

# Robust and efficient frontier pipelines for complex knowledge intensive tasks in the era of LLMs

---

CNI Seminar series, IISC

Venktesh Viswanathan

Postdoctoral Researcher, TU Delft

# Knowledge Intensive Language Tasks

**Efficient**

**Effective**



Conversational AI

Fact Checking Articles

Web Search

Knowledge-base  
Construction

Financial Auditing

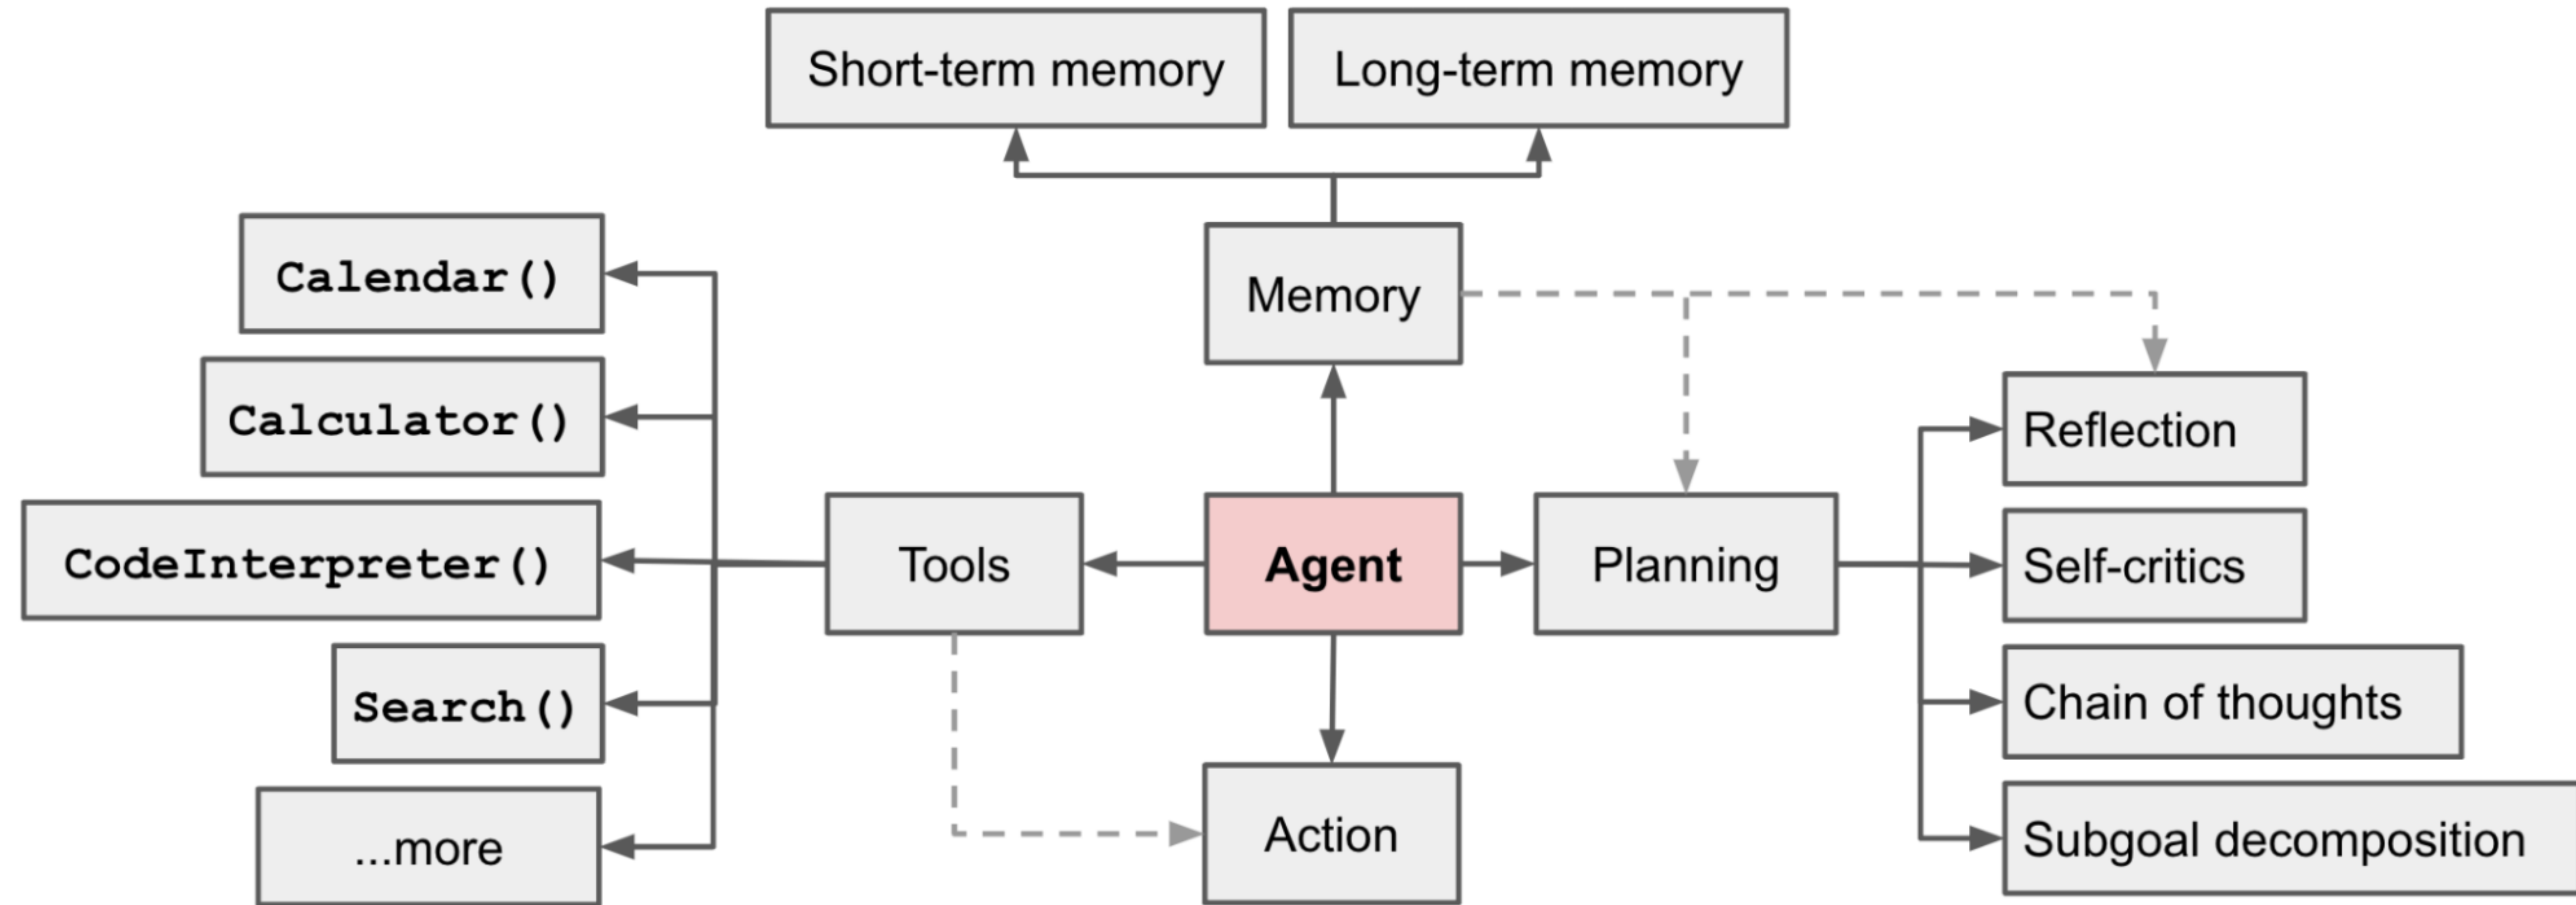
Live fact checking

Factual Article Generation

**Indexing and  
Precompute**

**Test time  
Reasoning**

# LLM Agents as general purpose solvers



# Hallucination in closed book setting



What causes Noonan syndrome?



Noonan syndrome is caused by a mutation in the **PTEN** gene. This gene is responsible for regulating cell growth and division, and when it is mutated, it can lead to the development of Noonan syndrome.

Now imagine a LLM citing fake cases when a resident is preparing his report  
Or a lawyer preparing his arguments

There's no provenance even if answer is correct.

# Hallucinations - catastrophic effects

TECH · LAW

## Humiliated lawyers fined \$5,000 for submitting ChatGPT hallucinations in court: 'I heard about this new site, which I falsely assumed was, like, a super search engine'

BY RACHEL SHIN  
JUNE 23, 2023 AT 9:41 AM PDT



Lawyers who filed legal documents with false citations generated by ChatGPT have been fined \$5,000 each. (Erik McGregor—LightRocket/Getty Images)

MIT  
Technology  
Review

Featured Topics Newsletters Events Podcasts

SIGN IN

ARTIFICIAL INTELLIGENCE

## Why Meta's latest large language model survived only three days online

Galactica was supposed to help scientists. Instead, it mindlessly spat out biased and incorrect nonsense.

By Will Douglas Heaven

November 18, 2022

## Air Canada must honor re-invented by airline's chatbot

Air Canada appears to have quietly killed its costly chatbot support.

ASHLEY BELANGER - 2/16/2024, 12:12 PM

How to mitigate hallucination and establish provenance?

# Ask LLM to explain itself ?

## Unfaithful Reasoning

**Input**

Q: John plans to sell all his toys and use the money to buy video games. He has 13 lego sets and he sells them for \$15 each. He ends up buying 8 videogames for \$20 each and has \$5 left. How many lego sets does he still have?

**CoT output**

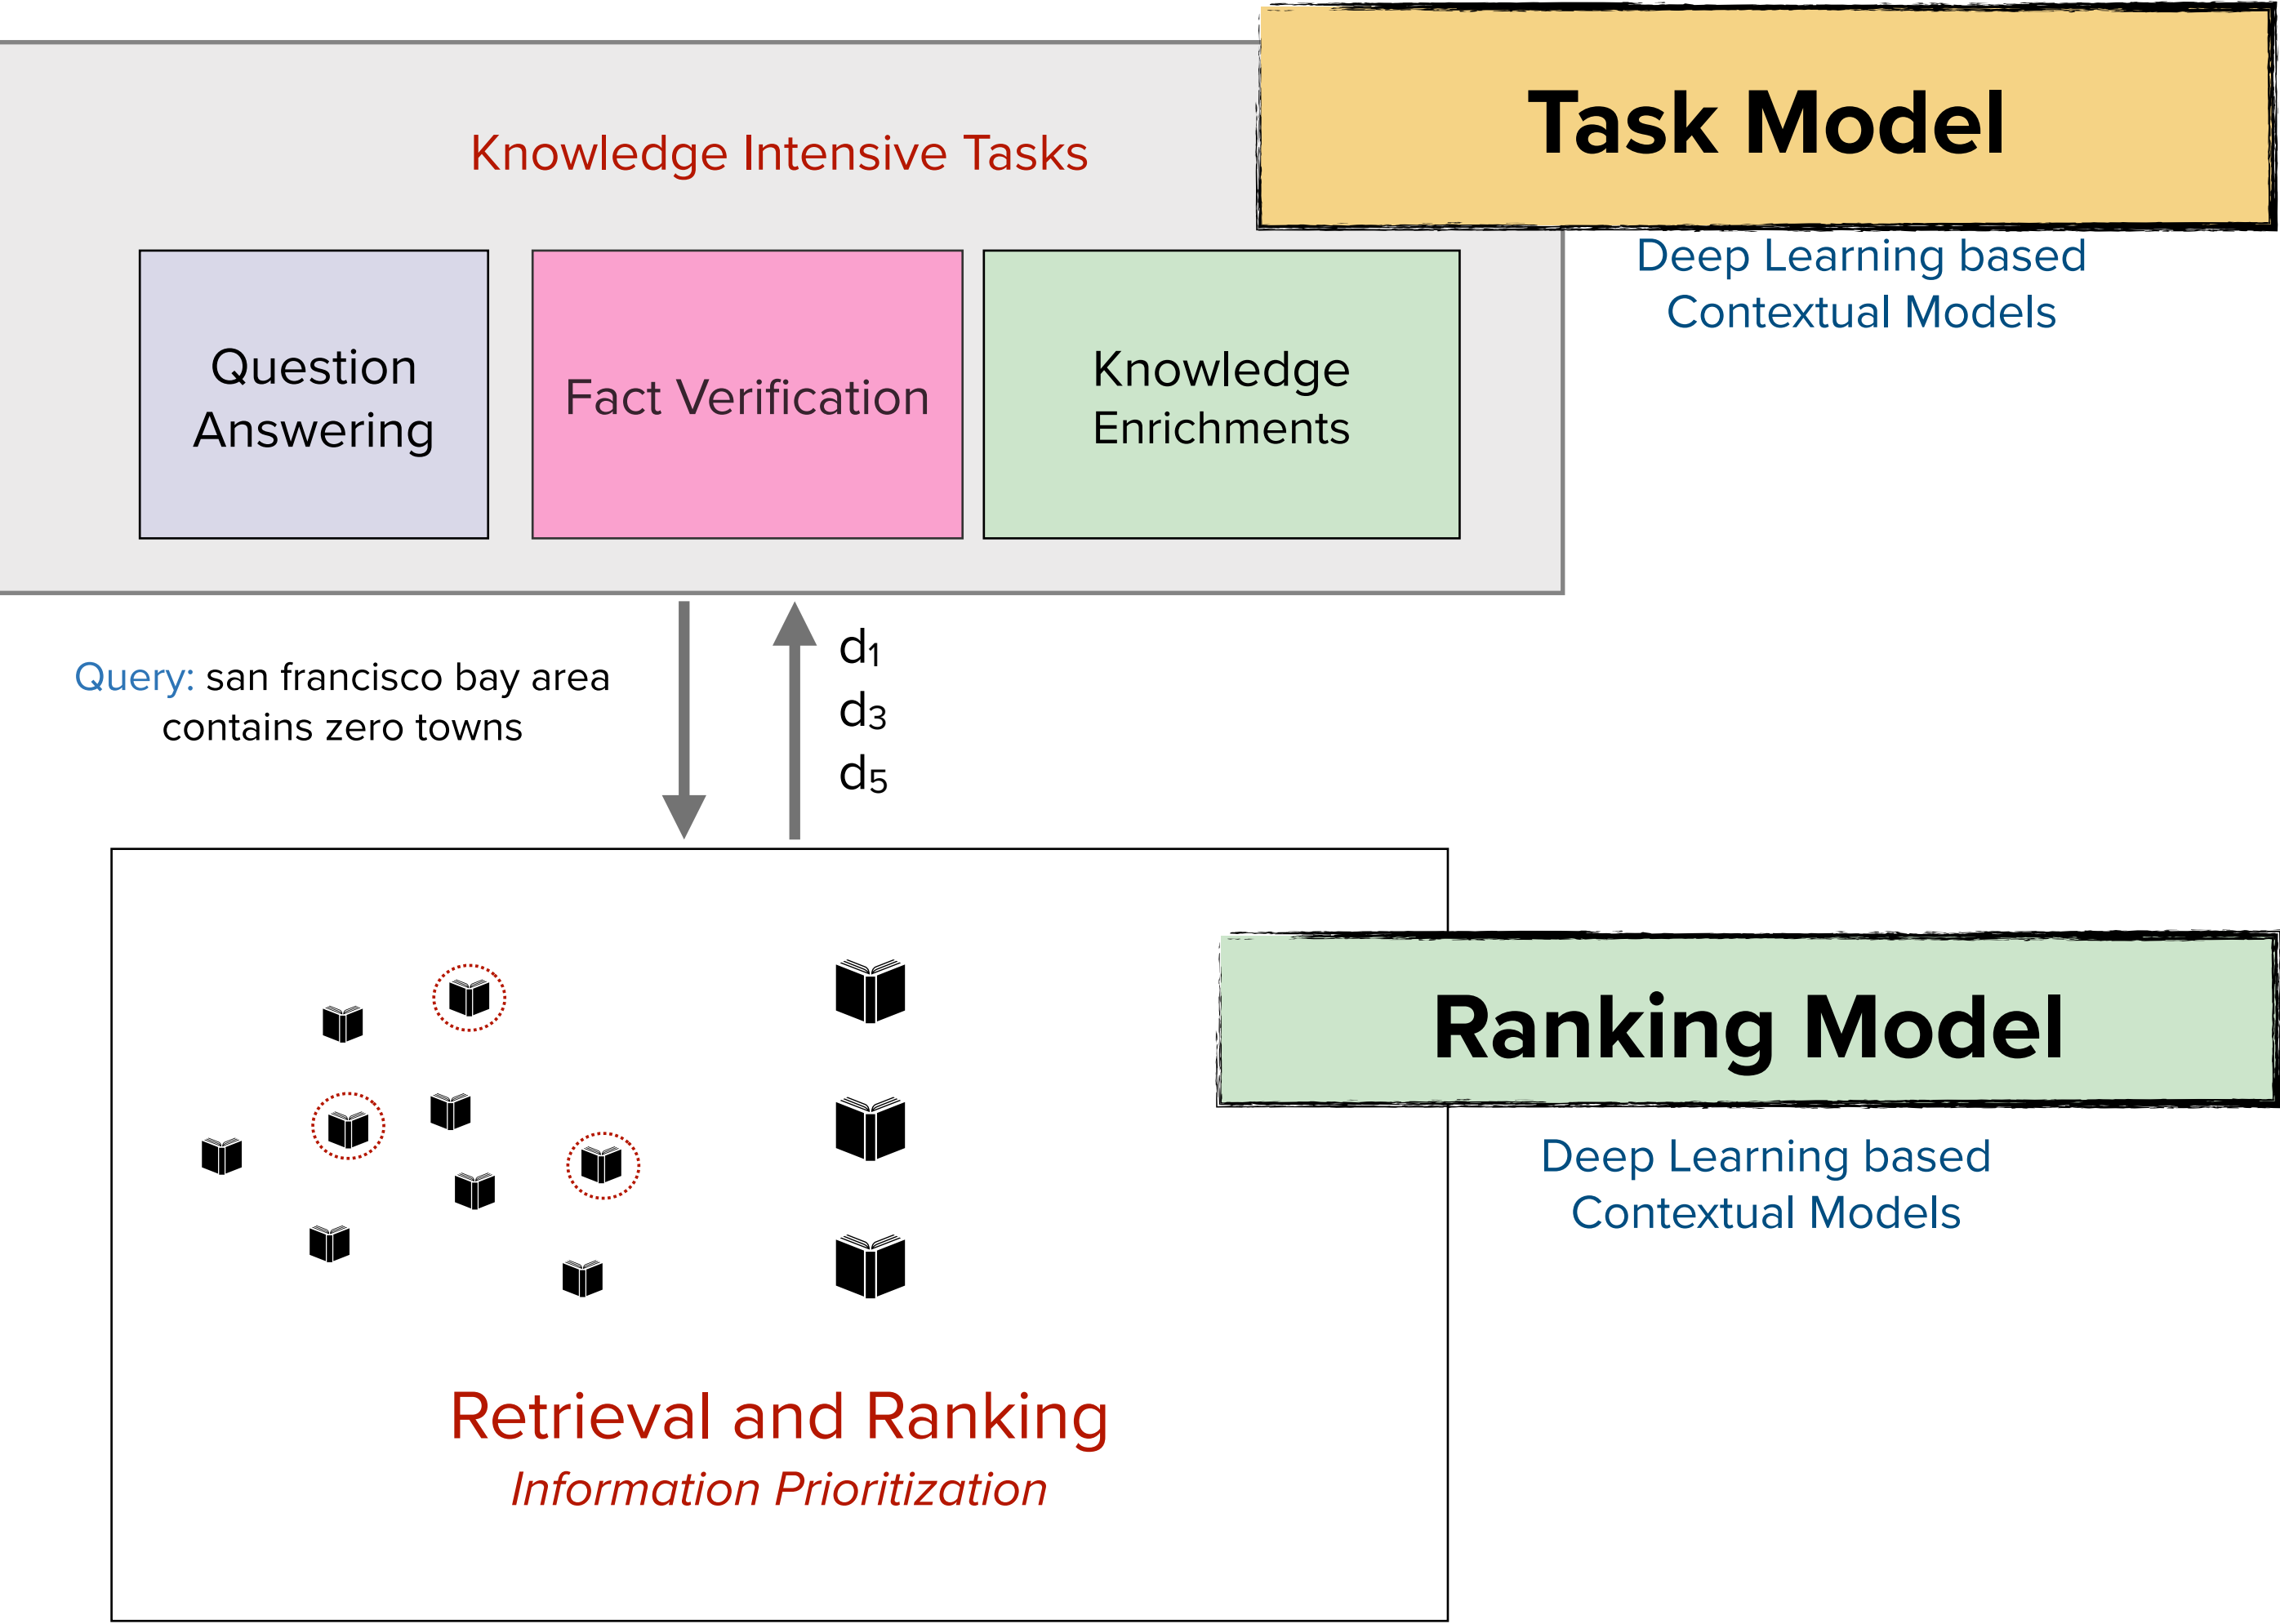
A: He started with 13 lego sets. He sold them for \$15 each. So he must have gotten  $13 * \$15 = \$195$  from selling them. He used this money to buy 8 video games for \$20 each. This means he spent  $8 * \$20 = \$160$  on the video games. He had \$5 left, so in total he must have had  $\$195 + \$5 = \$200$ . The answer is 0.

## Knowledge Gaps

**Question:** What is the mouth of the river which serves as the mouth of the Bumping River?

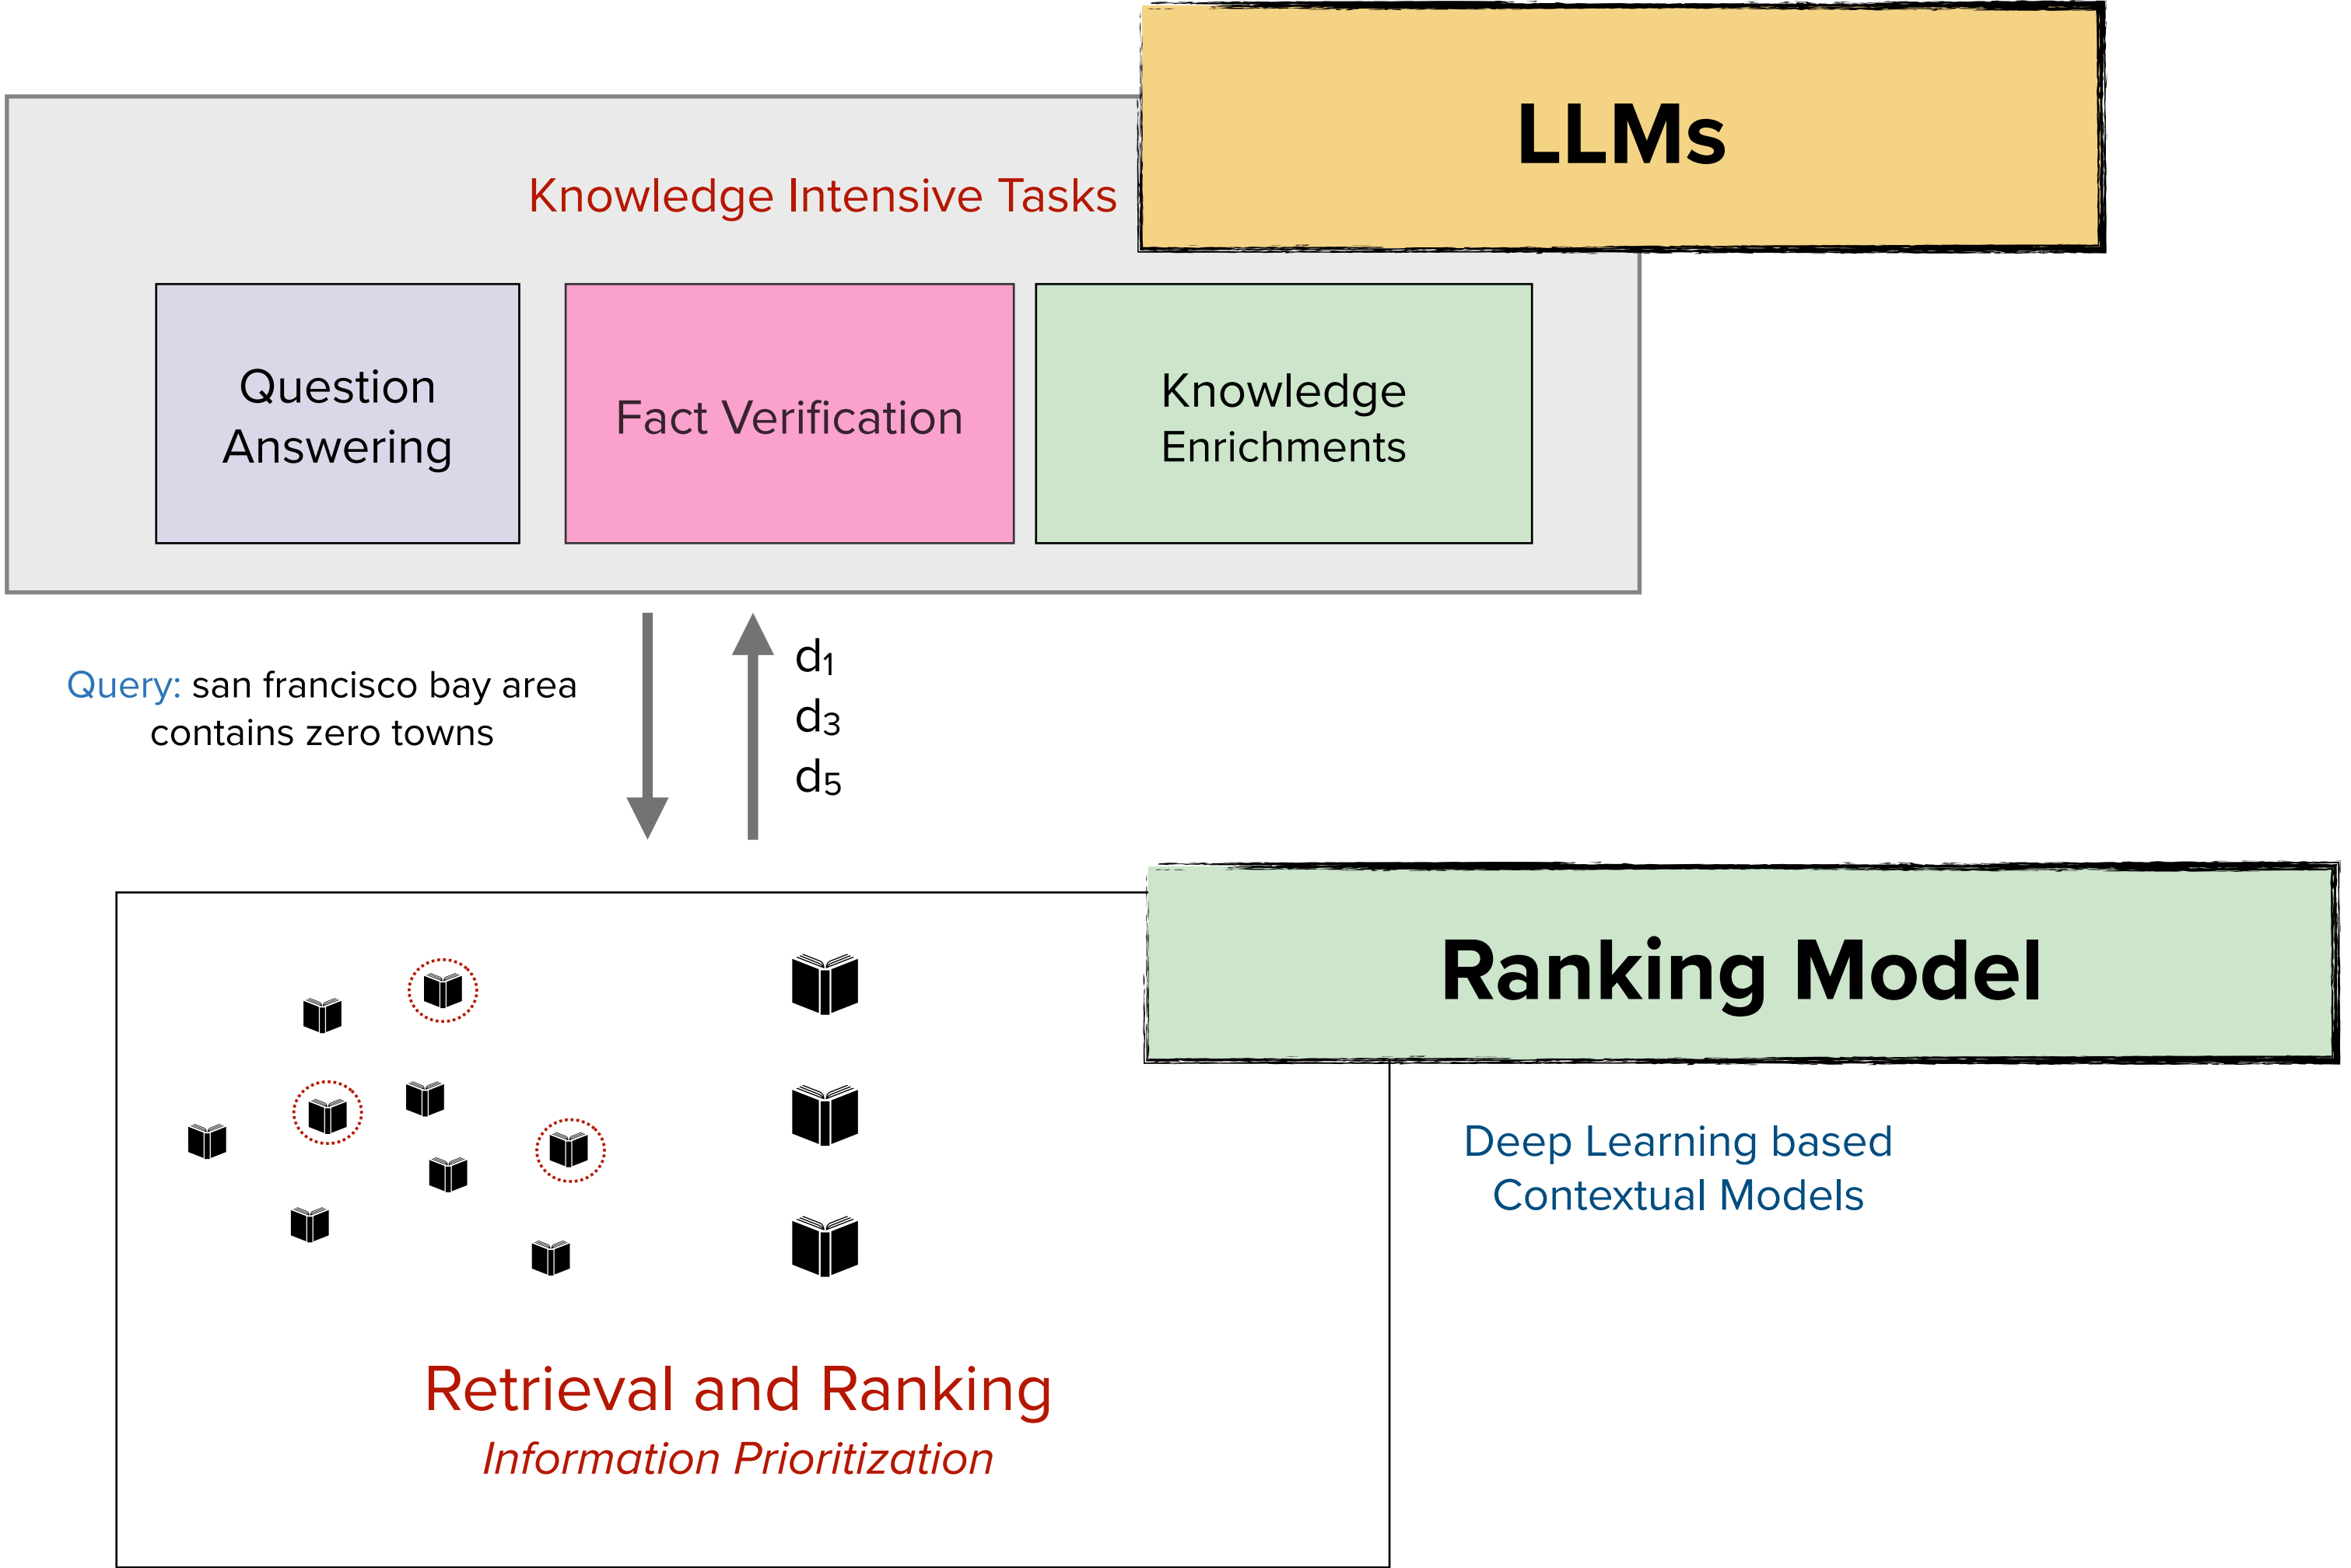
**FEW-SHOT-COT.** : [Answer]: **There is no river named Bumping River.**

# RAG to the rescue

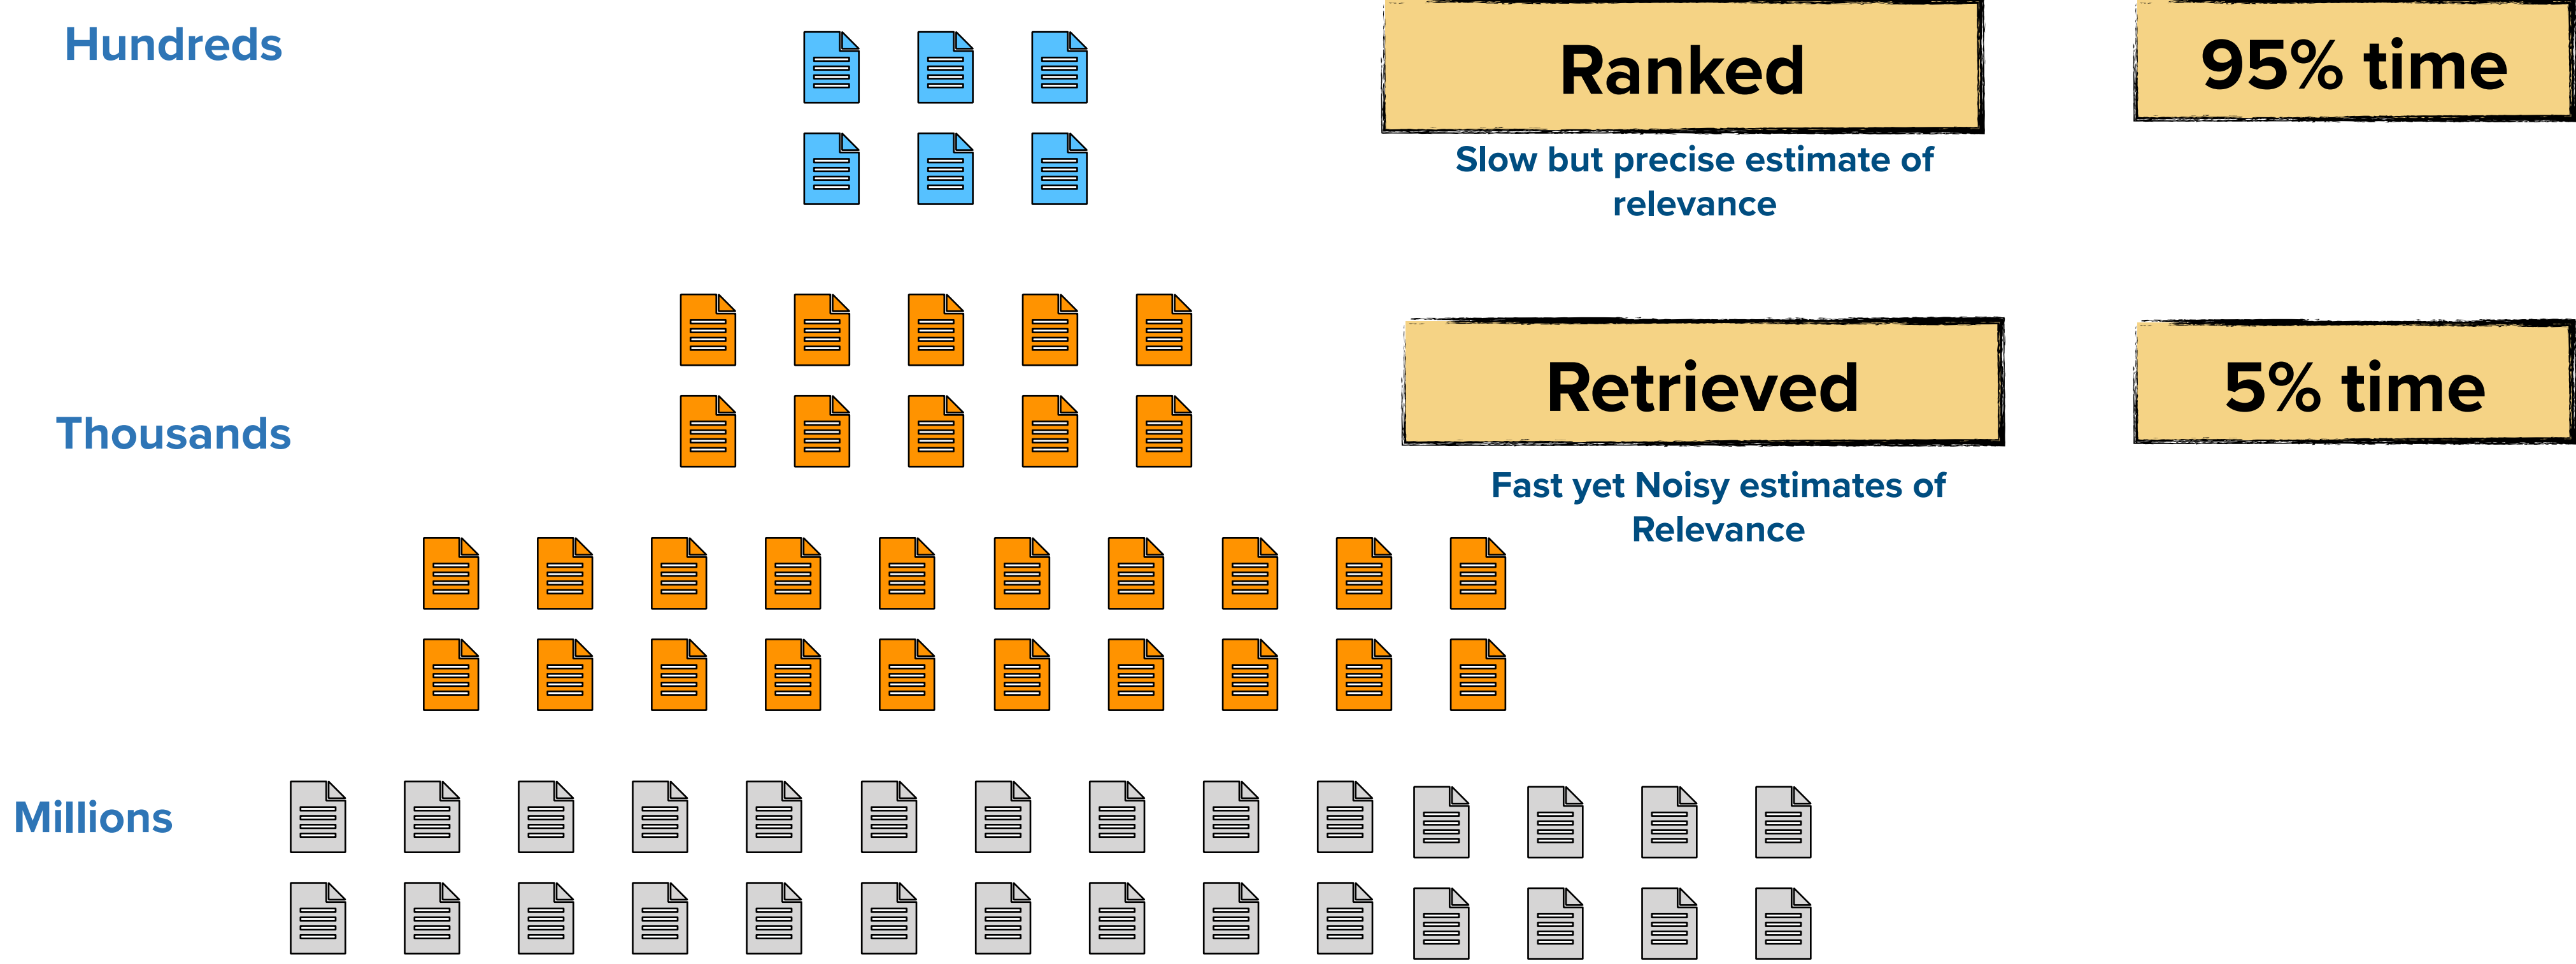




# RAG to the Rescue




# Telescoping view of retrieve-rerank pipelines



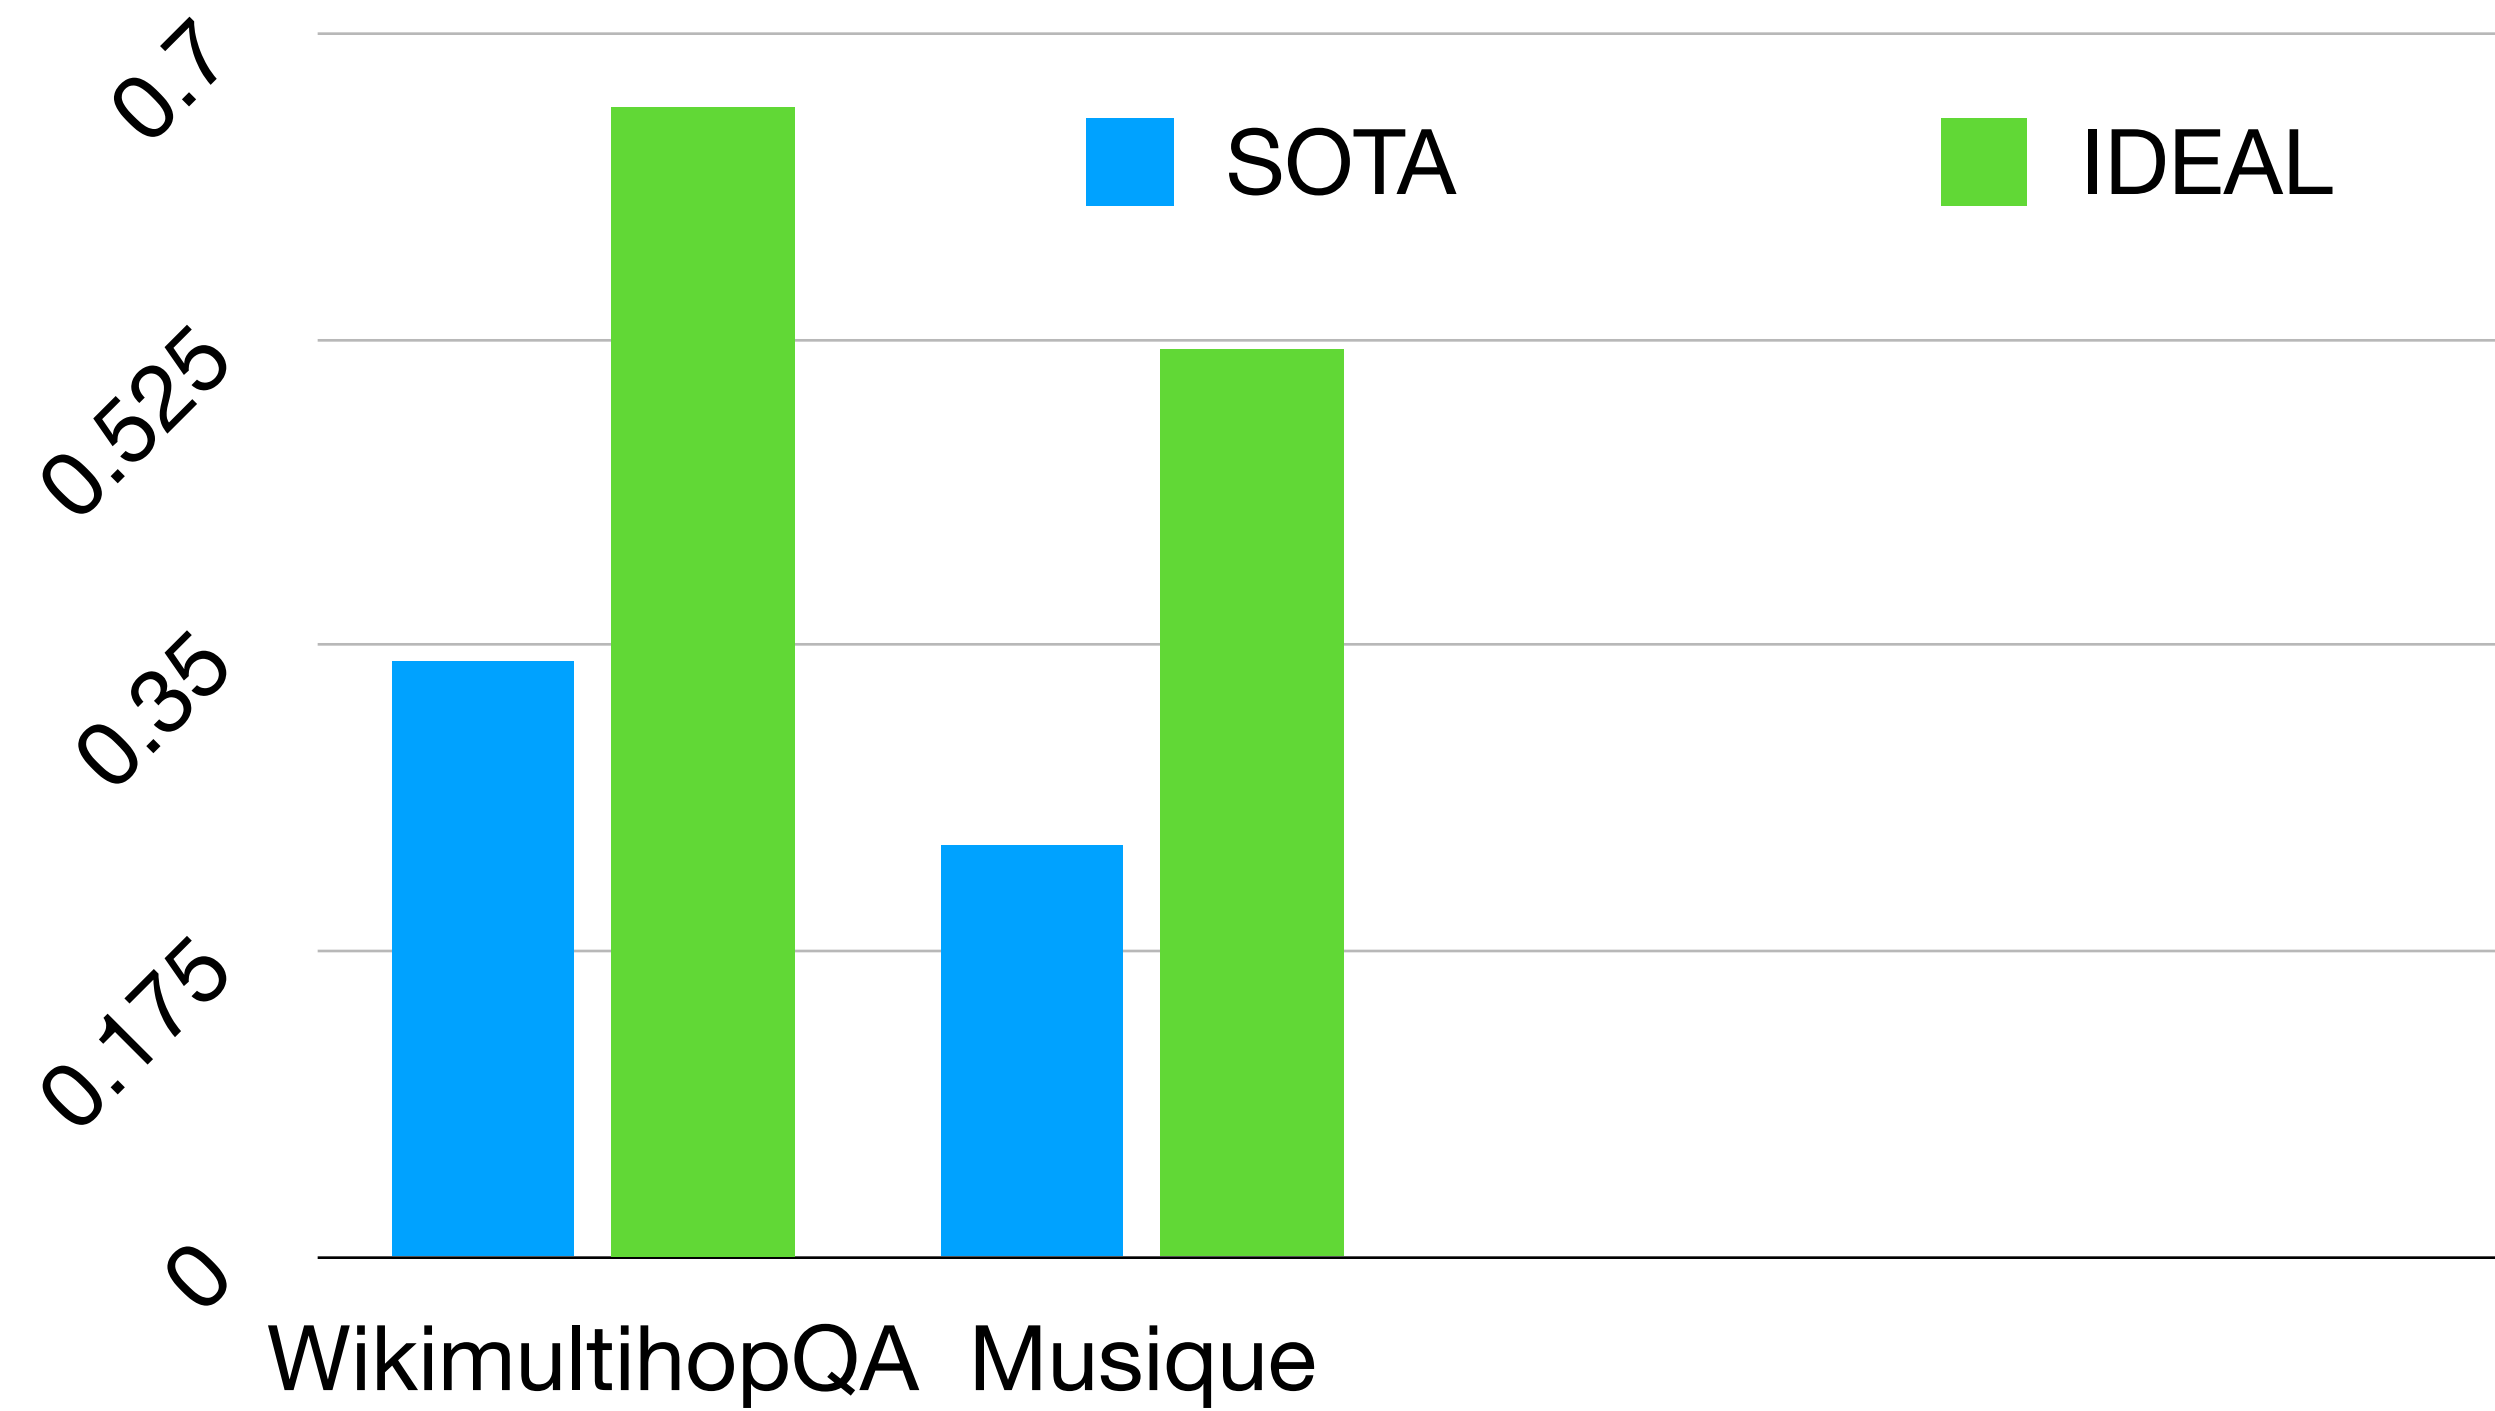
# Bounded recall problem

- RAG pipelines require the most relevant document to appear within top-5 or top-10 to fit i context of most affordable LLMs.
- Classical re-ranking approaches are limited by recall of first-stage retrieval.
- How do we capture more relevant documents ?
- How do we ensure the relevant documents are ranked higher and answer the question ?




# Retrieval and Reasoning Gap in complex QA

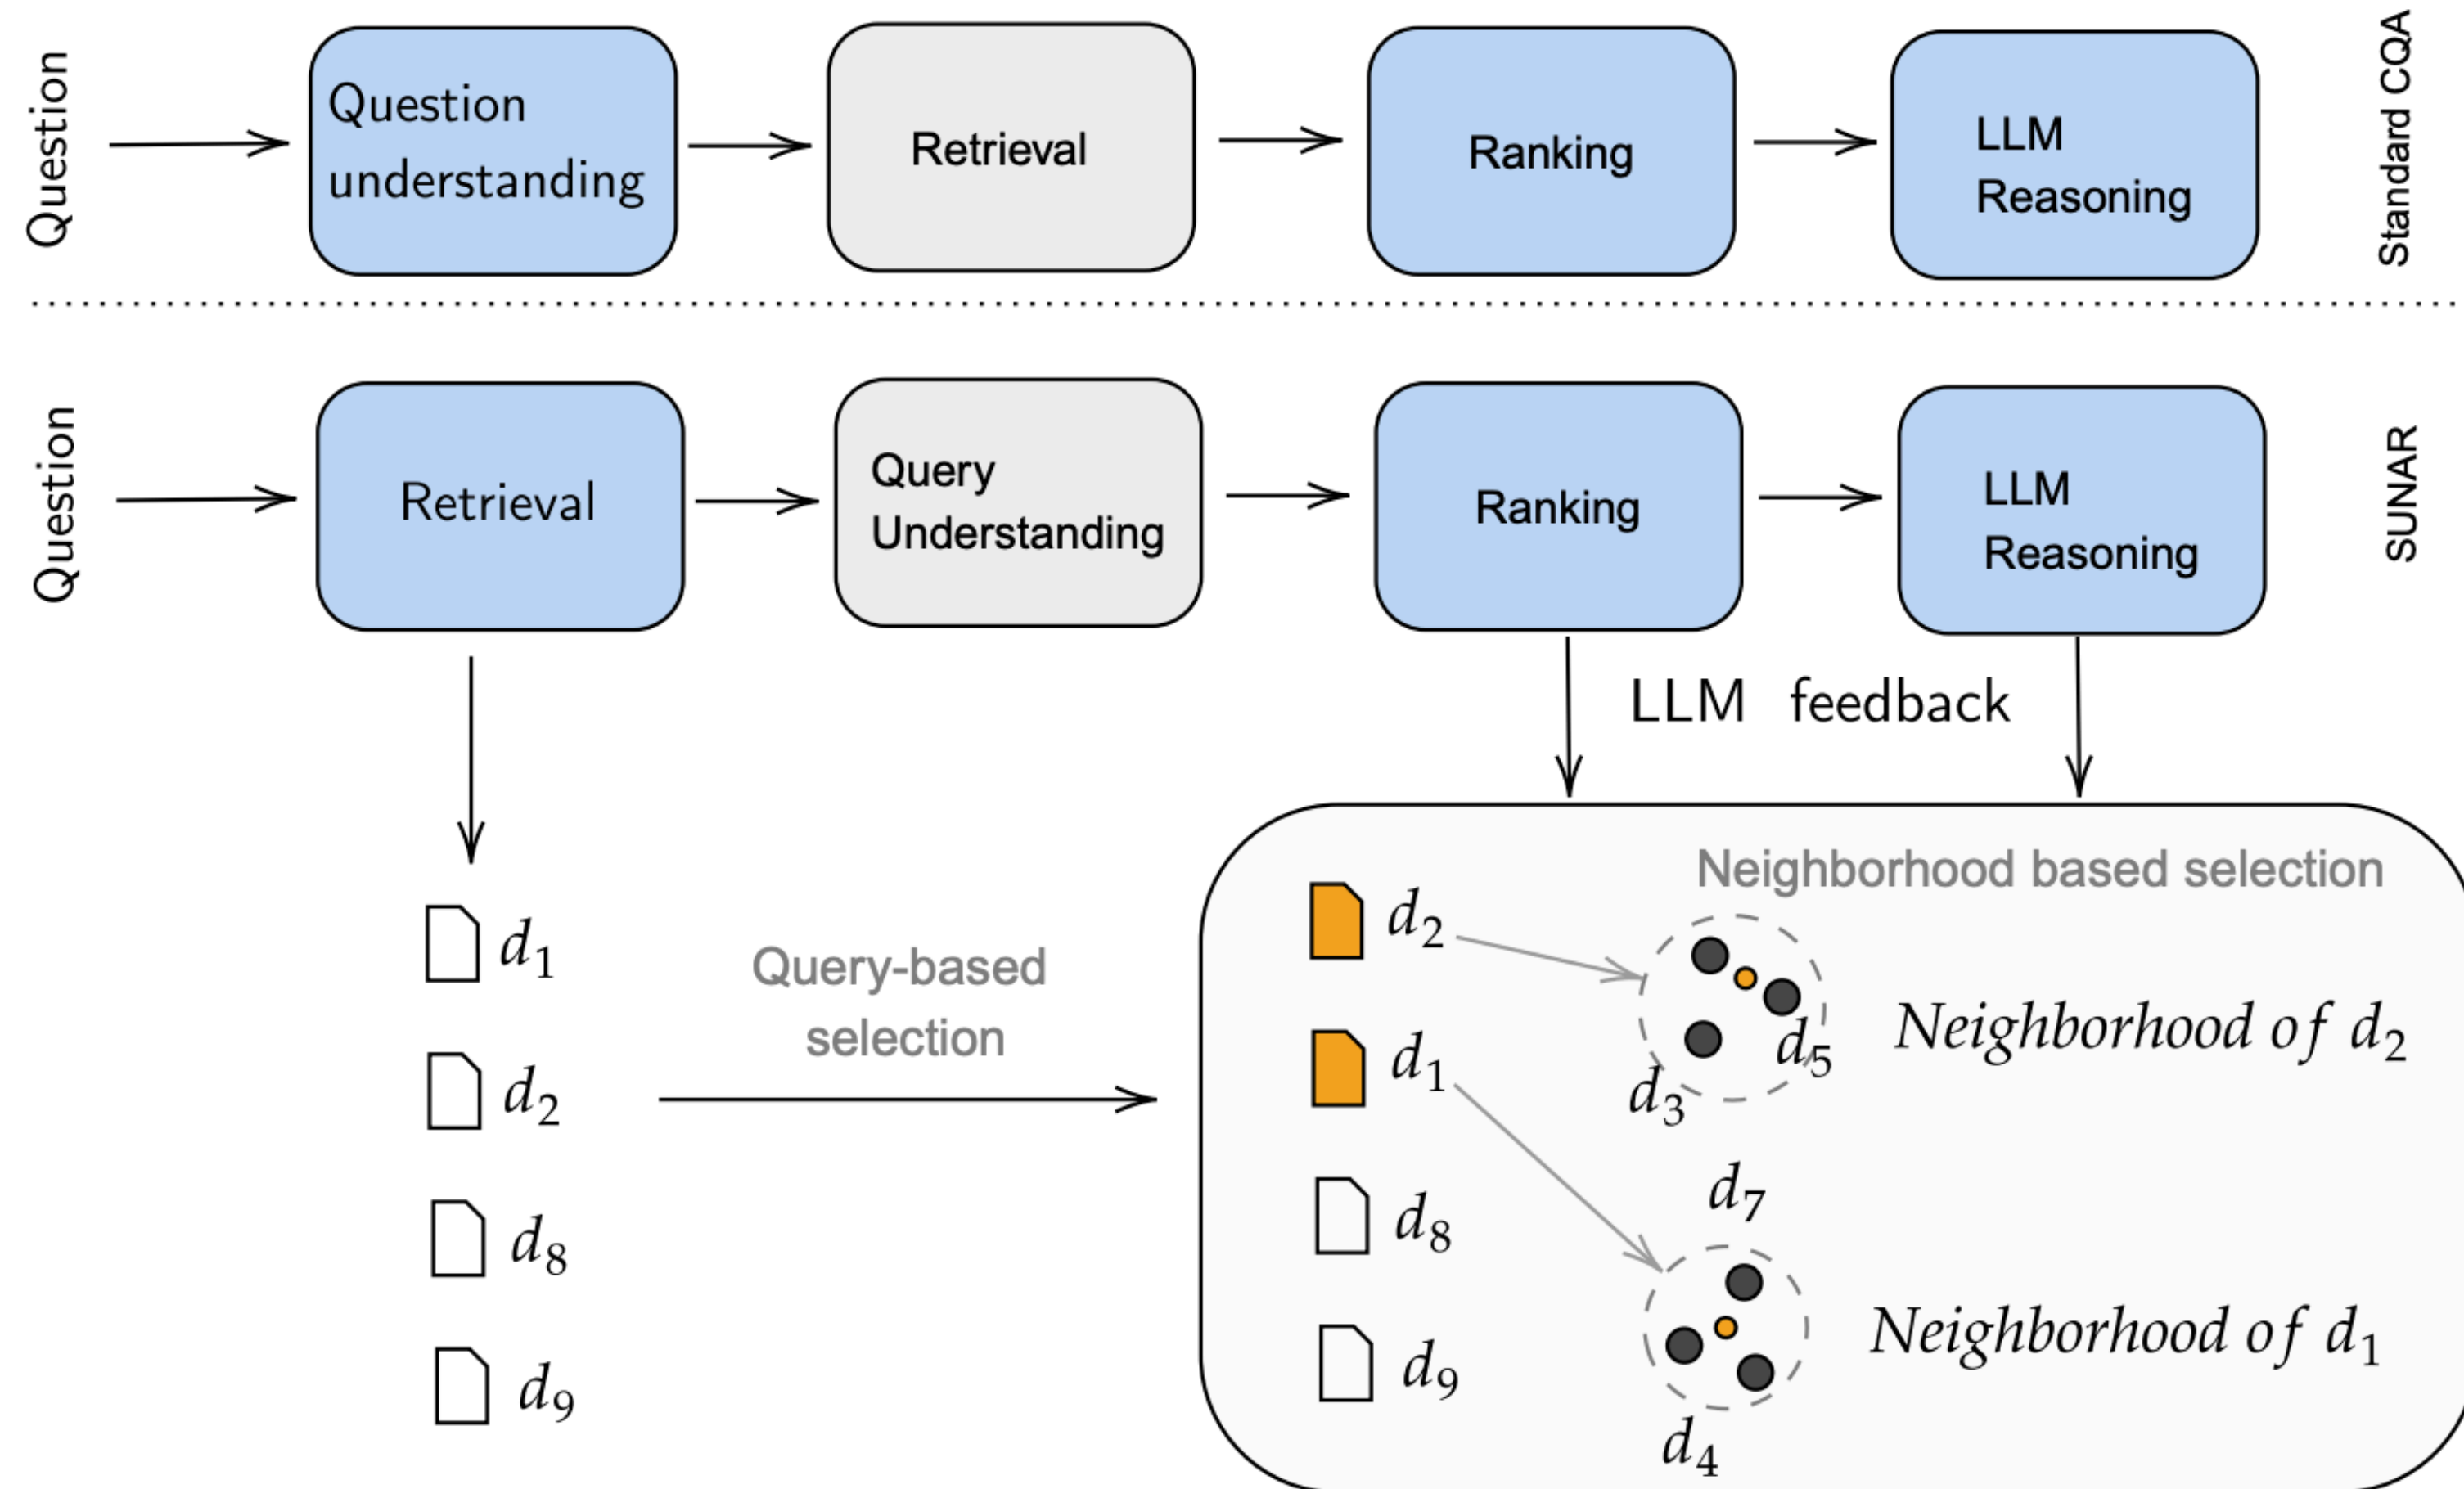
## Retrieval Gap



## Reasoning gap



# Semantic-Uncertainty based Neighborhood Aware Retrieval



# SUNAR- Deep Dive

---

## Algorithm 1 The SUNAR Algorithm

---


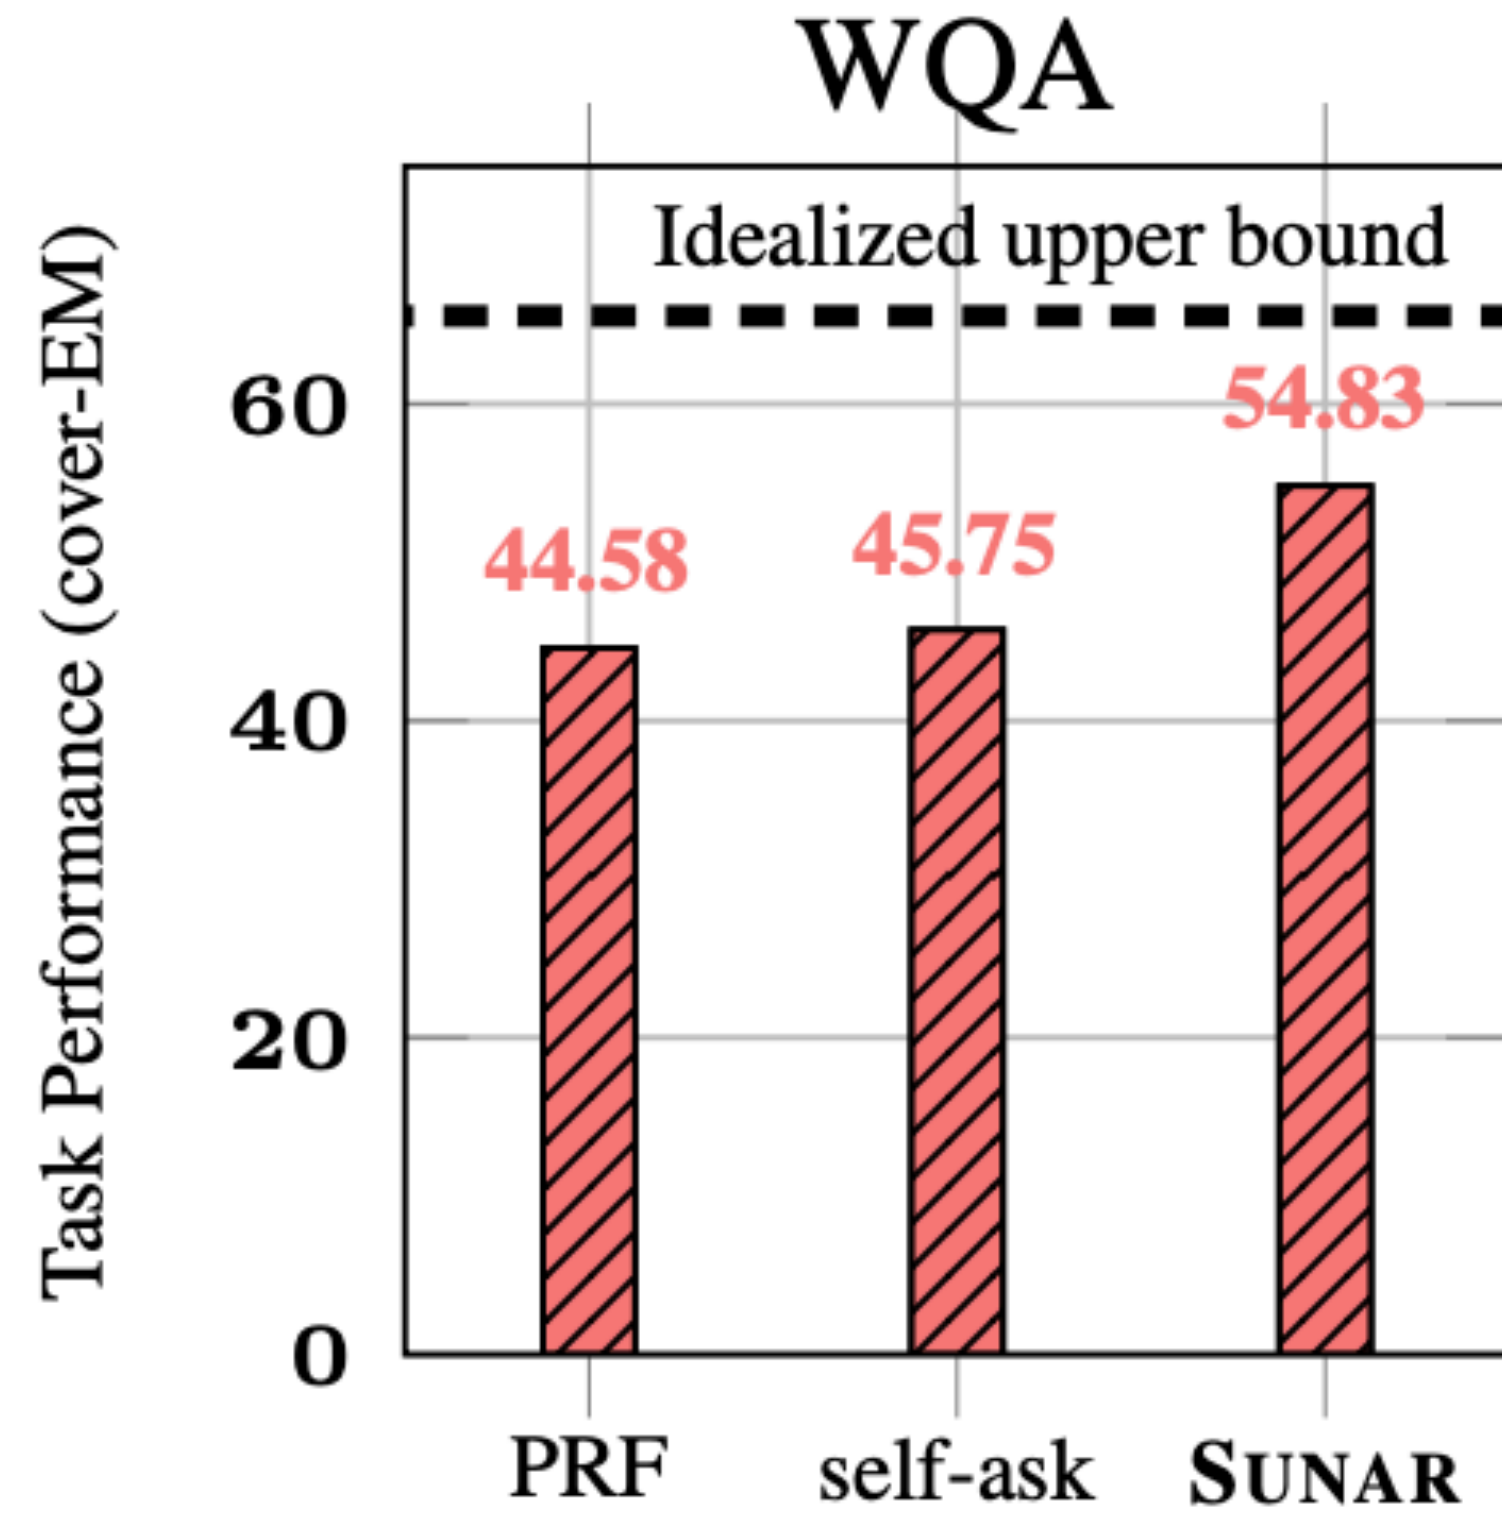
**Input:** Initial retrieved list  $R$ , batch size  $b$ , re-ranking budget  $c$ , document graph  $G$

**Output:** Re-Ranked pool  $R^+$

```
1:  $R^+ \leftarrow \emptyset$  ▷ Re-Ranking results
2:  $C \leftarrow R$  ▷ Re-ranking pool
3:  $N \leftarrow \emptyset$  ▷ Neighbor pool
4: do
5:    $B \leftarrow \text{SCORE}(\text{top } b \text{ from } P, \text{ subject to } c)$ 
6:    $\{sa_1 \dots sa_m\} \leftarrow \phi(\mathbb{P}_{LLM}(sq_1, B))$ 
7:    $\{ac_1 \dots ac_s\} \leftarrow \sigma(sa_1 \dots sa_m)$  ▷ Clustering
8:
9:    $B \leftarrow \text{RESCORE}(B, 1/s)$  ▷ Rescore batch
10:   $R^+ \leftarrow R^+ \cup B$  ▷ Add batch to results
11:
12:  // Discard Batches
13:   $R \leftarrow R \setminus B$ 
14:   $N \leftarrow N \setminus B$ 
15:   $N \leftarrow N \cup (\text{NEIGHBOURS}(B, G) \setminus R^+)$ 
16:
17:  // Alternate  $R$  and  $N$ 
18:   $C \leftarrow \begin{cases} R & \text{if } C = F \\ N & \text{if } C = N \end{cases}$ 
19: while  $|R^+| < c$ 
```

---

# Bridging retrieval gap and downstream reasoning enhancement



# Outperforms existing state-of-the-art approaches & LLM agnostic

Method	MQA	WQA
<b>Methods (w/o query understanding)</b>		
ZERO-SHOT-COT (Kojima et al., 2023)	8.62	30.42
FEW-SHOT-COT (Wei et al., 2023)	15.02	32.83
FEW-SHOT-COT +PRF (Li et al., 2022)	16.69	35.55
SUNAR <sub>R</sub> (ours)	21.32	40.96
<b>Methods (w/ query understanding)</b>		
Self-RAG (Asai et al., 2024)	17.80	35.25
ReAct (Yao et al., 2023)	21.41	43.25
Decomp (Khot et al., 2023)	21.01	44.08
SearChain (Xu et al., 2024)	21.72	44.42
SELF-ASK +PRF (Li et al., 2022)	20.28	44.58
SELF-ASK (Press et al., 2023)	24.84	45.75
<b>NAR (w/ query understanding) (ours)</b>		
SUNAR <sub>R</sub>	28.11	47.67
SUNAR	<b>32.75 †</b>	<b>54.83 †</b>
<b>Golden Evidence (Ideal Upper Bound)</b>		
FEW-SHOT-COT	44.28	65.55

Method	MQA	WQA
<b>gpt-40-mini</b>		
SELF-ASK	26.76	37.33
SUNAR	<b>32.19</b>	<b>48.16</b>
<b>Llama 3.1 (8B)</b>		
SELF-ASK	5.43	25.83
SUNAR	<b>13.82</b>	<b>39.52</b>
<b>Mistral v0.2 (7B)</b>		
SELF-ASK	7.84	27.72
SUNAR	<b>26.12</b>	<b>40.23</b>




# SUNAR helps tackle hallucination and knowledge gaps

Method	Evidences
<b>Question</b>	<b>Where was the director of film Ronnie Rocket born? [Dataset: WQA]</b>
SELF-ASK	[Evidence 1]: <b>This is a list of film series by director.</b> [Evidence 2]: <b>This is a list of notable directors in motion picture and television arts.</b> [Final Answer]: <b>Unknown</b>
SUNAR (ours)	[Evidence 1]: <b>Ronnie Rocket is an unfinished film project written by David Lynch, who also intended [...].</b> [Evidence 2]: <b>David Keith Lynch was born in Missoula, Montana, on January 20, 1946. His father [...].</b> [Final Answer]: <b>Missoula, Montana</b>
<b>Question</b>	<b>Who did the screenwriter for Good Will Hunting play in Dazed and Confused? [Dataset: MQA]</b>
SELF-ASK	[Evidence 1]: <b>Damon begins working alongside his younger brother, Stefan Salvatore, to resist greater [...].</b> [Evidence 2]: <b>Damon Salvatore is a fictional character in The Vampire Diaries. He is portrayed by Ian Somerhalder in the television.</b> [Final Answer]: <b>Damon Salvatore</b>
SUNAR (ours)	[Evidence 1]: <b>Damon and Ben Affleck wrote Good Will Hunting(1997), a screenplay[...].</b> [Evidence 2]: <b>Benjamin Affleck- Boldt( born August 15, 1972) is an American actor . He later appeared in the independent coming- of- age comedyDazed and Confused as Fred O'Bannion [...]"</b> [Final Answer]: <b>Fred O'Bannion</b>

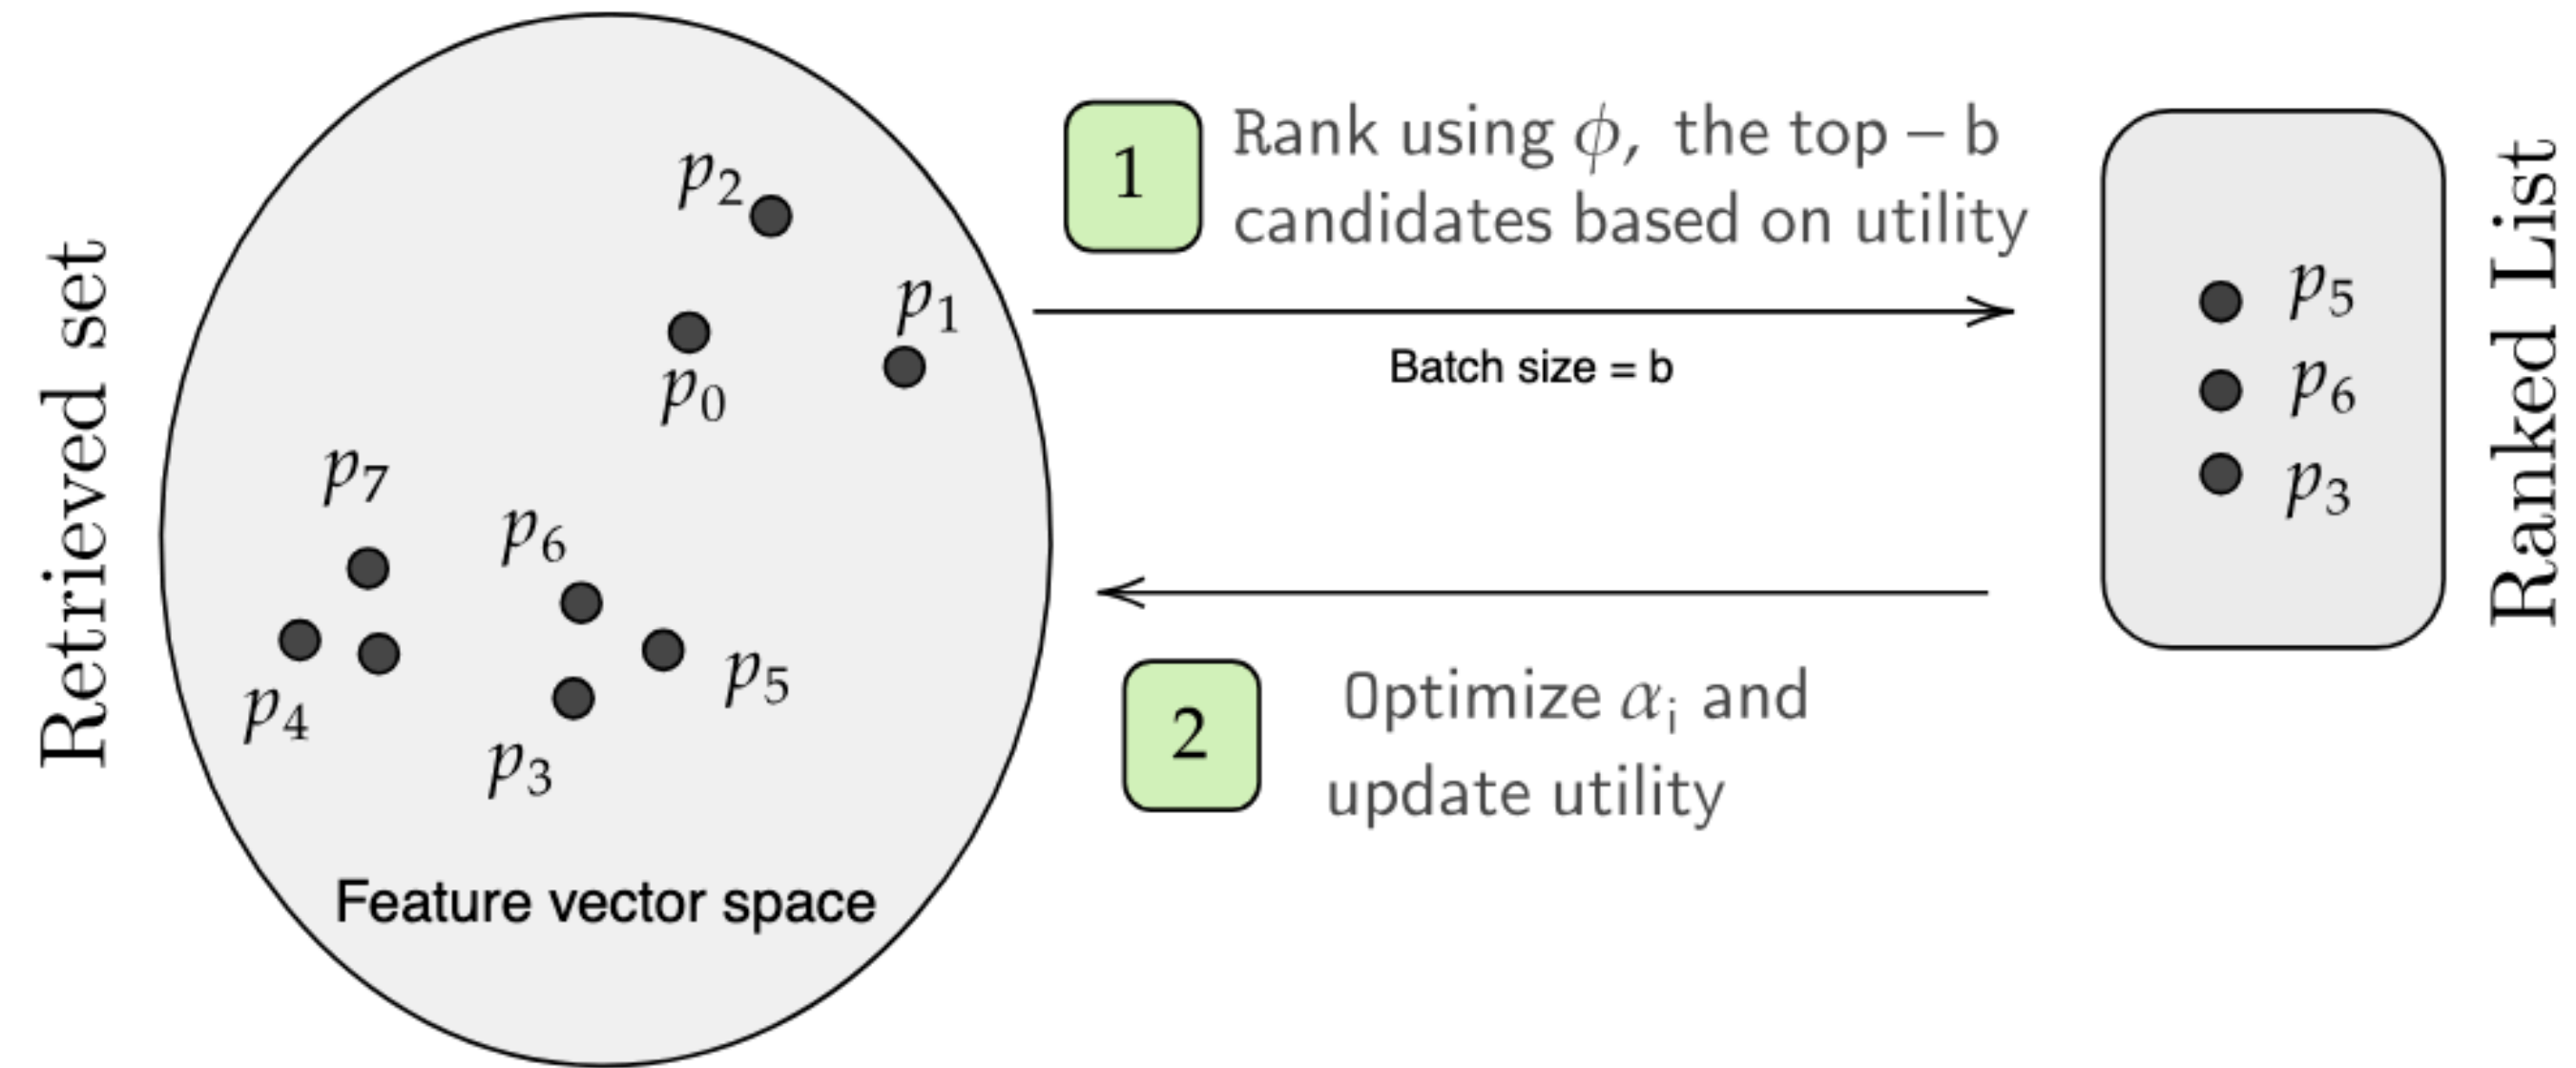
# Online Relevance Estimation

# Telescoping systems and drawbacks



- Telescoping approaches involve progressive filtering of documents through less-precise retrieval methods
- Key is capturing relevant documents with low retrieval scores that current approaches ignore.

# Online Relevance Estimation

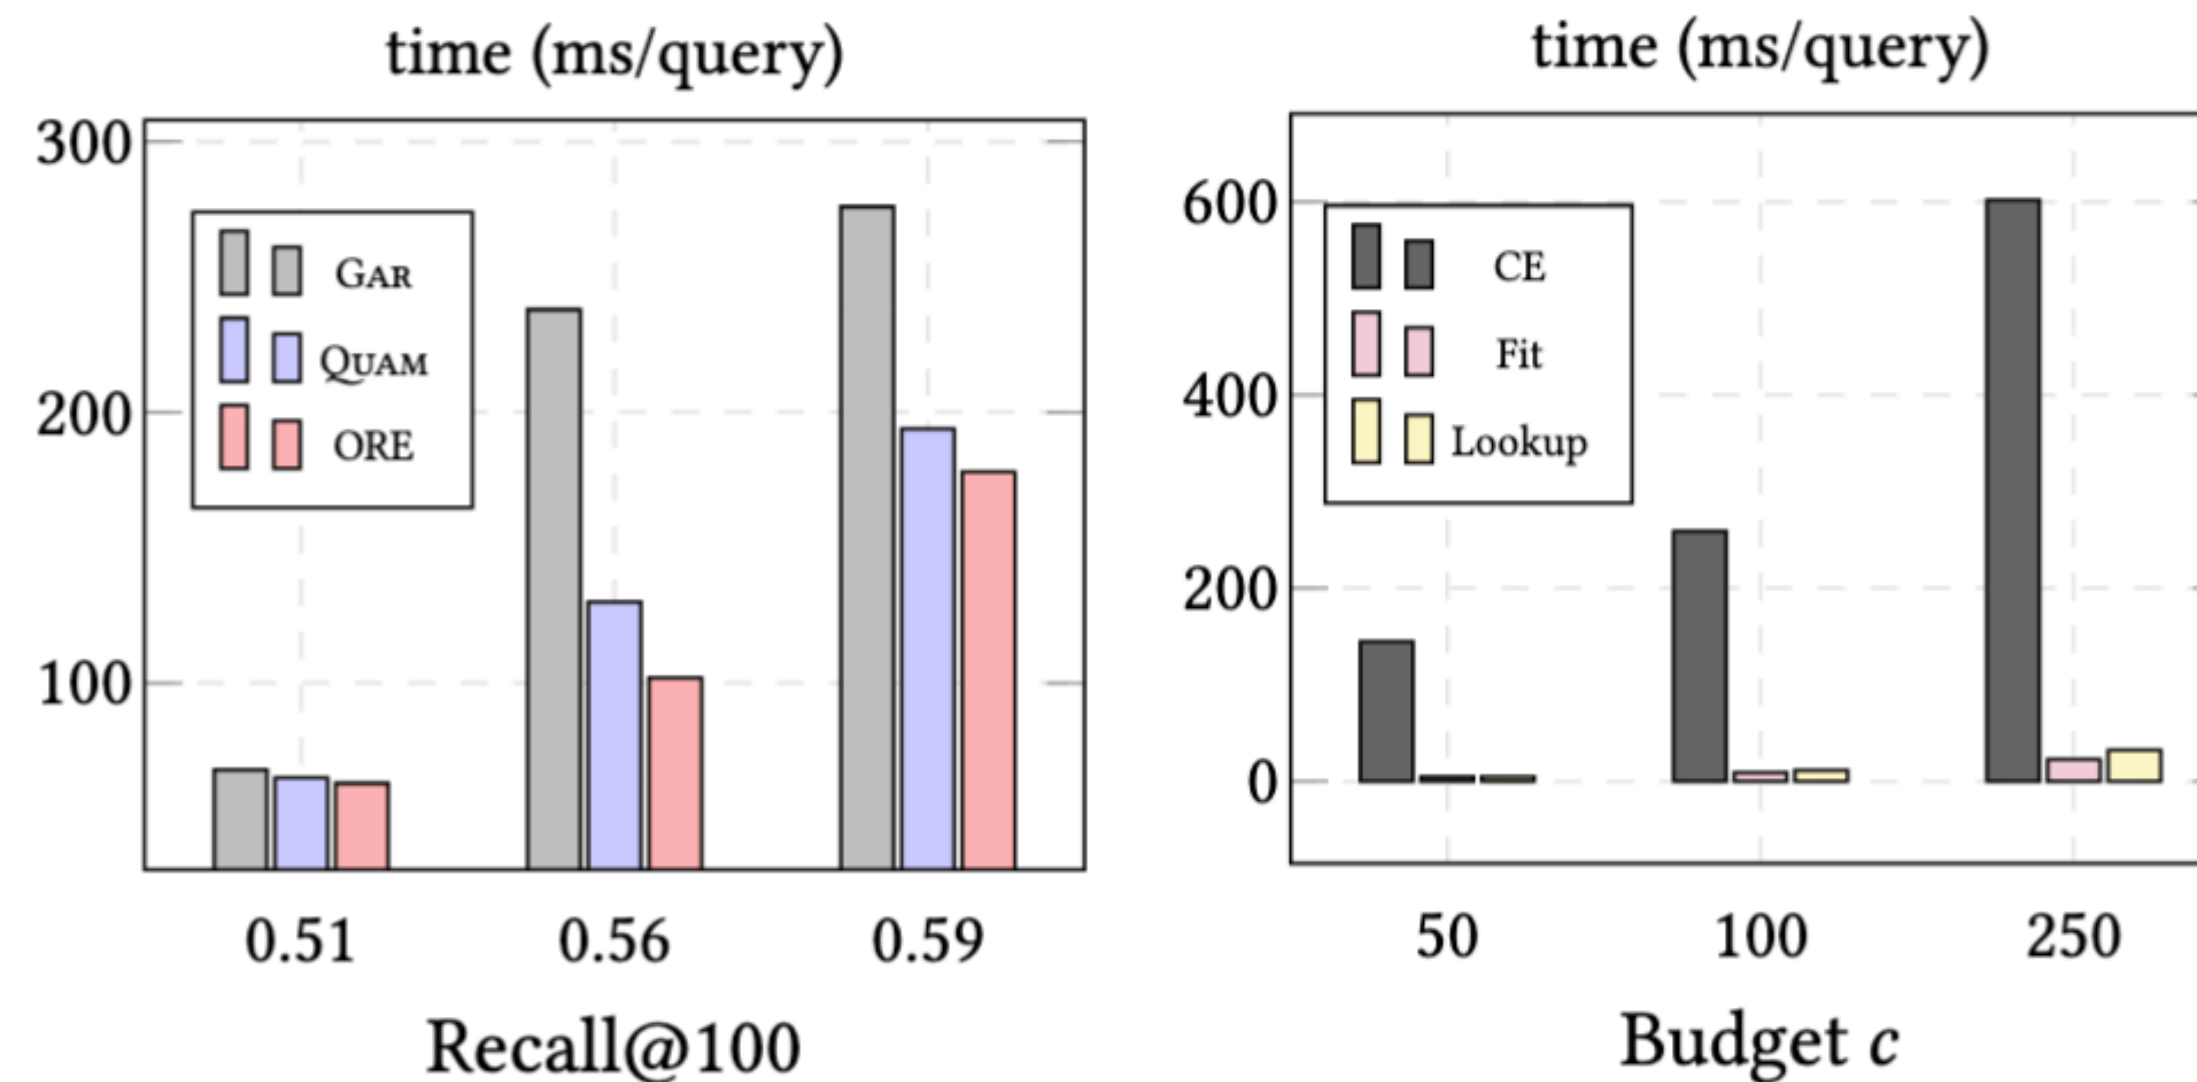


# Features are flexible

Feature	Notation	Taxonomy	Source		Description
			Offline	Online	
$x_1$	$BM25(q, d)$	Q2DAFF		✓	Lexical similarity between query and document.
$x_2$	$TCT(q, d)$	Q2DAFF		✓	Semantic similarity between query and document.
$x_3$	$RM3(q', d)$	D2DAFF		✓	Lexical similarity between expanded query using RM3 and document.
$x_4$	$BM25(d, d')$	D2DAFF	✓		Lexical similarity between pair of documents.
$x_5$	$TCT(d, d')$	D2DAFF	✓		Semantic similarity between pair of documents.
$x_6$	$LAF(d, d')$	D2DAFF	✓		Learnt affinity or similarity between pair of documents [34].

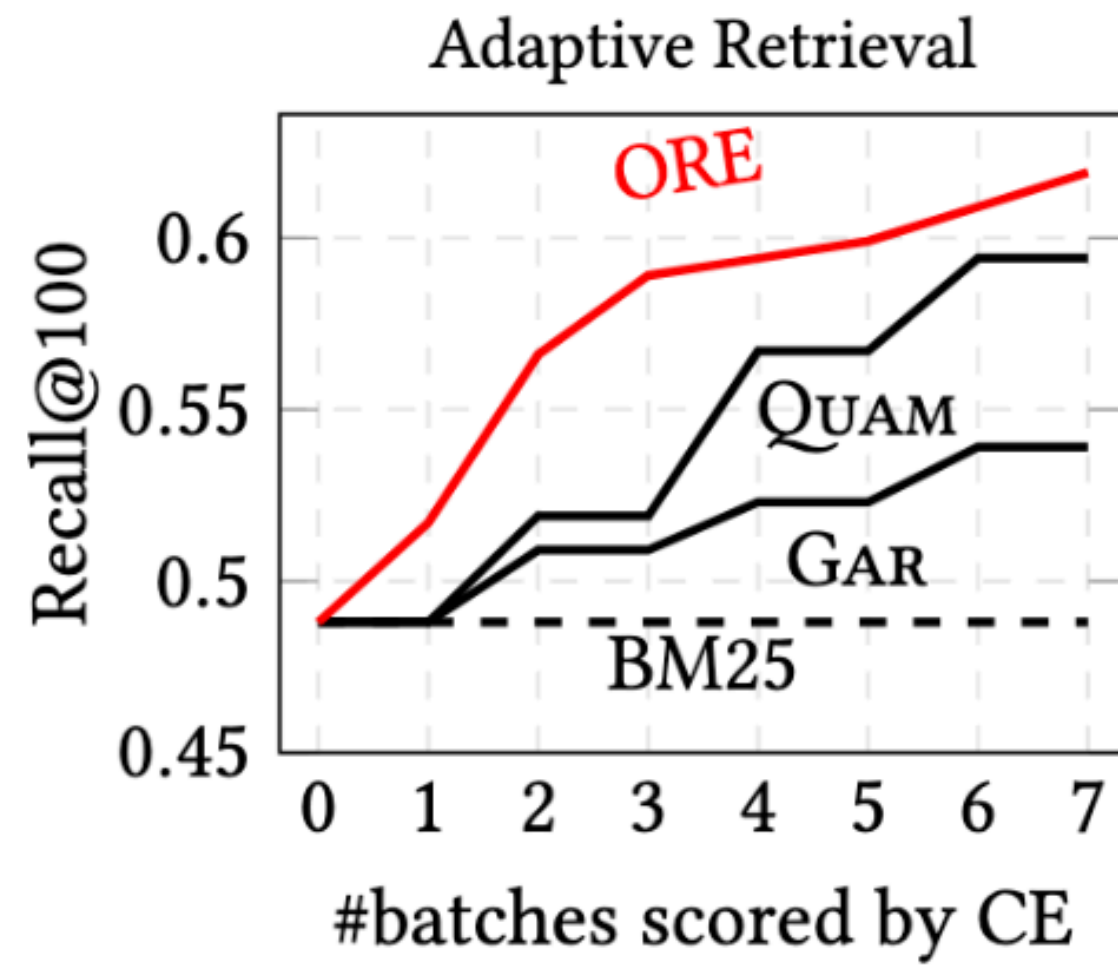
# Latency and Computational Efficiency

ORE offers 2x-7x speedup over SOTA based on ranker employed




The online estimation component takes **10x** less time than ranker calls



# Sample Efficiency of ORE



(a)



(b)



# Impressive Performance Gains

Dataset	Pipeline	$c = 50$		$c = 100$	
		nDCG@c	Recall@c	nDCG@c	Recall@c
DL21	<b>HYBRID</b>				
	RRF»MonoT5 [R]	0.576	0.401	0.558	0.520
	CC»MonoT5 [C]	0.584	0.419	0.569	0.545
	ORE	$R$ <b>0.604</b>	$R$ <b>0.444</b>	$RC$ <b>0.609</b>	$RC$ <b>0.609</b>
	<b>ADAPTIVE</b>				
	BM25»MonoT5 [B]	0.436	0.242	0.433	0.331
	w/ GAR <sub>BM25</sub> [G]	0.457	0.290	0.465	0.414
	w/ QUAM <sub>BM25</sub> [Q]	0.478	0.310	<b>0.499</b>	0.454
	w/ ORE <sub>BM25</sub>	$G_B^Q$ <b>0.503</b>	$G_B^Q$ <b>0.364</b>	$B$ 0.481	$G_B$ <b>0.463</b>
	w/ GAR <sub>TCT</sub> [G]	0.502	0.331	<b>0.520</b>	0.489
	w/ QUAM <sub>TCT</sub> [Q]	0.491	0.311	0.518	0.477
	w/ ORE <sub>TCT</sub>	$G_B^Q$ <b>0.532</b>	$G_B^Q$ <b>0.406</b>	$B$ 0.512	$B$ <b>0.502</b>
	<b>HYBRID</b>				
	RRF»MonoT5 [R]	0.452	0.260	0.430	0.341
CC»MonoT5 [C]	0.459	0.278	0.433	0.362	
ORE	$RC$ <b>0.481</b>	$R$ <b>0.297</b>	$RC$ <b>0.459</b>	$RC$ <b>0.389</b>	
DL22	<b>ADAPTIVE</b>				
	BM25»MonoT5 [B]	0.290	0.115	0.275	0.164
	w/ GAR <sub>BM25</sub> [G]	0.287	0.121	0.290	0.191
	w/ QUAM <sub>BM25</sub> [Q]	<b>0.308</b>	0.135	<b>0.303</b>	<b>0.196</b>
	w/ ORE <sub>BM25</sub>	<b>0.292</b>	<b>0.137</b>	0.284	0.195
	w/ GAR <sub>TCT</sub> [G]	0.329	0.157	<b>0.348</b>	0.256
	w/ QUAM <sub>TCT</sub> [Q]	0.329	0.155	0.334	0.237
	w/ ORE <sub>TCT</sub>	$G_B^Q$ <b>0.364</b>	$G_B^Q$ <b>0.206</b>	$B$ 0.342	$B$ <b>0.260</b>



## The Retrieval Gap




## The Reasoning Gap



# Numerical and Compositional Reasoning


**Claim:** Repealing the sales tax on boats in Rhode Island has spawned 2,000 companies, 7,000 jobs and close to \$2 billion a year in sales activity



How many  
companies were  
there before the  
tax ?

# Numerical and Compositional Reasoning

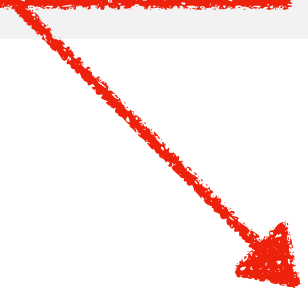
**Claim:** Repealing the sales tax on boats in Rhode Island has spawned 2,000 companies, 7,000 jobs and close to \$2 billion a year in sales activity



How many jobs  
were there  
before the tax ?

# Numerical and Compositional Reasoning

**Claim:** Repealing the sales tax on boats in Rhode Island has spawned 2,000 companies, 7,000 jobs and close to \$2 billion a year in sales activity



What was the  
annual sales there  
before the tax ?

# Numerical and Compositional Reasoning

**Claim:** Repealing the sales tax on boats in Rhode Island has spawned 2,000 companies, 7,000 jobs and close to \$2 billion a year in sales activity

Need More than Prompting LLMs



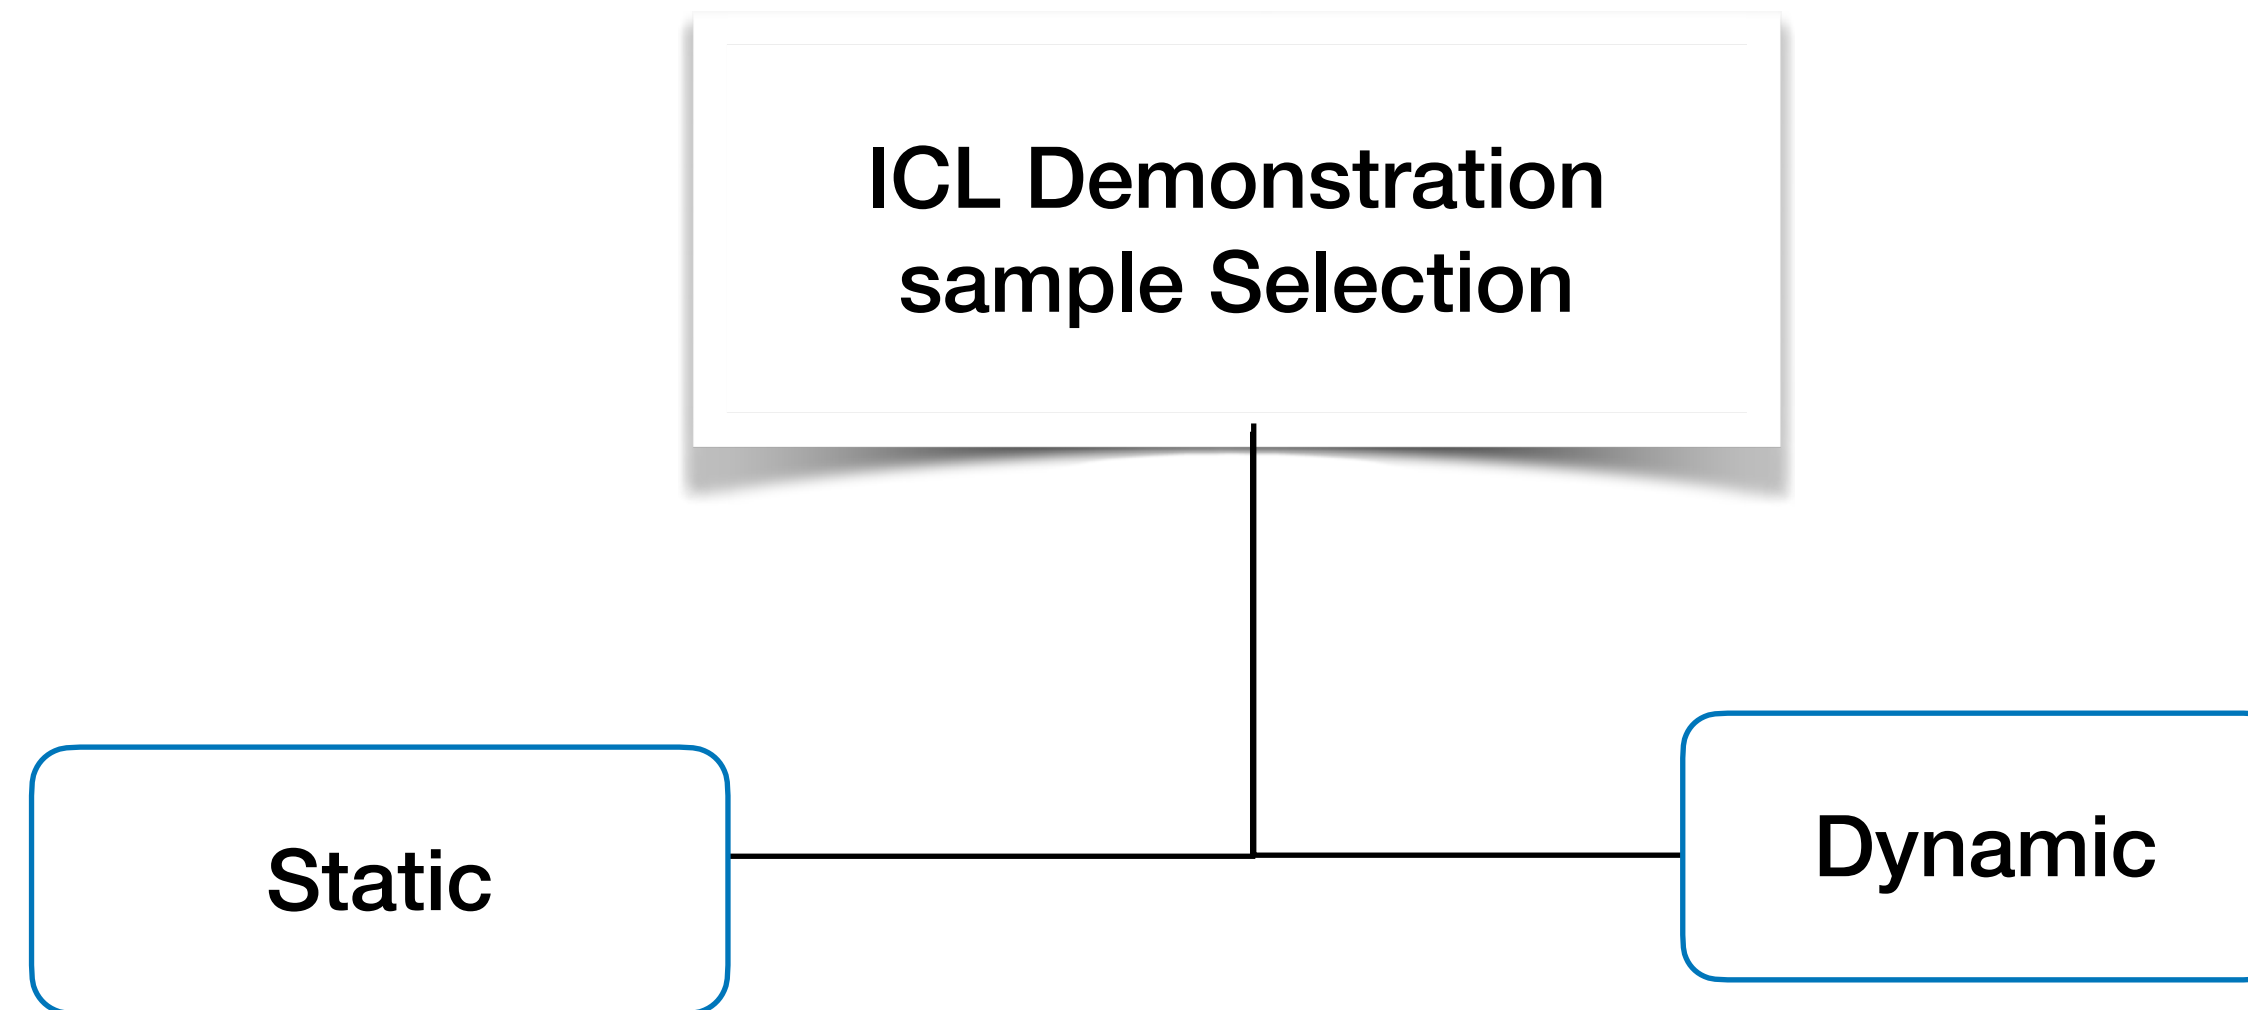
Rationales or Explanations

Need to compose abilities required to solve the task

Could be achieved through fine-tuning on required abilities.  
Result: Smaller 1M param models outperform larger 1B param models

Or skill composition through In-Context Learning

# Demonstration Samples is all you need ?



```
FinQA Prompt

Instruction:You are a helpful, respectful and honest assistant helping to solve math word problems or tasks requiring reasoning or math, using the information from given table and text.

Exemplars :
Read the following table, and then answer the question:
[Table]: Year | 2016 | 2015 | 2014 |
share-based compensation expense | 30809 | 21056 | 29793 |
income tax benefit | 9879 | 6907 | 7126 |
[Question]: how much percent did the income tax benefit increase from 2014 to 2016?
[Explanation]:  $x_0 = (9879 - 7126)$ ,
ans=(  $x_0/7126$  )
[Answer]: The answer is increased 38.6%
...
...

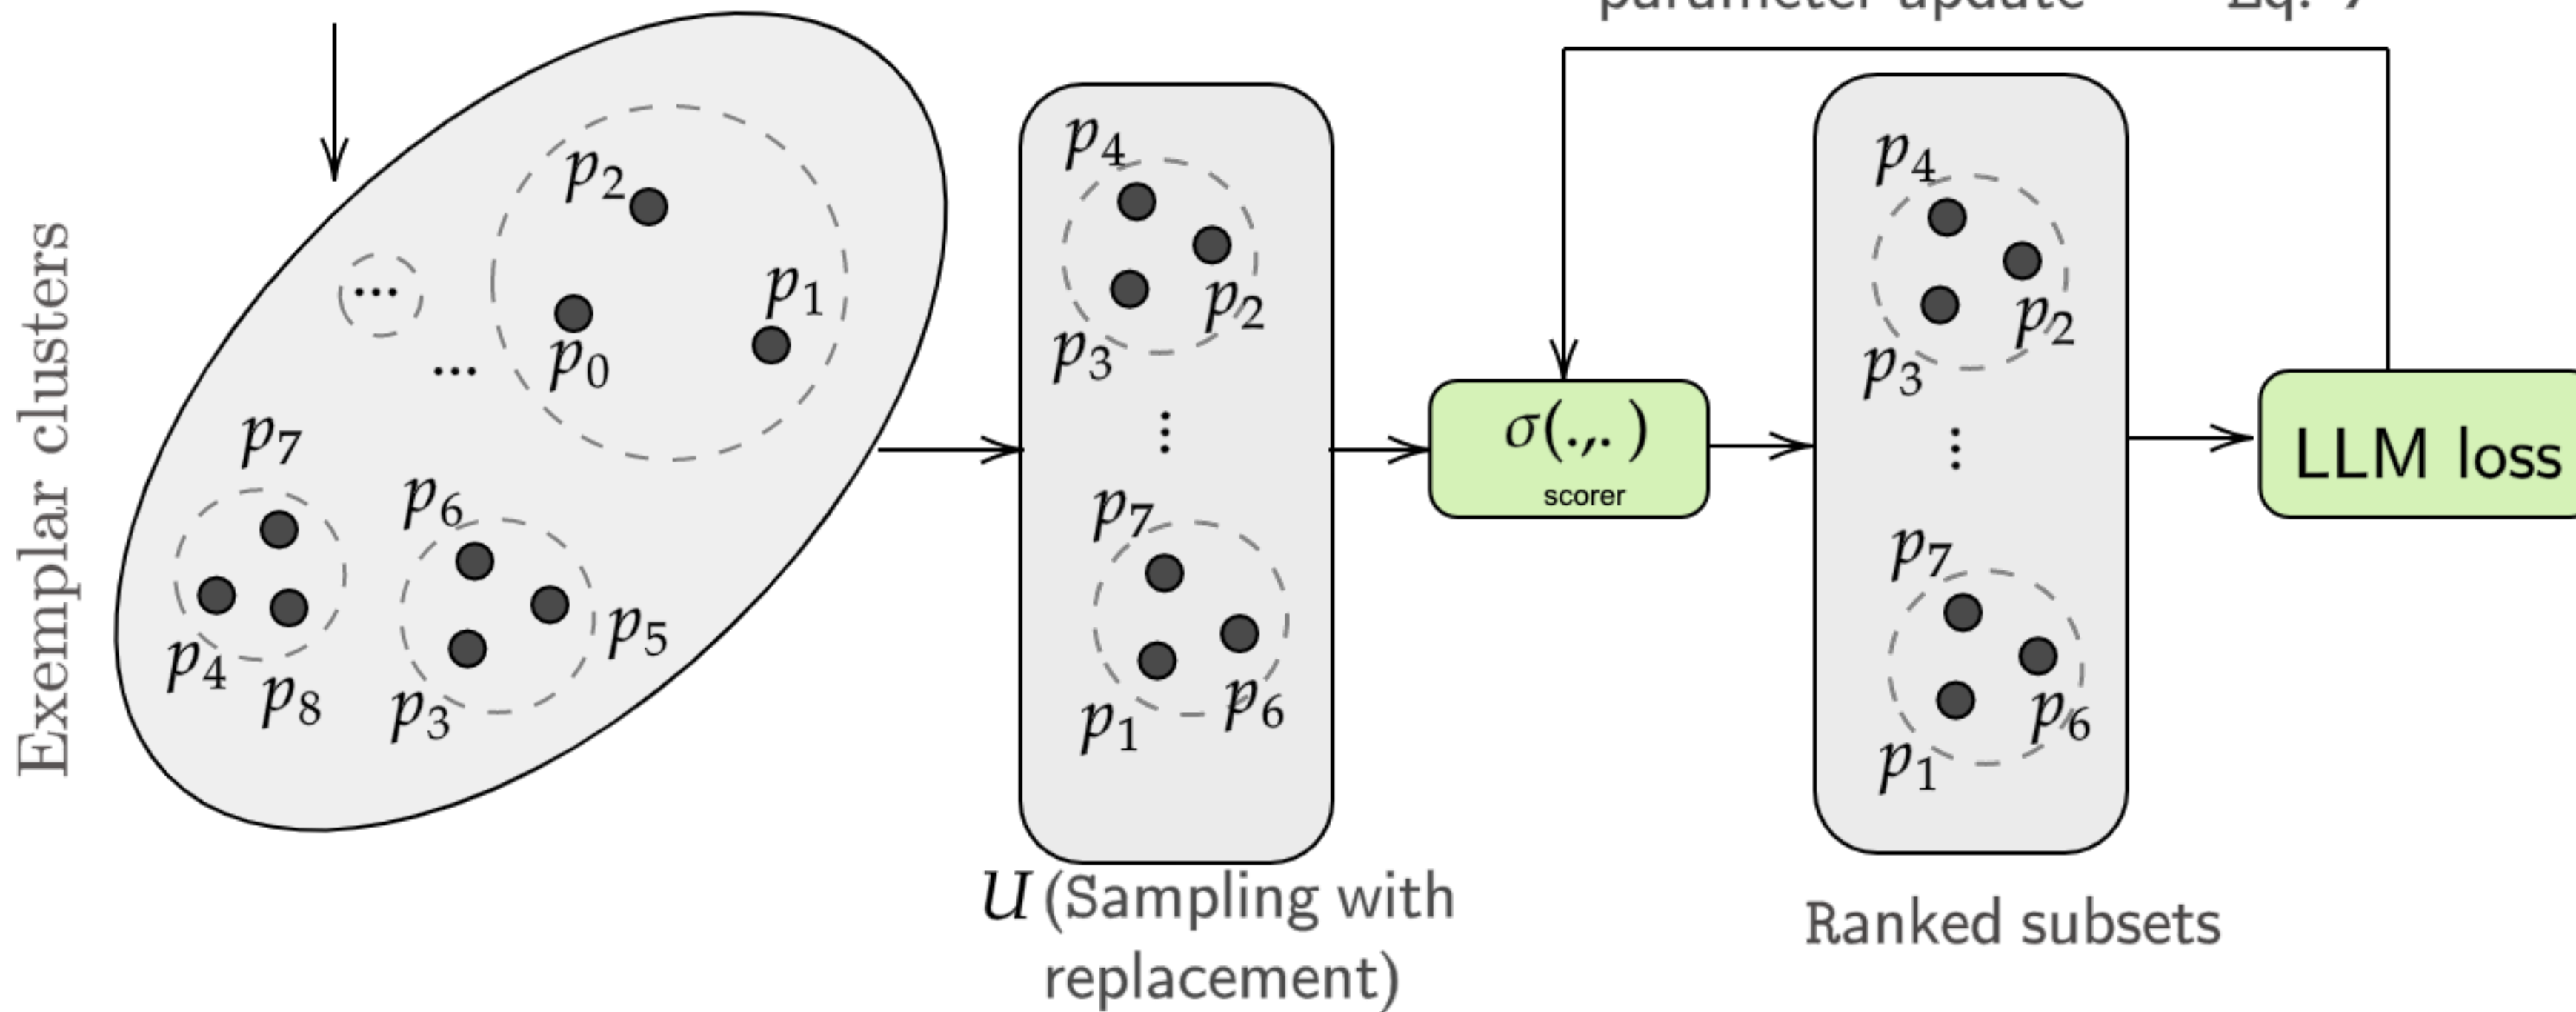
Test Input : Read the following table, and then answer the question: Table: Question:
Explanation: [INS] Answer: [INS]
```

# Smart Exploration and Exploitation for ICL Exemplars

## Exemplars

$p_1$  : While purchasing groceries ram bought **5 apples** ...

$p_2$  : Ephraim has **two machines** that make necklaces ...



# Loss Modeling (Approximation) for efficient selection

Subset of k Exemplars ( $S \subseteq \mathcal{S}$ )

Loss modelling function;  
Approximating  $L(S, \mathcal{V})$

$$\sigma(\vec{\alpha}, S) = \frac{1}{m} \sum_{j=1}^m \sum_{i=1}^n \alpha_i (x_i \in S) E_{ij} \quad (1)$$

$\alpha_i$  ith exemplars contribution, low if important exemplar

Any transformer based encoder

$$E_{ij} = \frac{\phi(x_i)^T \phi(u_j)}{\|\phi(x_i)\| \|\phi(u_j)\|}$$

$\phi(x_i)$  ith exemplar,  $x_i \in S$

$\phi(u_j)$  ith validation sample,  $u_i \in \mathcal{V}$



# Efficient Estimation of parameters

Update parameters to reduce approximation error

set of  $l$  subsets at timestep  $t$  with lowest validation loss

Validation Set

$$\mathcal{L}(\vec{\alpha}; U_t, V_t) = \sum_{S \in U_t} (L(S, \mathcal{V}) - \sigma(\vec{\alpha}, S))^2 + \sum_{S' \sim V_t} (L(S', \mathcal{V}) - \sigma(\vec{\alpha}, S'))^2$$

Remaining subsets;  $V_t \leftarrow \mathcal{U} \setminus U_t$ , where  $\mathcal{U} \subset \mathcal{S}$

Negative samples from high-loss set  $V_t$

Estimating loss here involves LLM calls and equivalent to arm pulling

# Summary

---

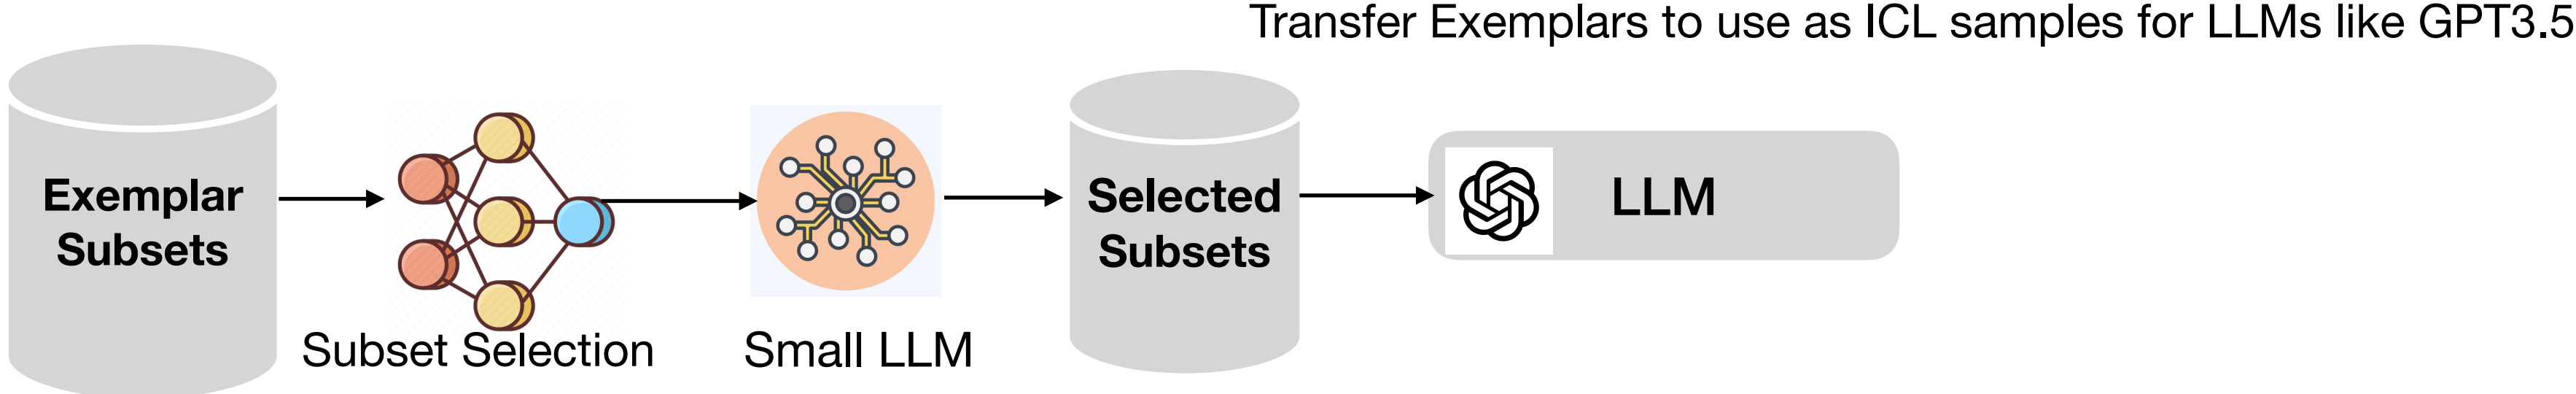
## Algorithm 1: EXPLORA

---

```
1 Input:  $\mathcal{U} \subseteq \mathcal{S}$ :  $\triangleright$  Initial exemplar subsets
2 Initialize:  $U_0 \leftarrow$  set of random  $l$  subsets from  $\mathcal{U}$ 
3  $t \leftarrow 0$ 
4  $\vec{\alpha} \leftarrow \mathcal{N}(0, 1)$   $\triangleright$  Sampling from a gaussian
5 while  $t < T$  do
6   Let  $V_t \leftarrow \mathcal{U} \setminus U_t$ 
7    $\vec{\alpha}_t \leftarrow \min_{\vec{\alpha}} \mathcal{L}(\vec{\alpha}, U_t, V_t)$   $\triangleright$  Eq. in previous slide
8    $S_t^* = \arg \min_{S \in V_t} \sigma(\vec{\alpha}_t, S)$   $\triangleright$  Lowest loss
   subset
9    $\tilde{S}_t = \arg \max_{S \in U_t} \sigma(\vec{\alpha}_t, S)$   $\triangleright$  Highest loss
   subset
10  if  $\sigma(\vec{\alpha}_t, S_t^*) < \sigma(\vec{\alpha}_t, \tilde{S}_t)$  then
11     $U_t \leftarrow U_t \setminus \{\tilde{S}_t\}$   $\triangleright$  Remove  $\tilde{S}_t$ 
12     $U_{t+1} \leftarrow U_t \cup \{S_t^*\}$   $\triangleright$  add  $S_t^*$ 
13  end
14   $t \leftarrow t + 1$ 
15 end
16 Output:  $U_T$ ; Set of  $l$  subsets from  $\mathcal{U}$  which have the
   lowest validation loss
```

---


# Tune and Transfer



# EXPLORA is Robust (Low Variance across test samples)

Datasets	GSM	Aqua	Tab	Fin
Zero-Shot COT	$\pm 5.18$	$\pm 7.08$	$\pm 1.84$	$\pm 4.50$
Few-Shot COT	$\pm 4.48$	$\pm 12.03$	$\pm 1.66$	$\pm 4.76$
KNN	$\pm 3.76$	$\pm 5.49$	$\pm 1.27$	$\pm 4.17$
MMR	$\pm 4.00$	$\pm 10.53$	$\pm 1.68$	$\pm 6.10$
Graph Cut	$\pm 6.38$	$\pm 8.18$	$\pm 2.03$	$\pm 5.29$
Facility Location	$\pm 4.23$	$\pm 6.71$	$\pm 1.74$	$\pm 4.94$
LENS	$\pm 5.04$	$\pm 6.67$	$\pm 1.72$	$\pm 5.81$
<b>EXPLORA</b>	<b><math>\pm 3.39</math></b>	<b><math>\pm 4.93</math></b>	<b><math>\pm 1.45</math></b>	<b><math>\pm 3.41</math></b>

# EXPLORA is Resource Efficient



# Results Transfers Well (L for Llama and M for Mistral)

<b>Method</b>	<b>T</b>	<b>GSM</b>	<b>Aqua</b>	<b>Tab</b>	<b>Fin</b>
<b>EXP</b>	L	79.07	53.94	79.65	54.66
	M	77.86	53.54	77.41	59.46
<b>EXP+SC</b>	L	85.82	63.78	86.76	61.16
	M	86.35	63.39	85.52	64.52
<b>EXP+KNN+SC</b>	L	85.89	64.17	85.74	63.64
	M	85.14	62.20	86.29	65.12
<b>EXP+MMR+SC</b>	L	86.20	62.99	87.81	64.60
	M	86.13	63.78	86.96	64.60

Prompts are transferred from Llama or Mistral to GPT3.5-turbo

# A Recap

- Efficiency and Effectiveness are critical for practical robust RAG pipelines.
- Telescoping systems are limited in efficiency and suffer from Recall Boundedness.
- LLMs are still limited in reasoning.
- Test Time scaling for Retrieval is central to robust pipelines for complex knowledge intensive tasks.
- Careful selection of exemplars help in transferring abilities to LLMs through ICL.

# Conclusion - Research Vision

- End-End Test Time Reasoning (TTR) has huge scope.
  - How do we incorporate Reasoning feedback (LLM) to improve retrieval
  - How can retrieval improve reasoning.
  - How to do this efficiently?
  
- How can we do this in a scalable manner ?