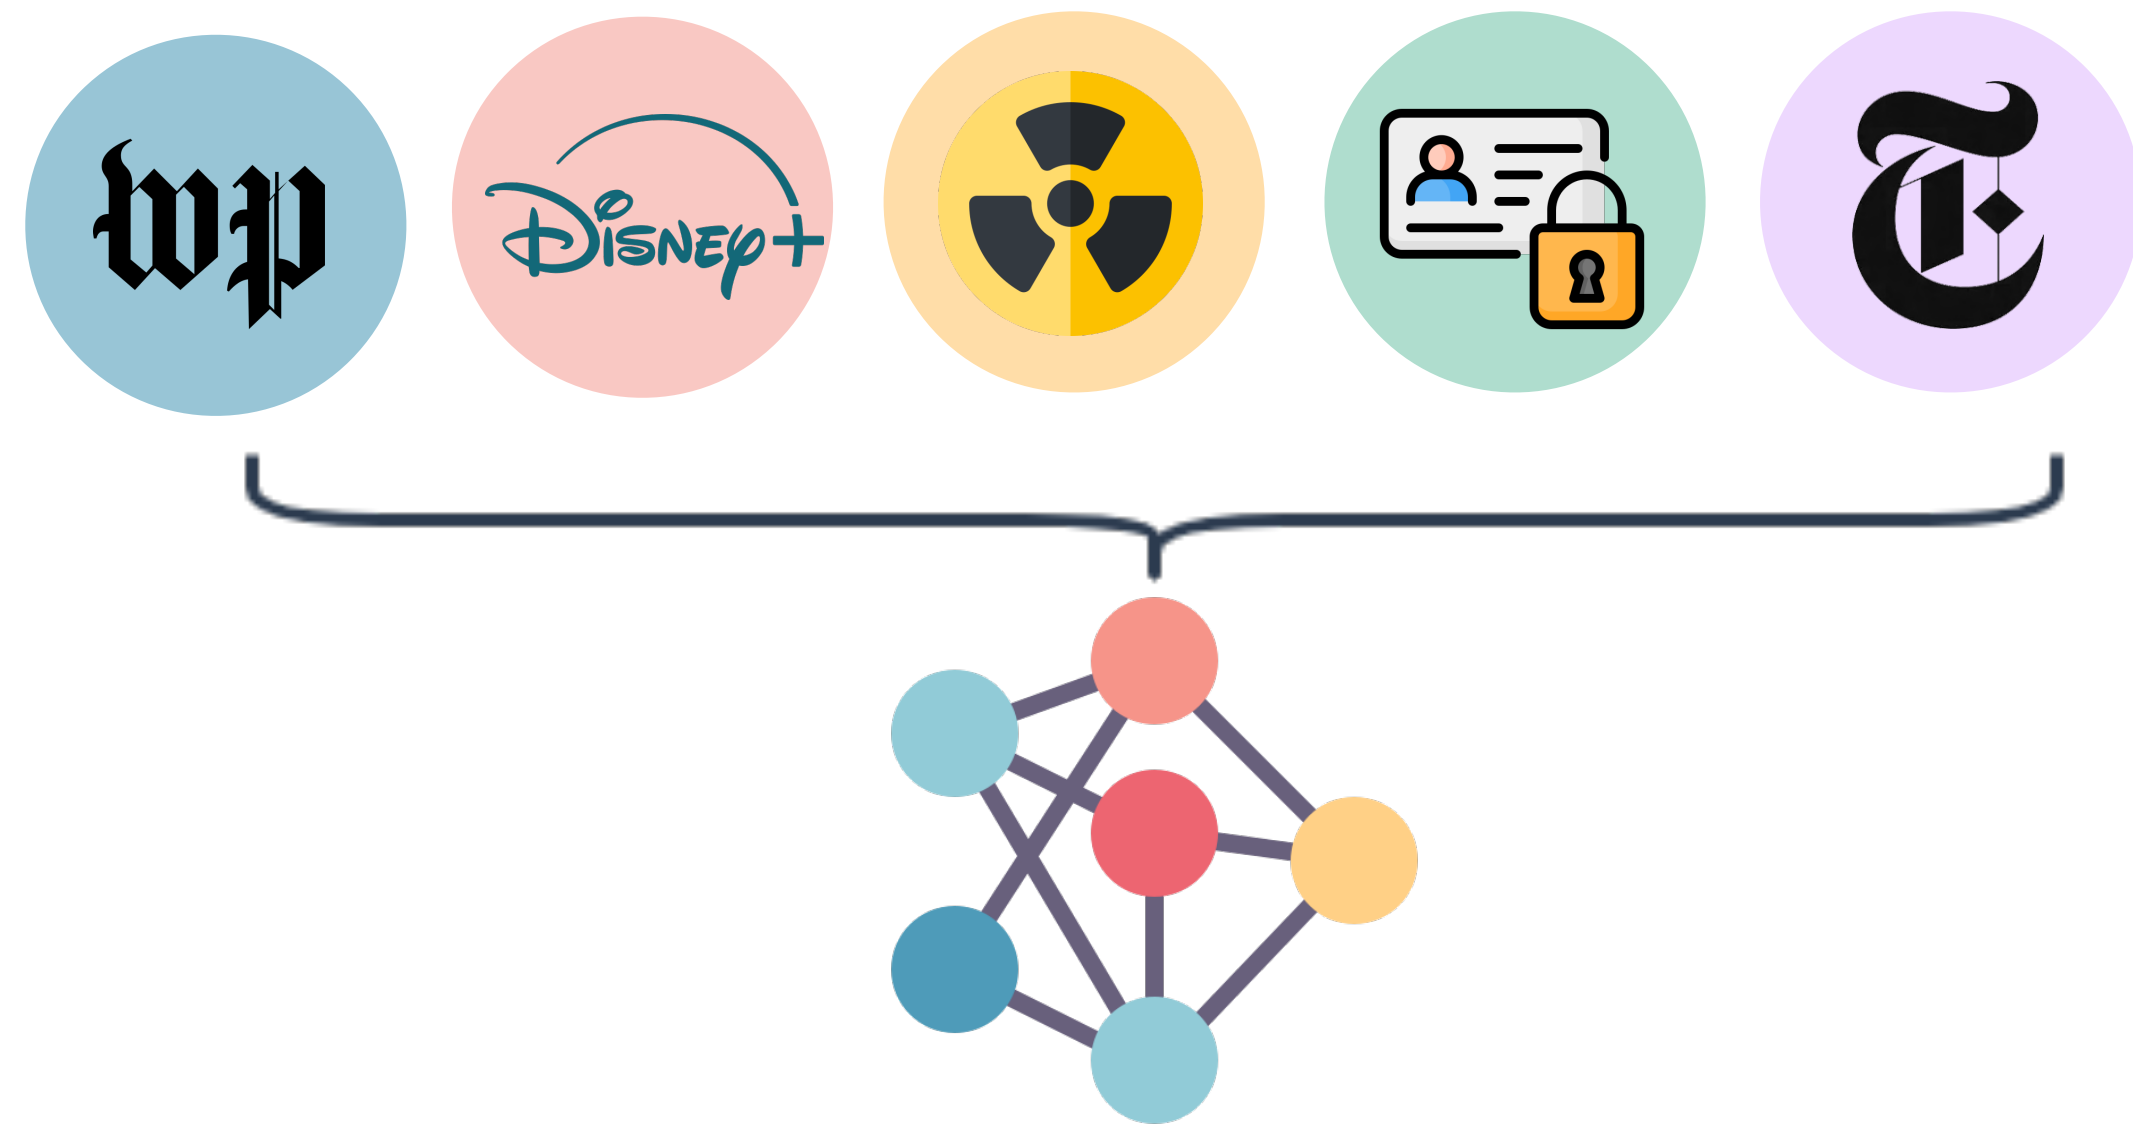
Decoding-time:  
Check & Resample

Training-based:  
Unlearning

**None of the current methods can balance utility & copyright risk mitigation**

How can we build *responsible* models?

# Proposal: Models with Data Provenance



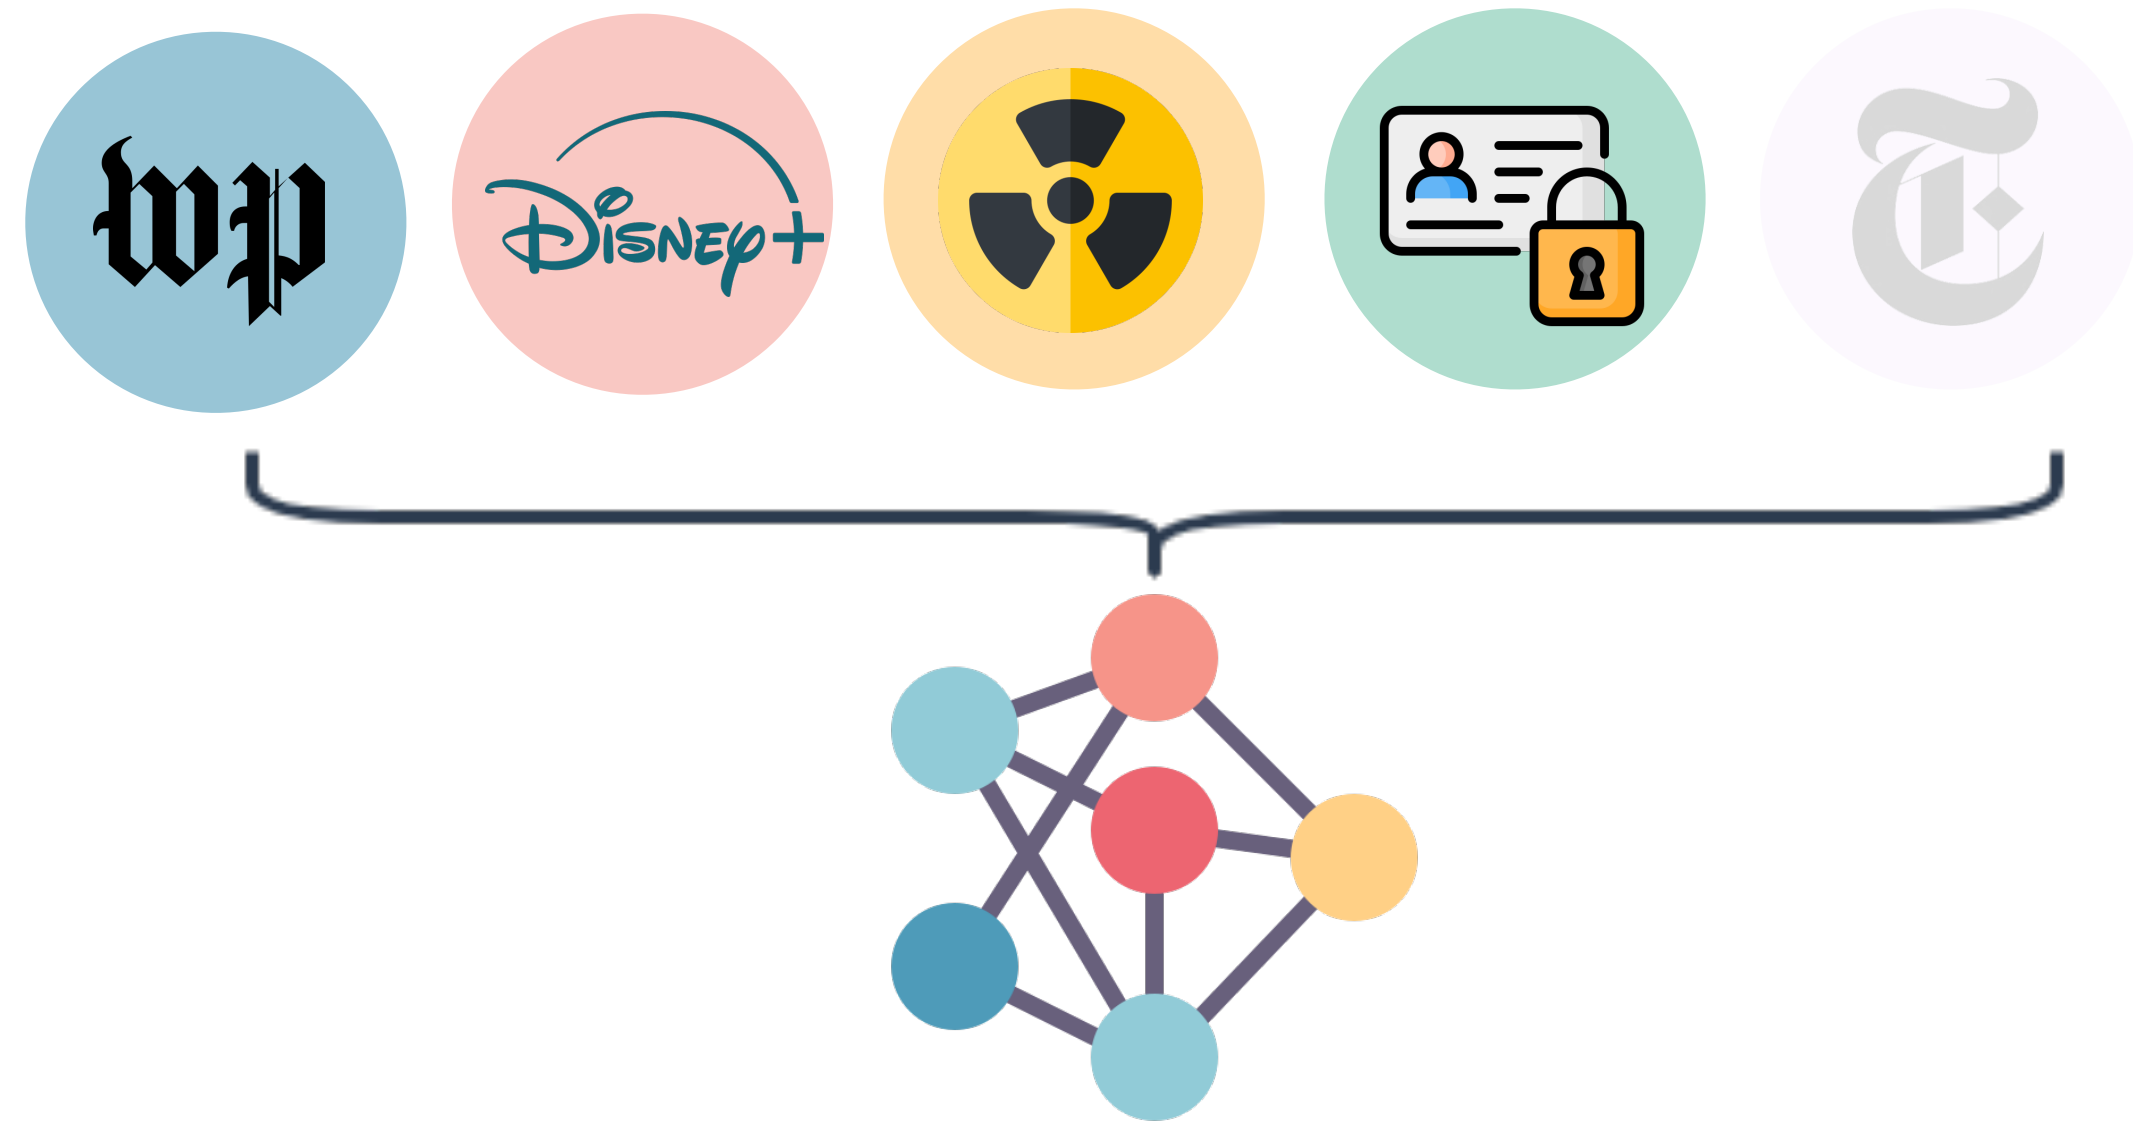
Models with **different components** trained on different **origins** of the data

These **origins (license, categories...)** are transparently tracked and documented.

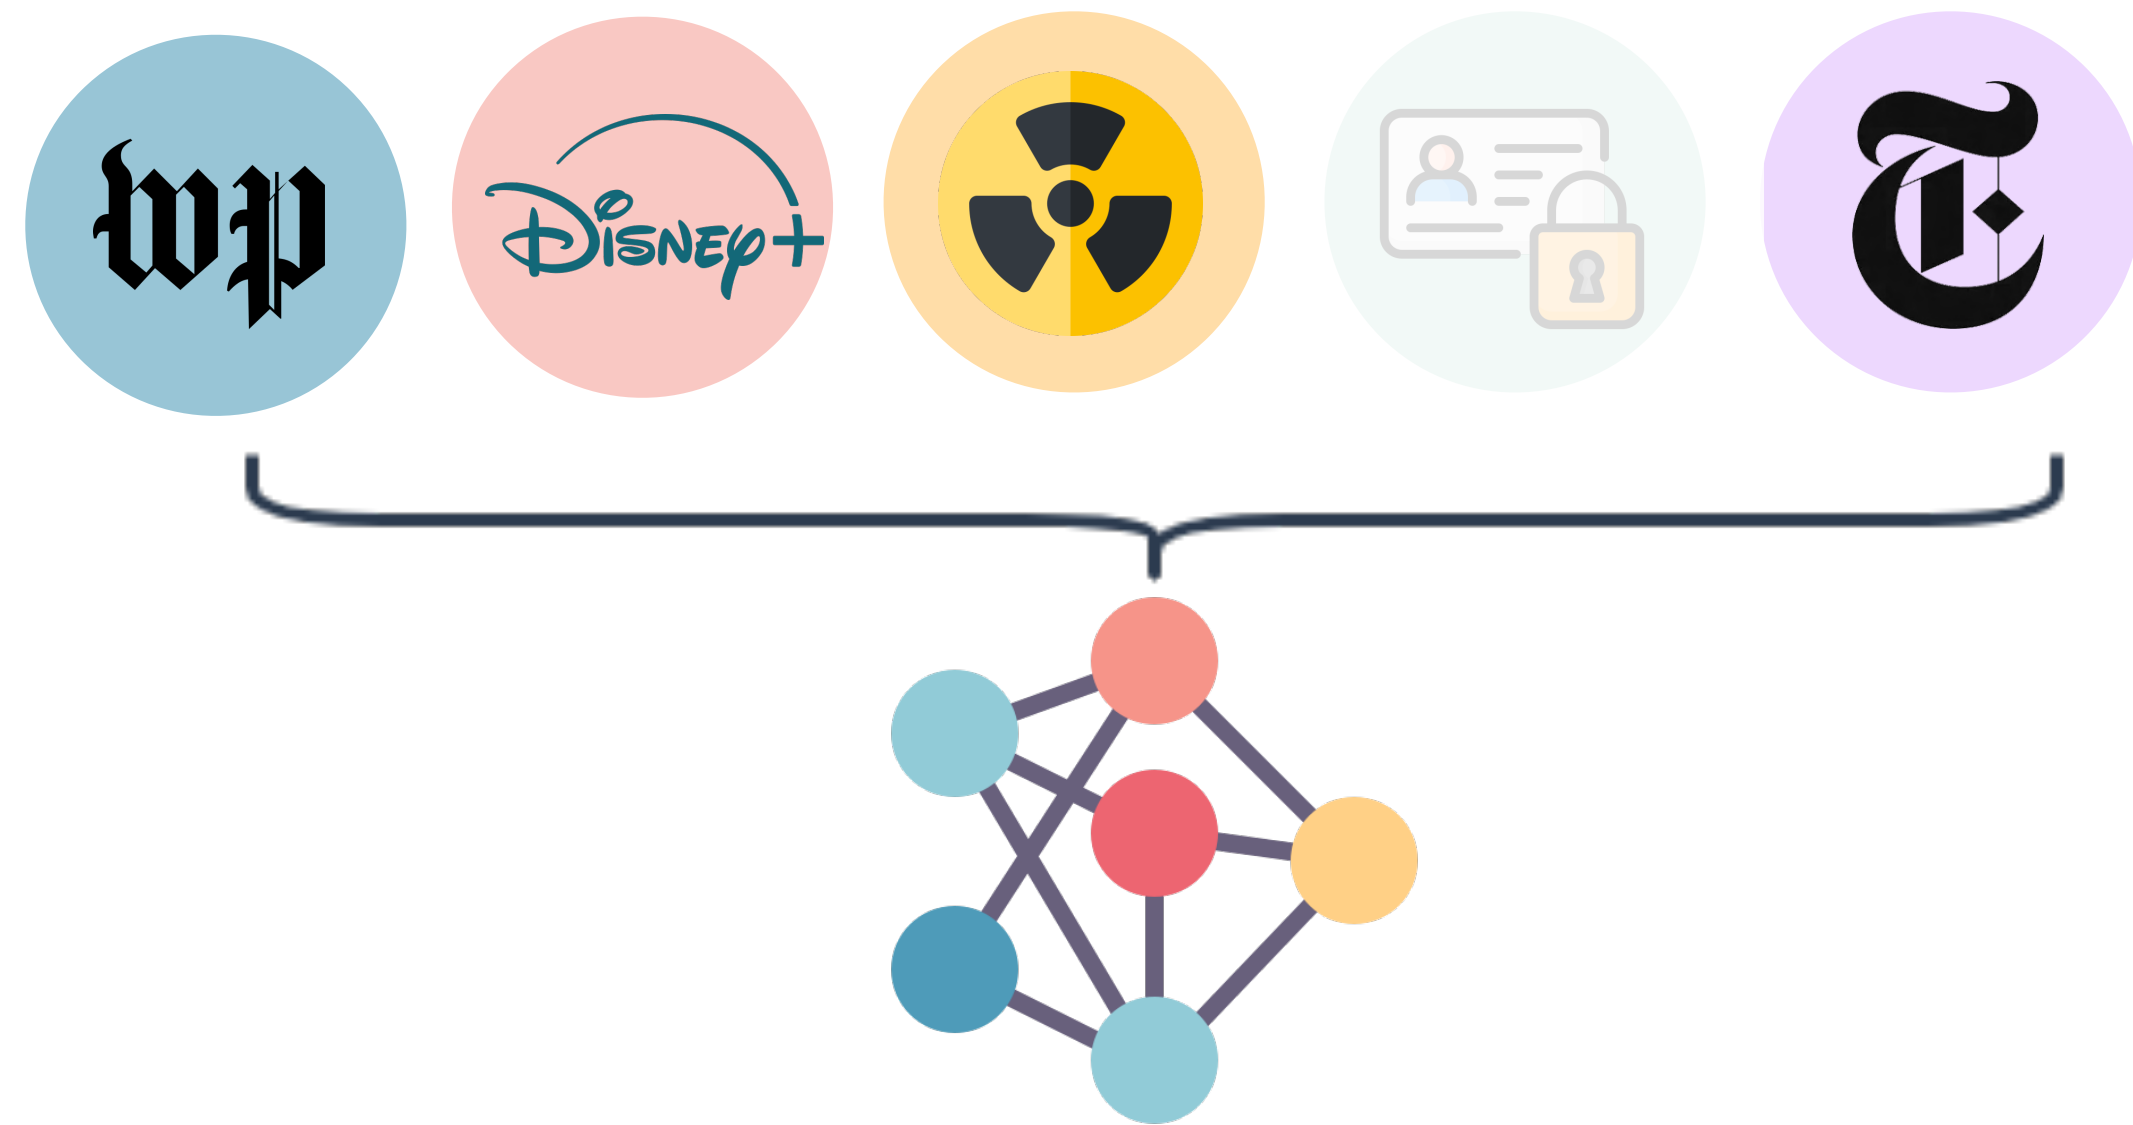
# Proposal: Models with Data Provenance



Copyright Takedowns



# Proposal: Models with Data Provenance

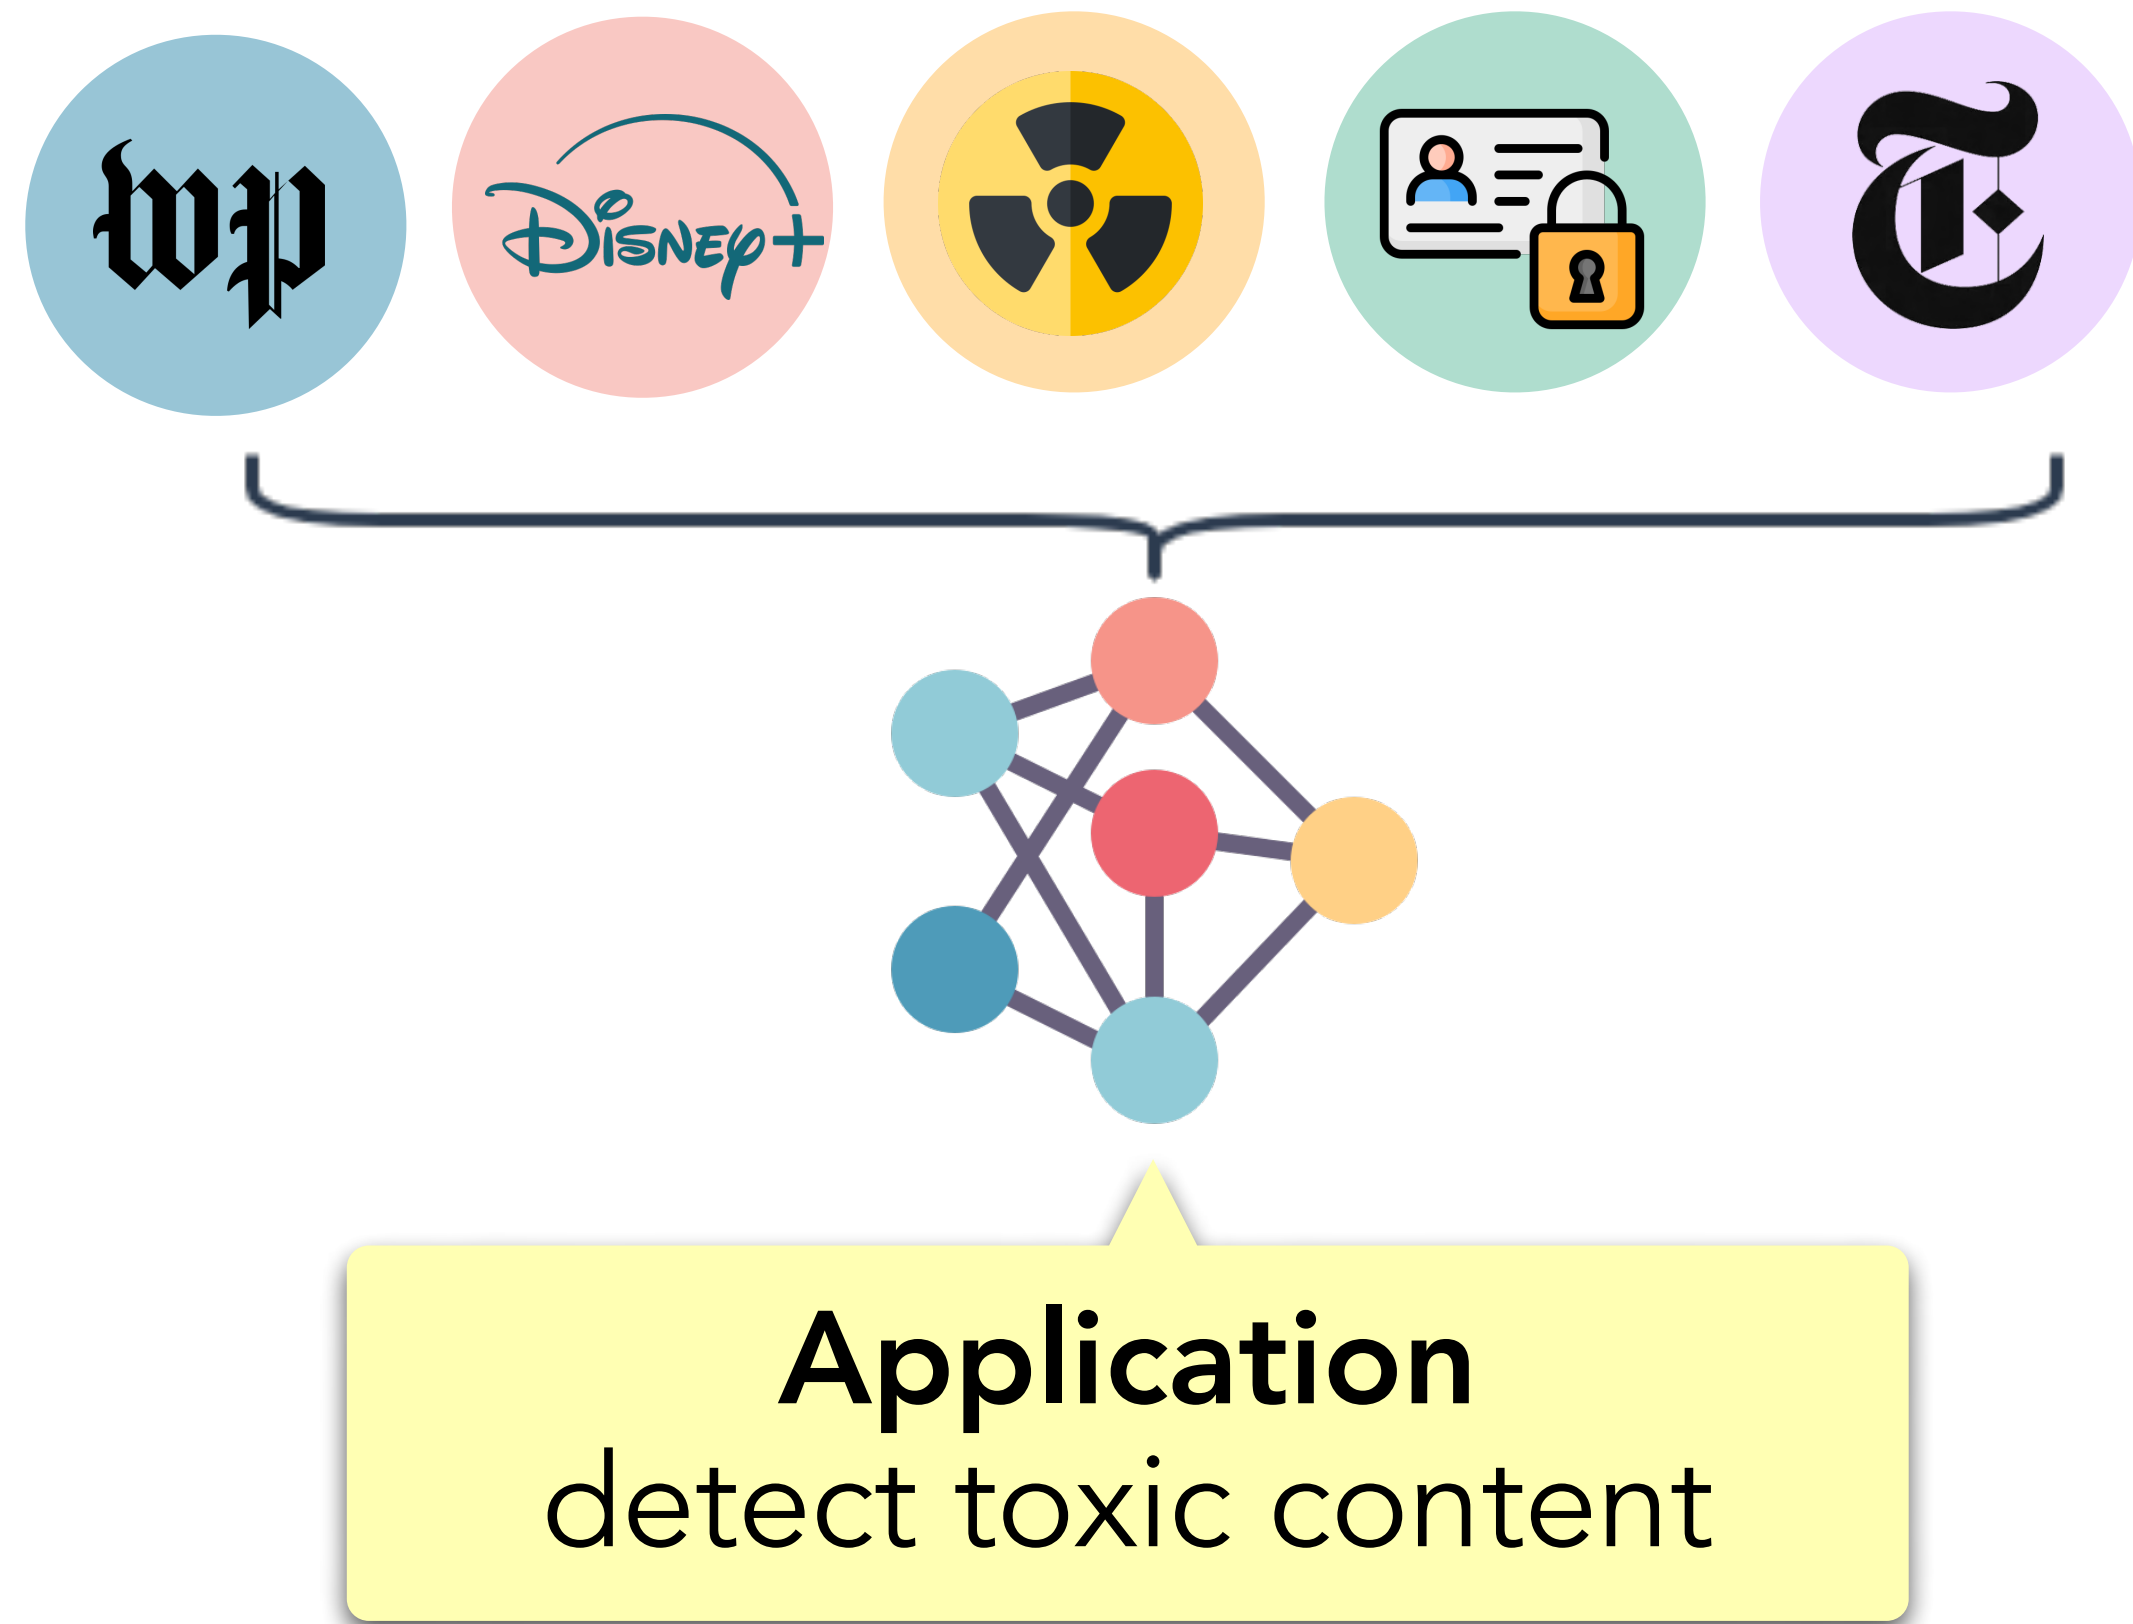


**Copyright Takedowns**



**Unlearning Private Data**

# Proposal: Models with Data Provenance

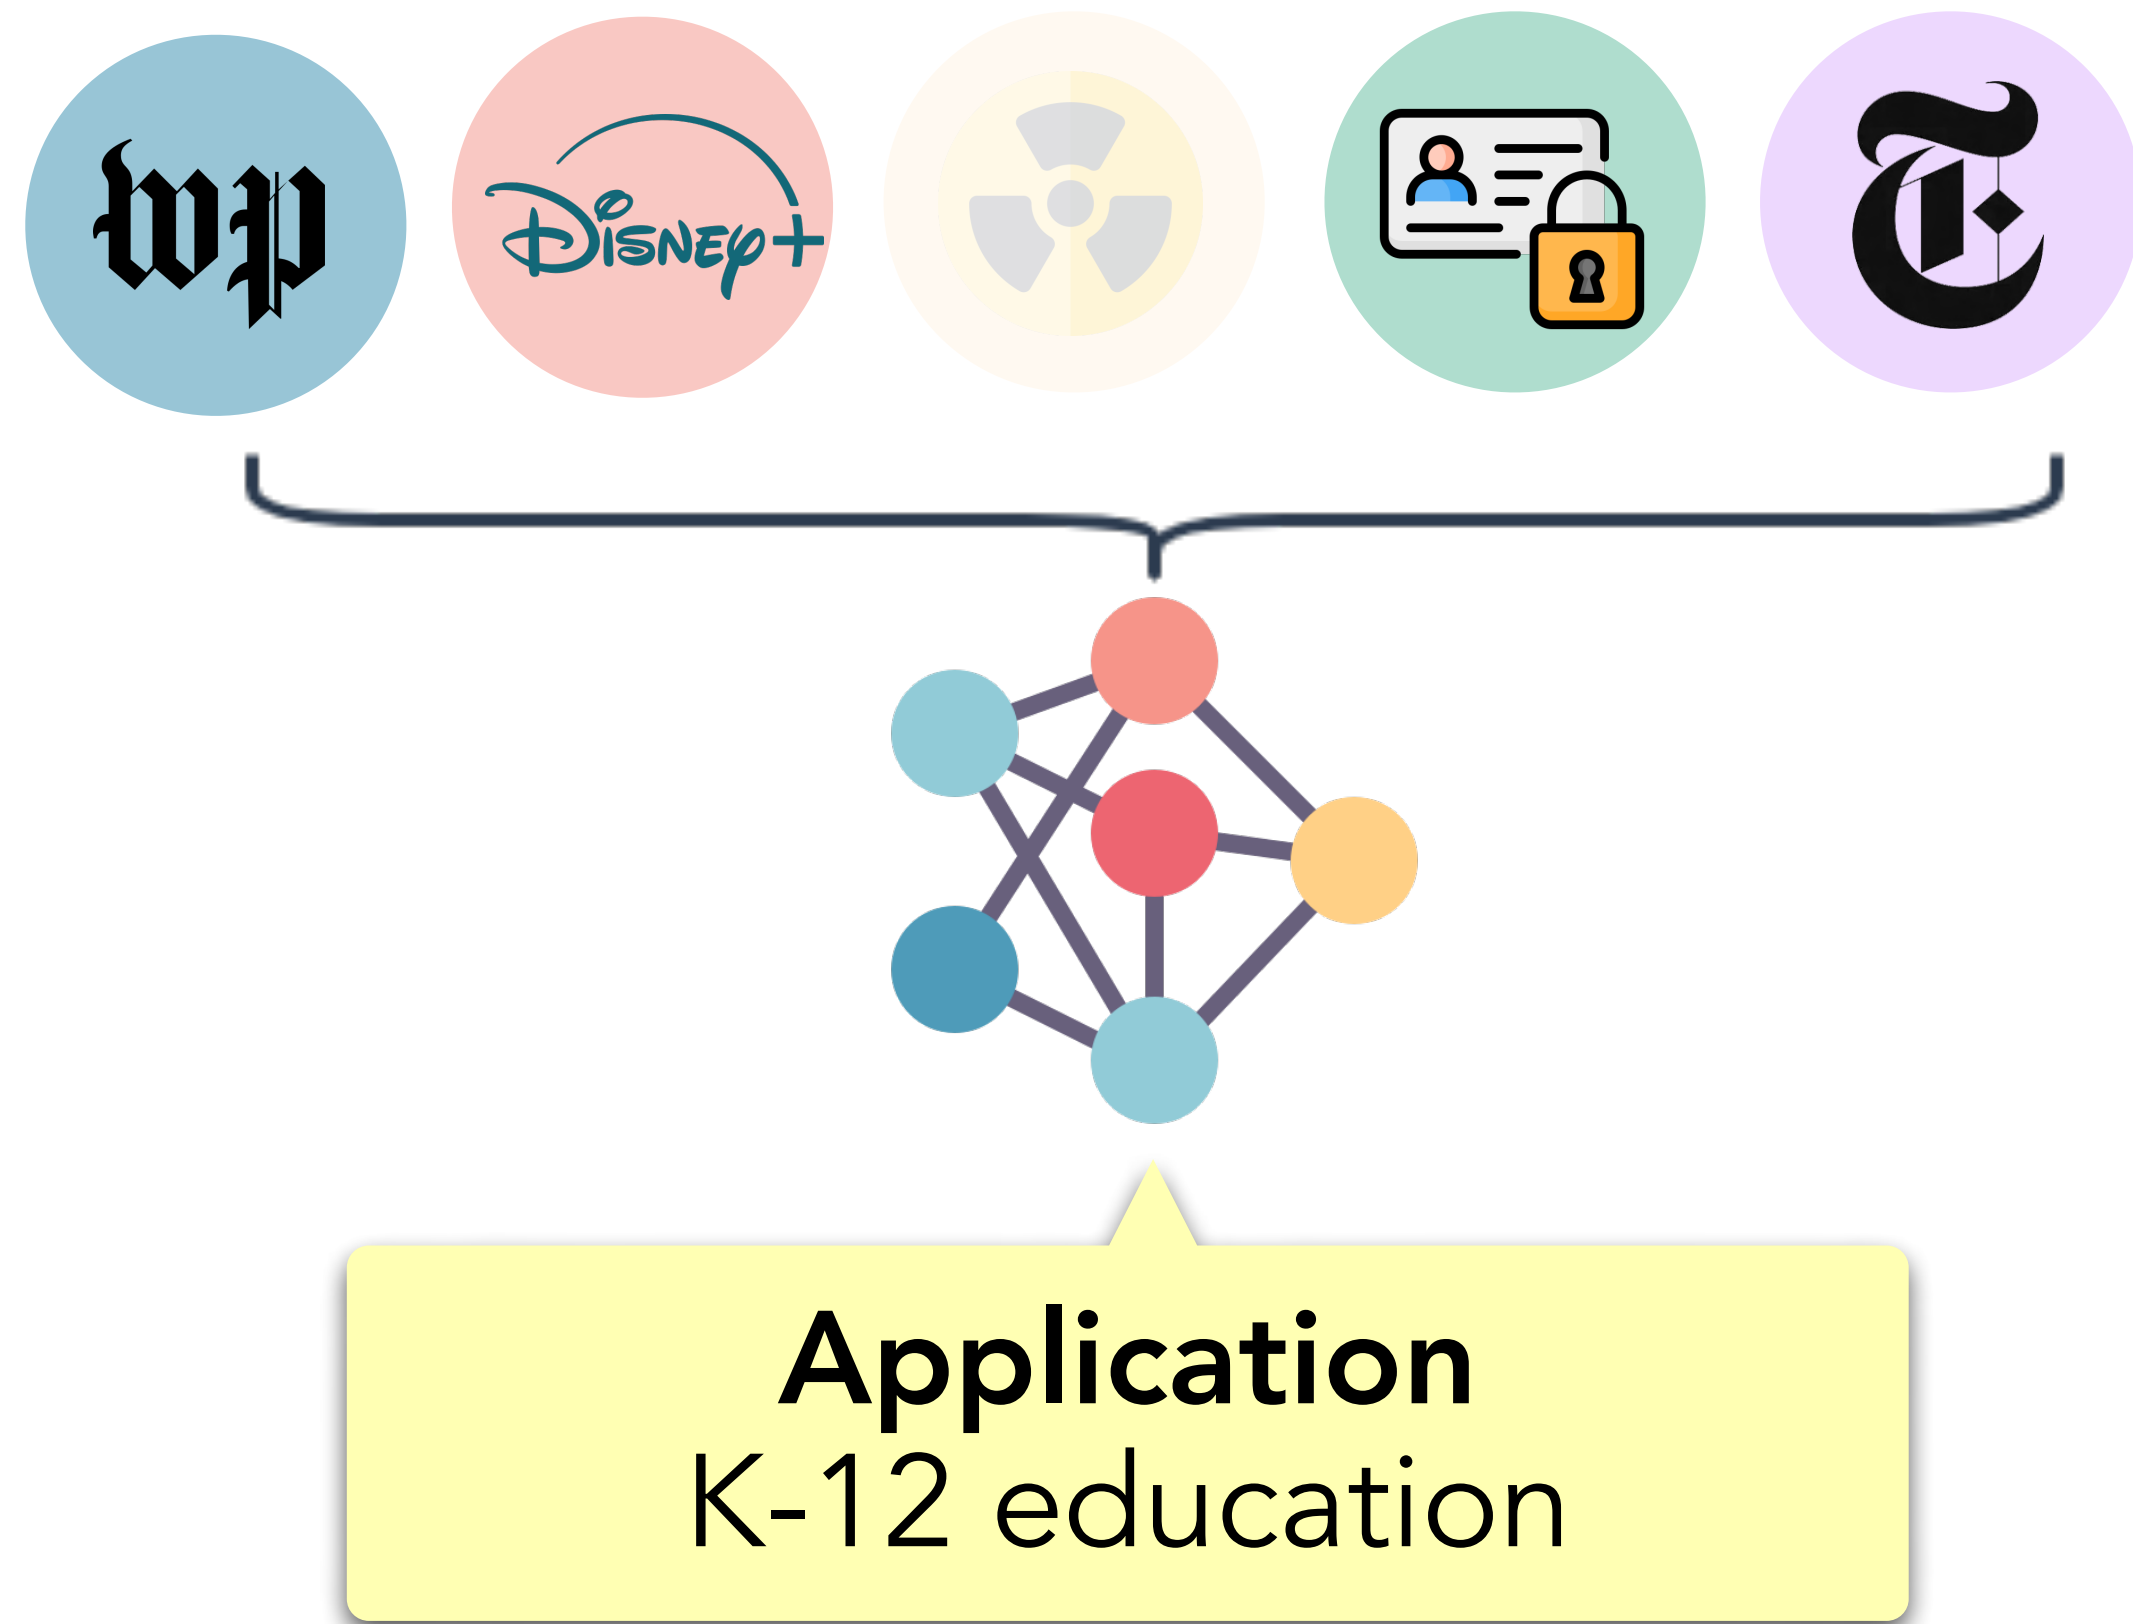


✓ Copyright Takedowns

✓ Unlearning Private Data

✓ Safe deployment

# Proposal: Models with Data Provenance



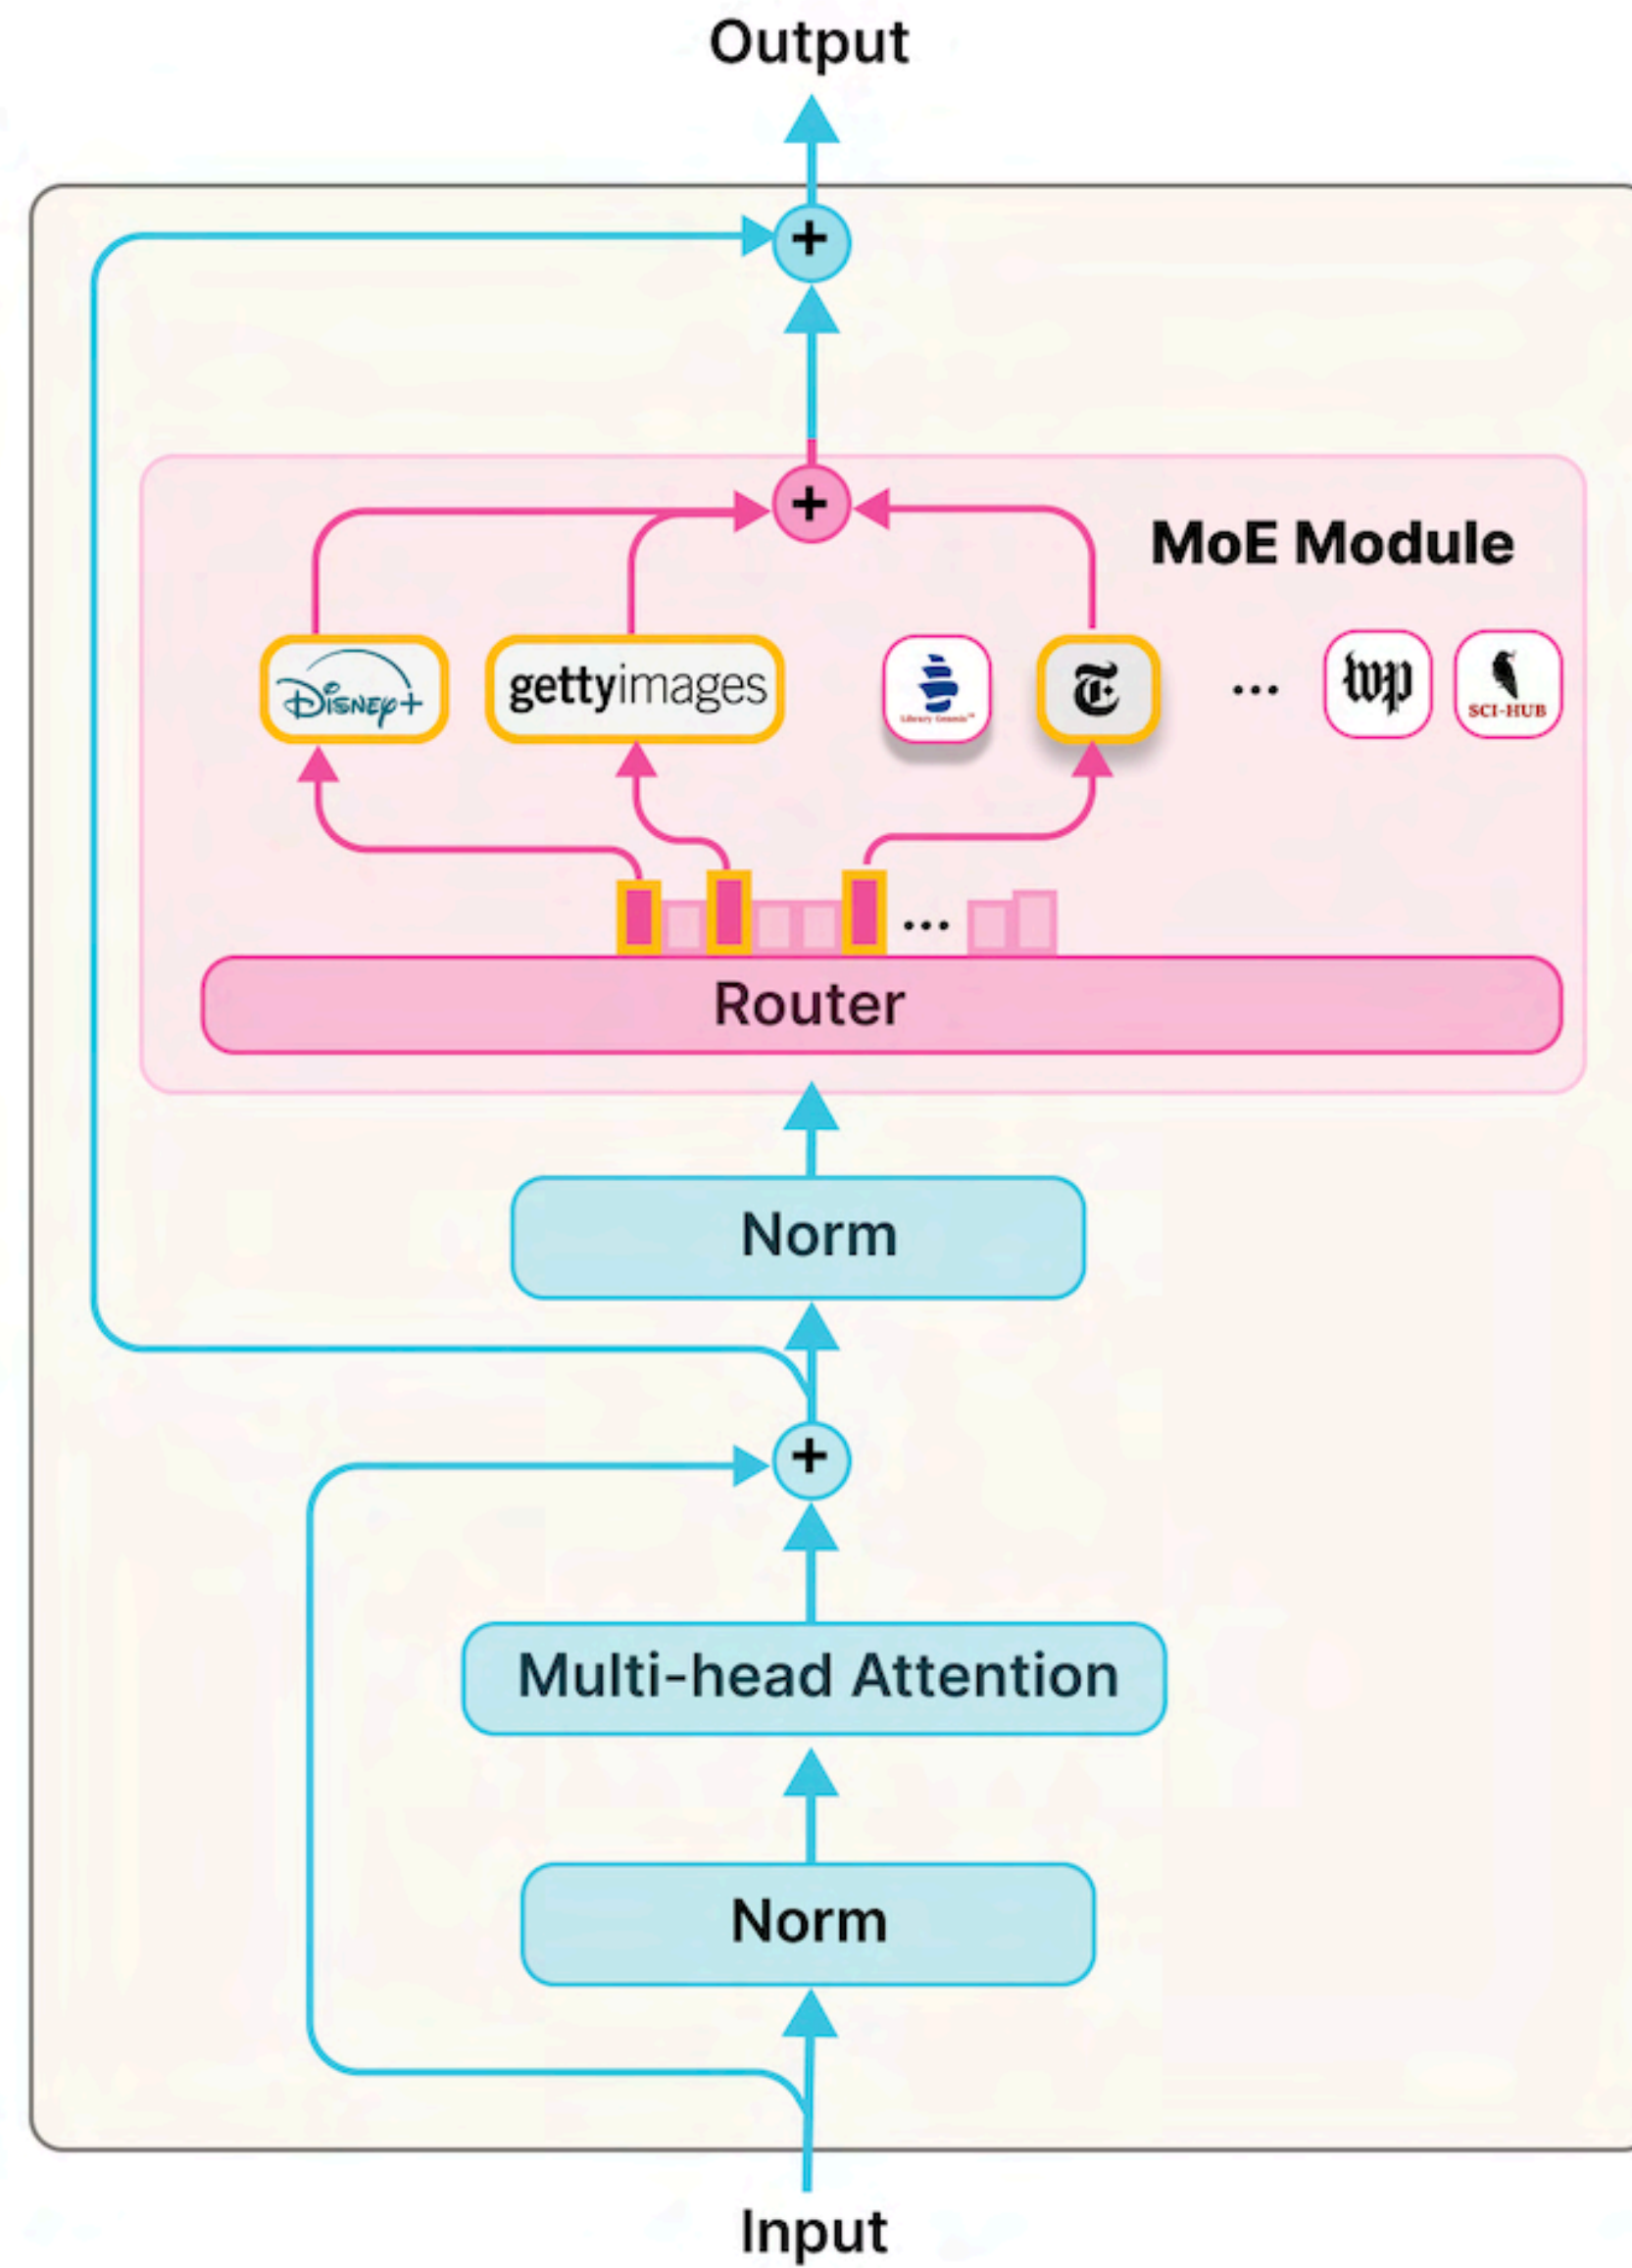
✓ Copyright Takedowns

✓ Unlearning Private Data

✓ Safe deployment

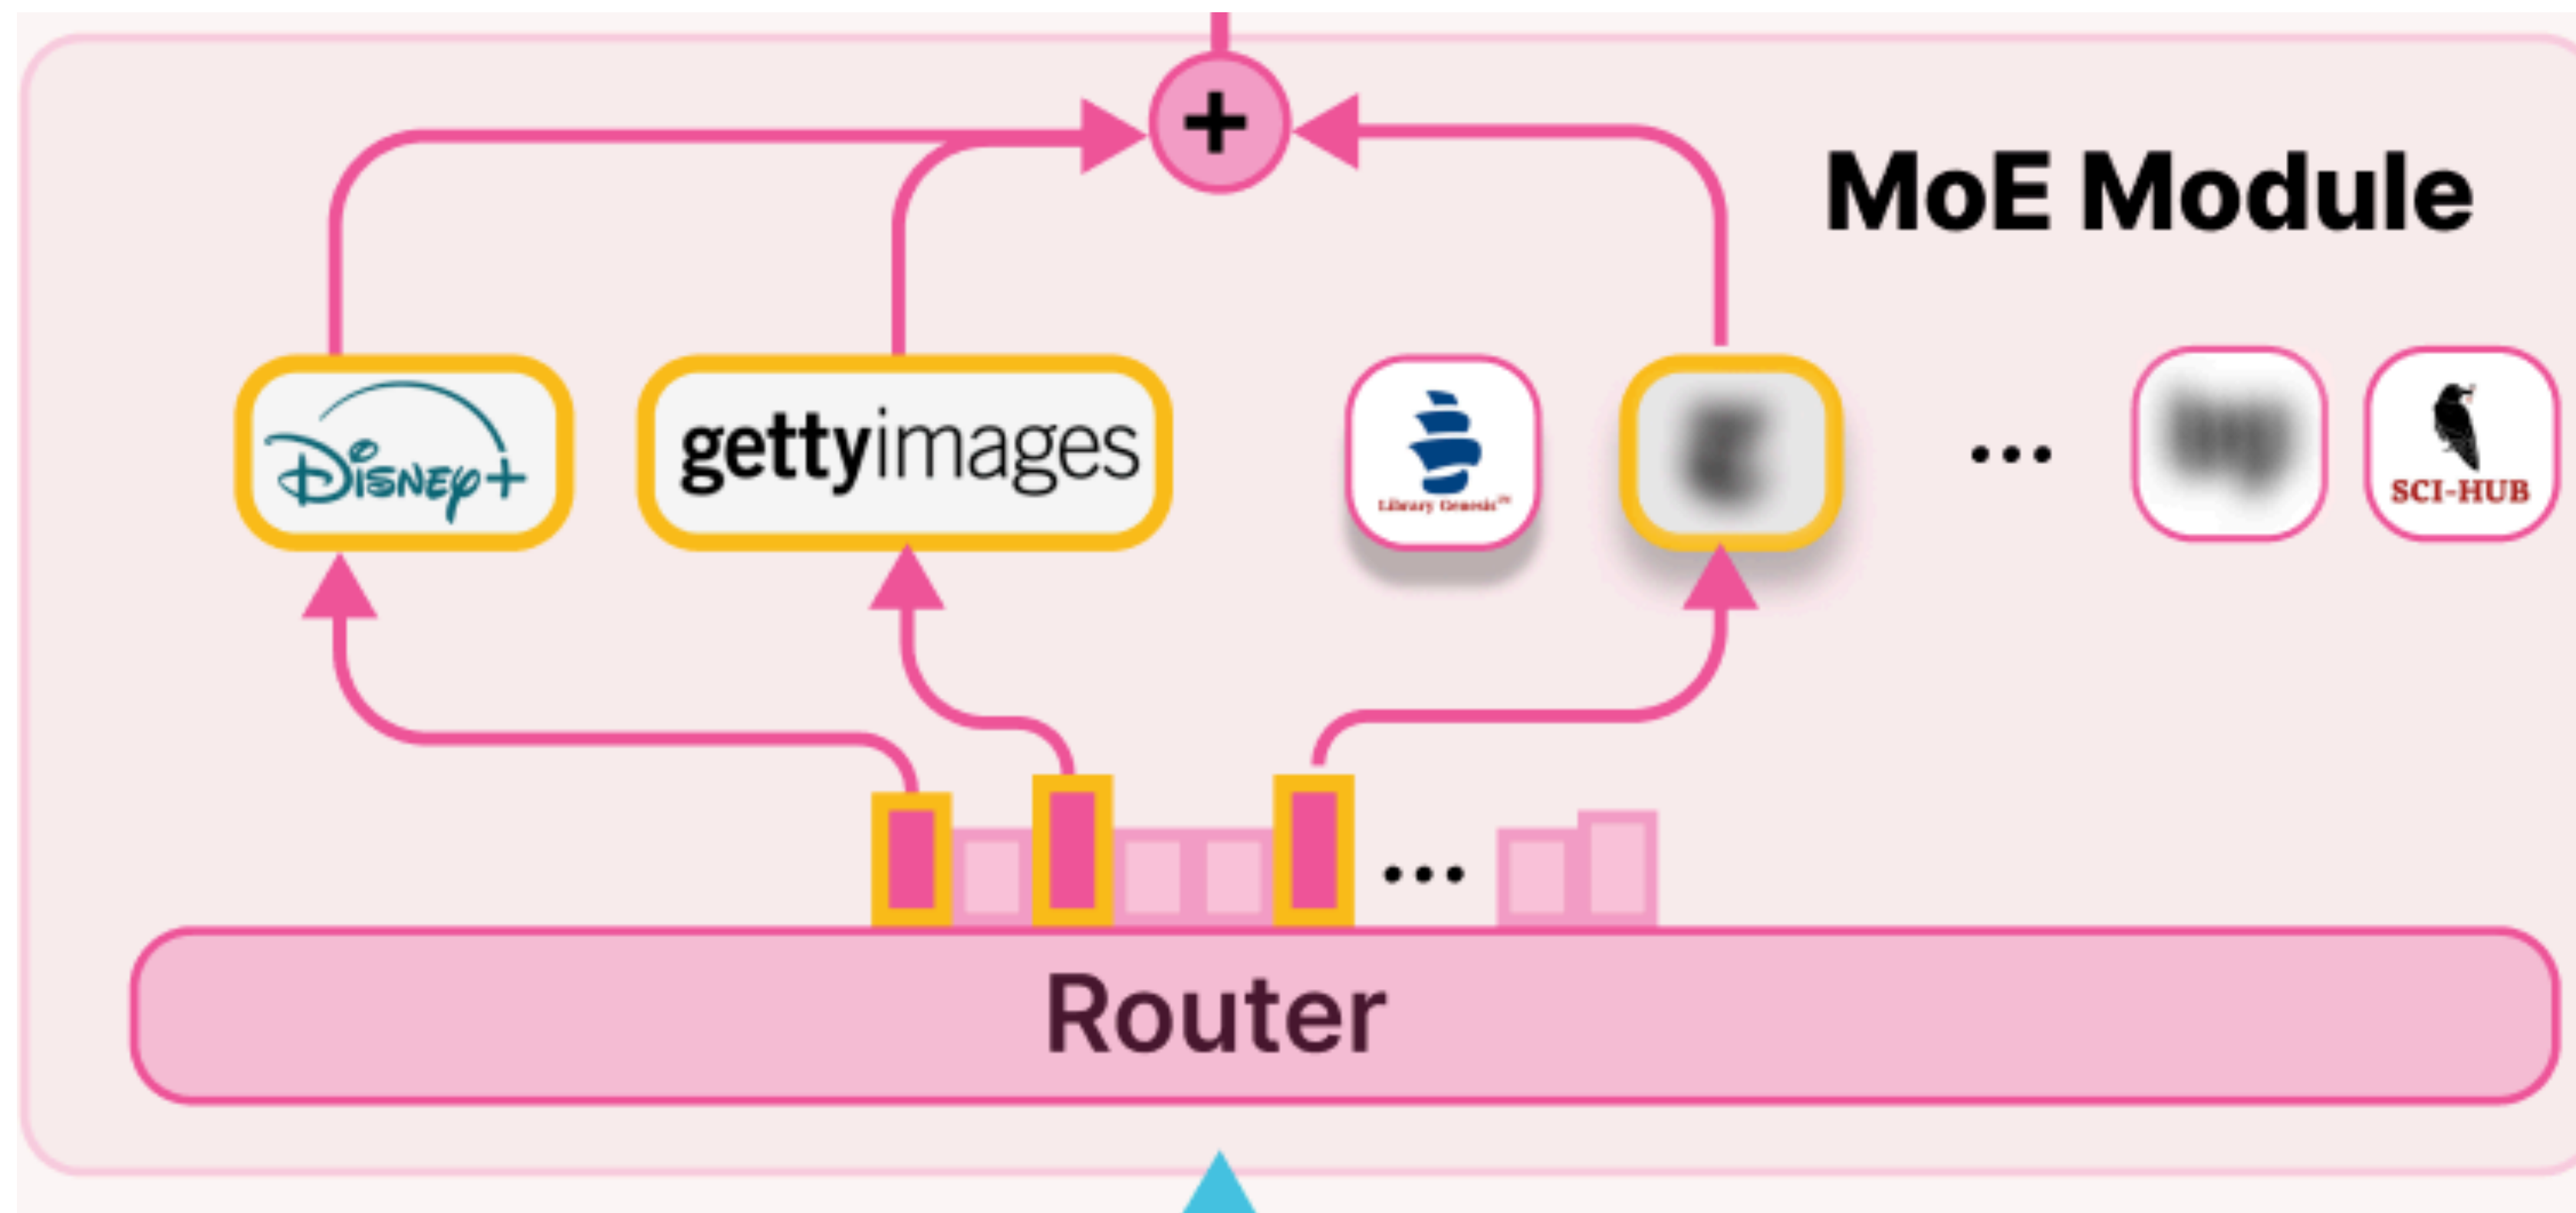


# How?



# How?

*Deactivation* of modules during inference based on the query



# Beyond Monolithic Language Models

*Augmented Models* 

*Data Modularity*



# *Modularity, not Monoliths*

# Thank You!

 swj0419@uw.edu

 @WeijiaShi2