

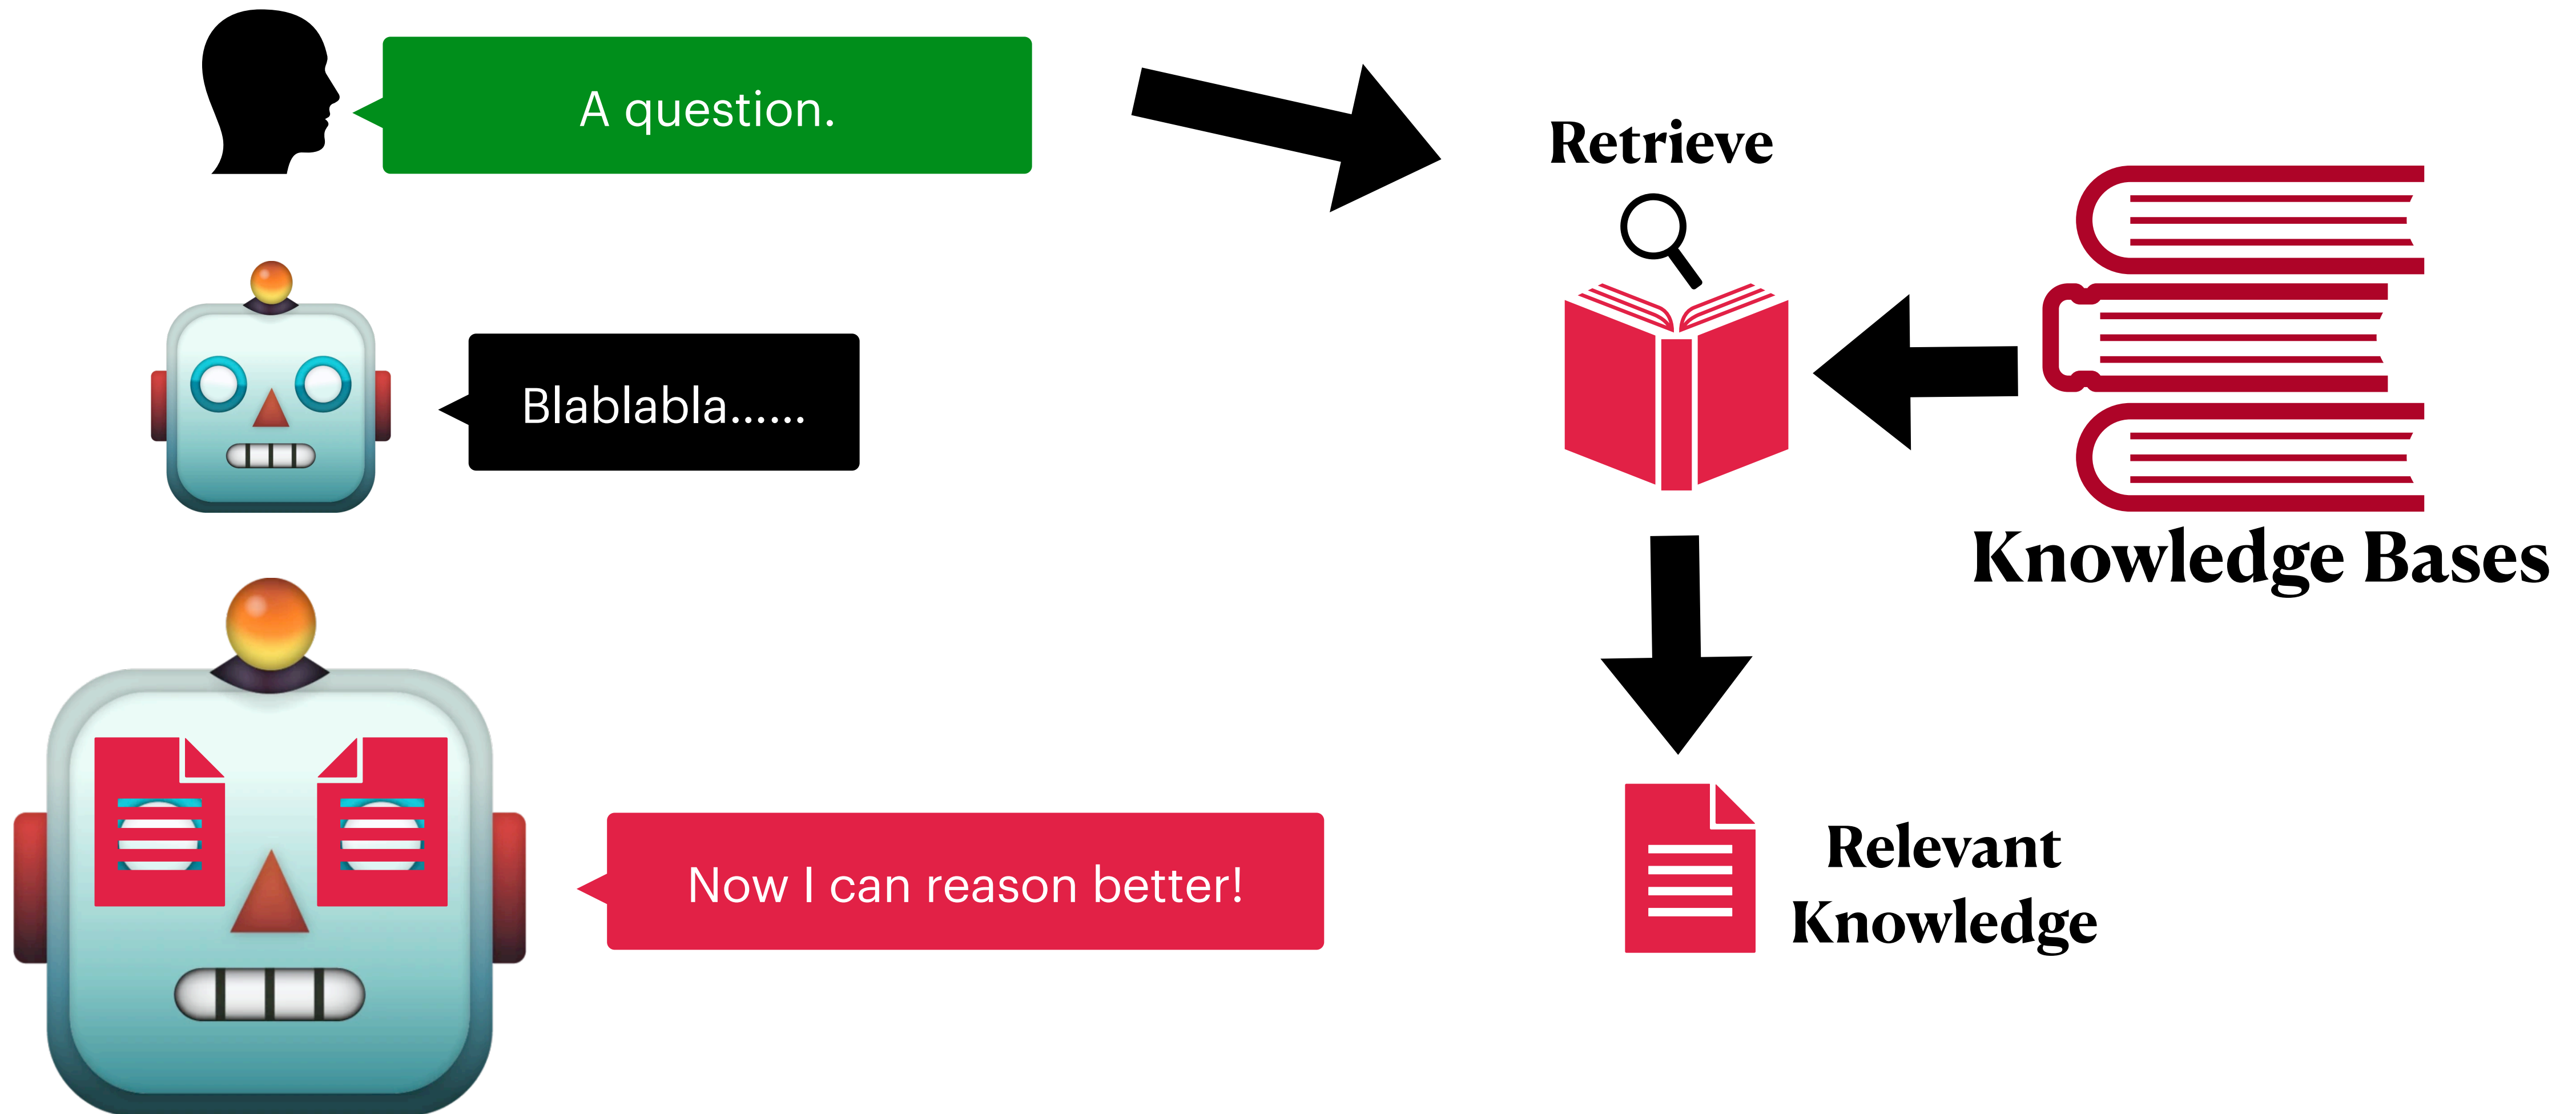
Knowledge Augmentation for Reasoning in Language

Part of the ACL 2023 Tutorial
“Complex Reasoning in Natural Language”

(Bill) Yuchen Lin

<https://yuchenlin.xyz>

Knowledge Augmentation for Reasoning




Knowledge Sources

🕸 Structured Knowledge



📖 Un/Semi-structured Knowledge

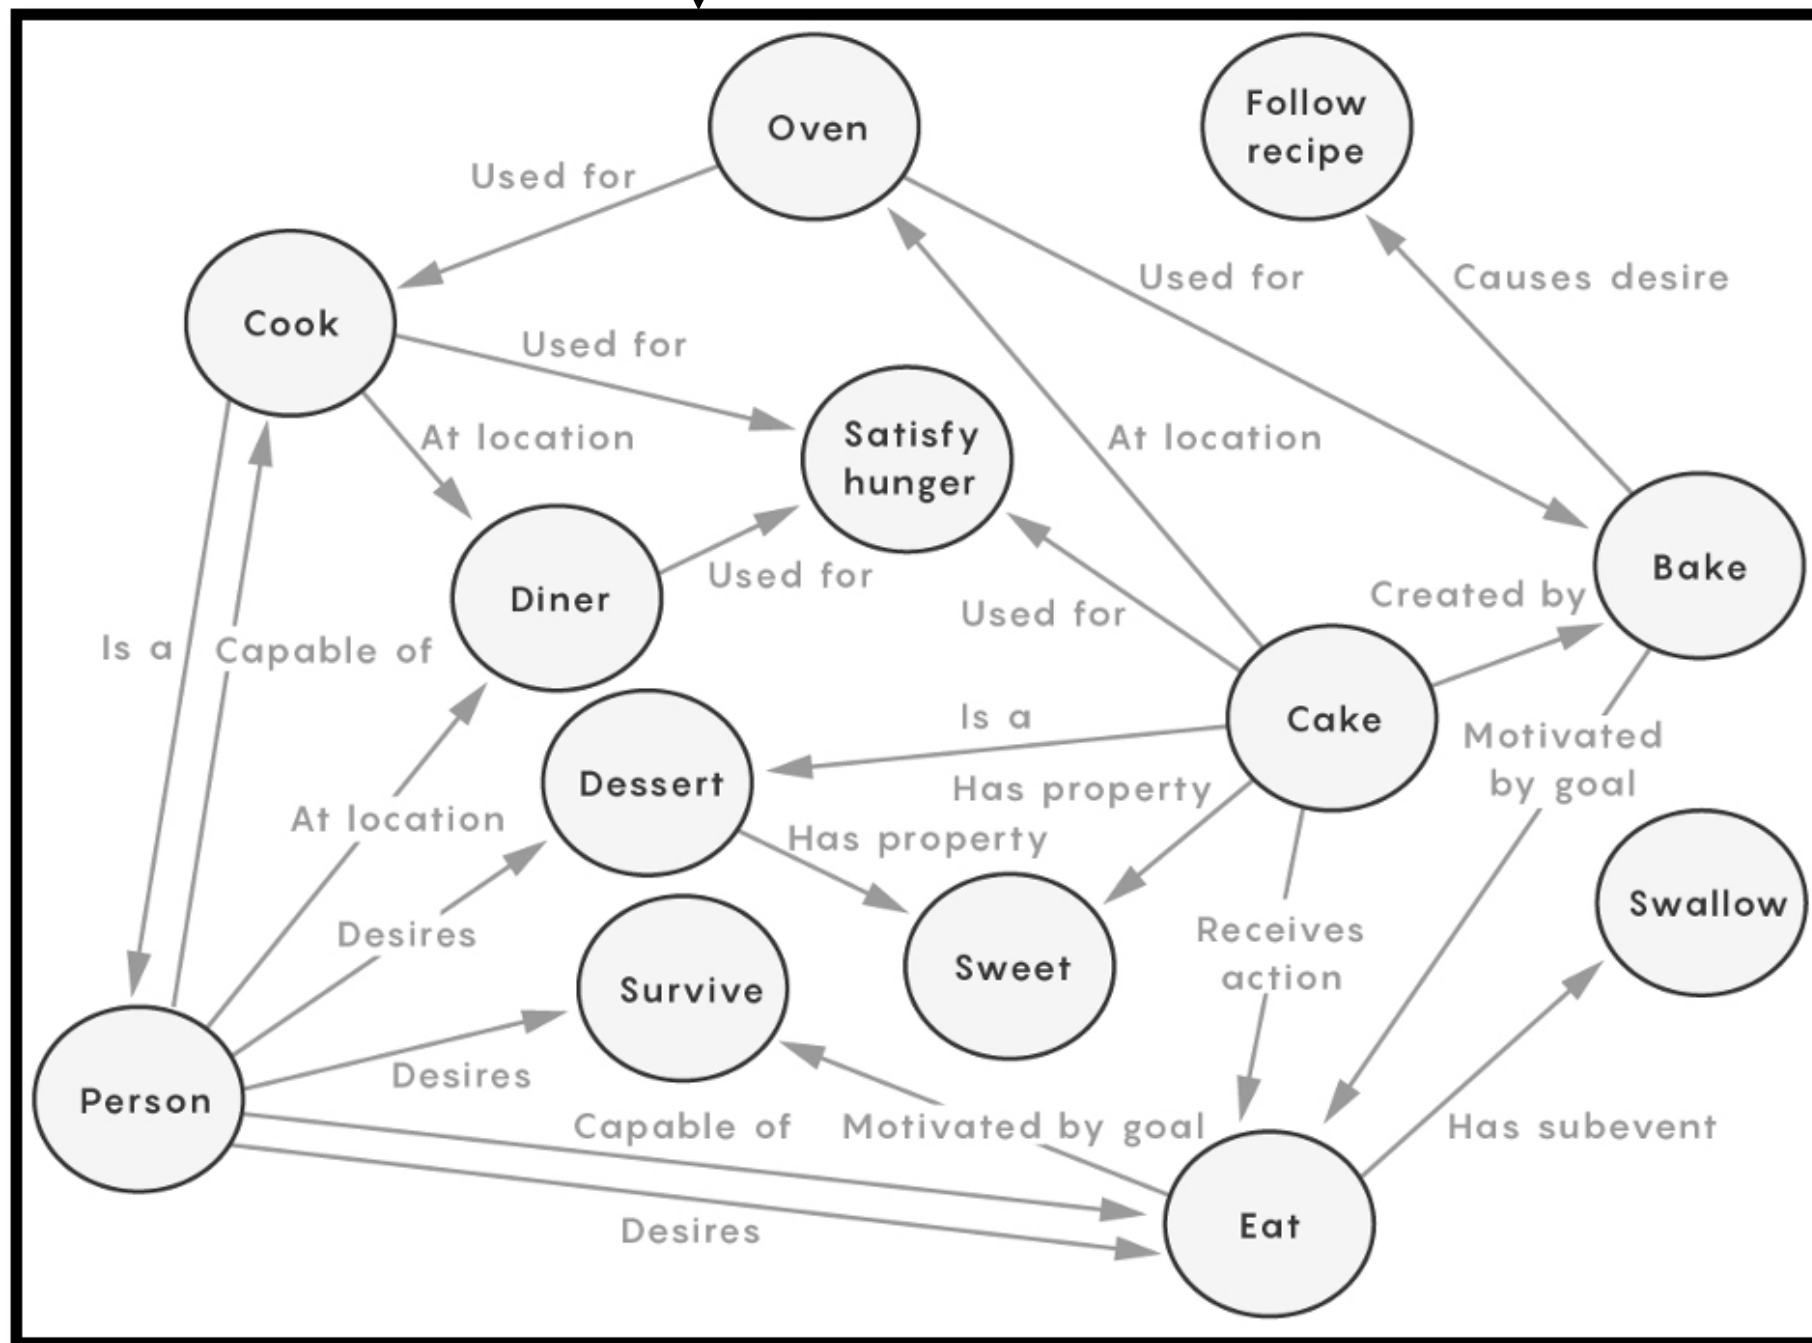


🤖 Parametric Knowledge



OpenAI GPTs

Structured Knowledge

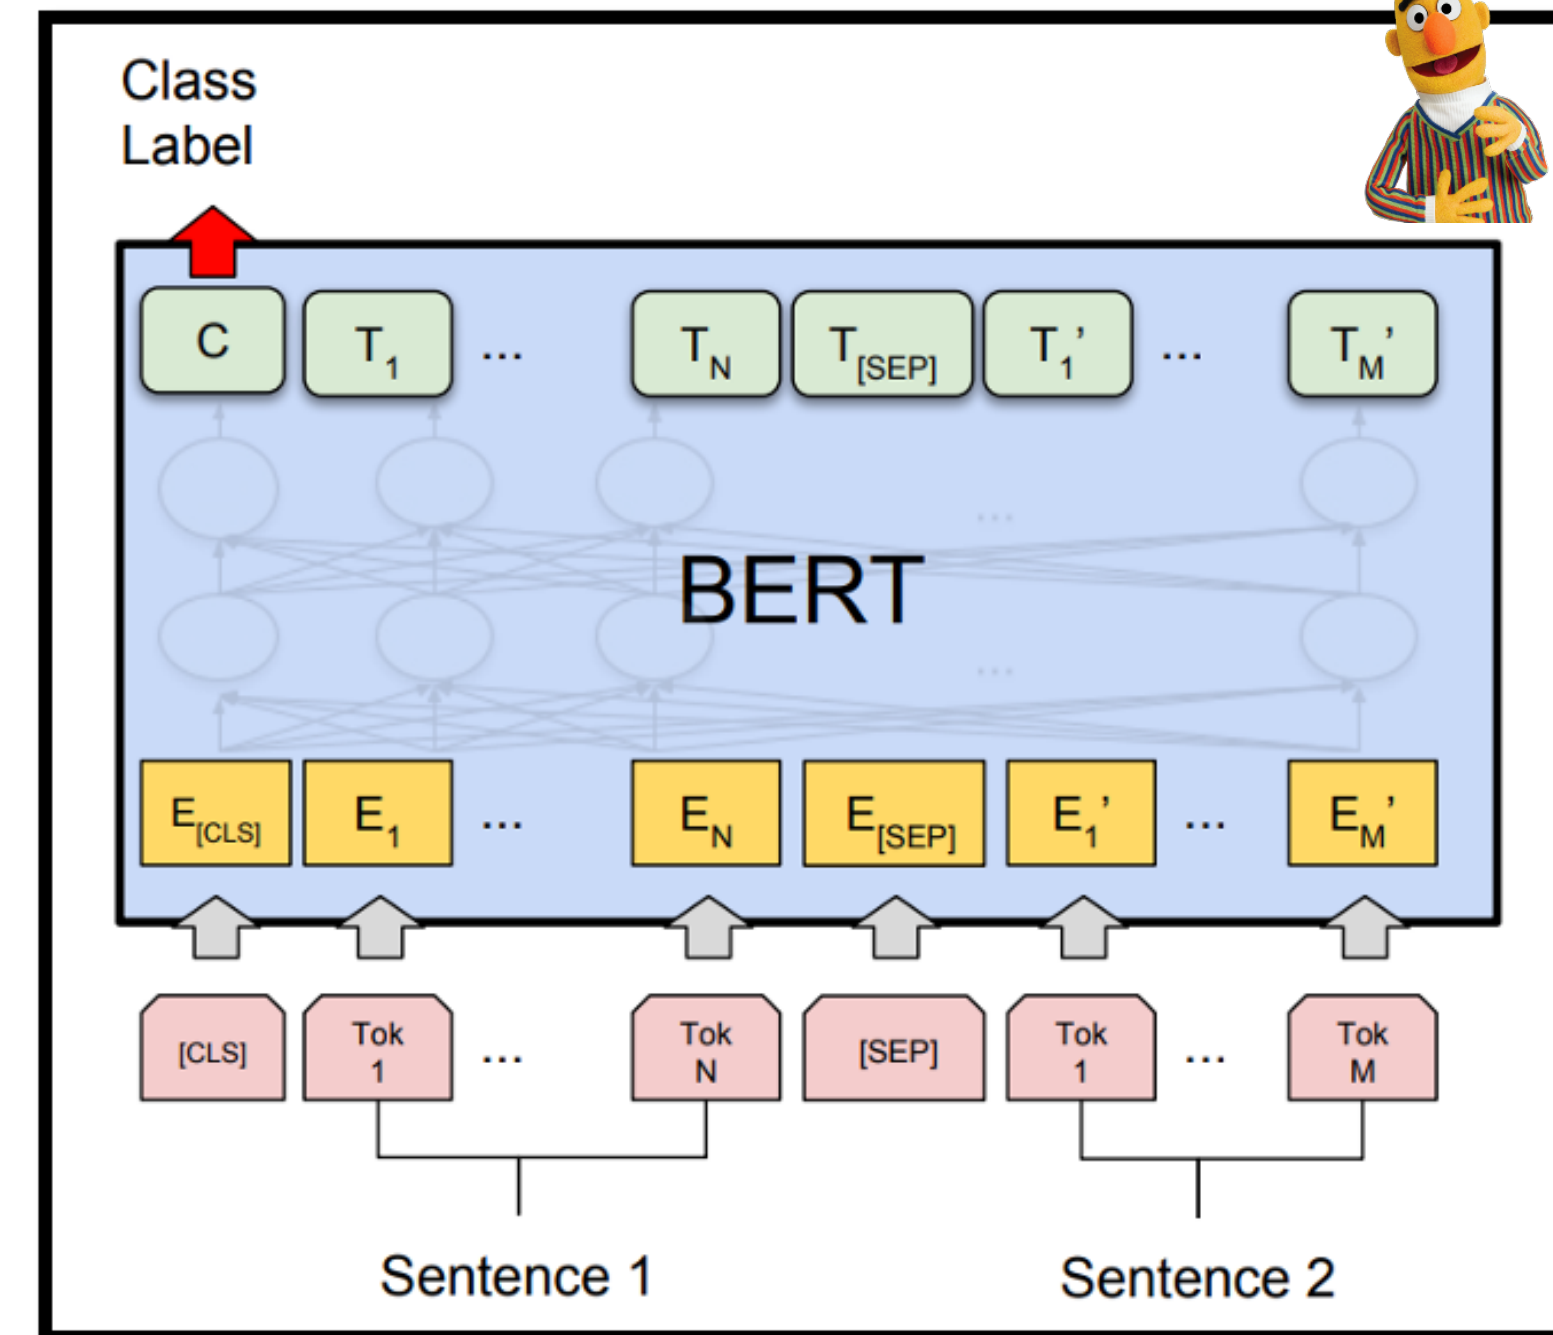


Symbolic Structures of Knowledge


Task: Multiple-Choice QA

Question: xxxxx?
Options: A) a_1, B) a_2, C) a_3

Am I too old?



Neural Language Models (e.g., BERT [9])



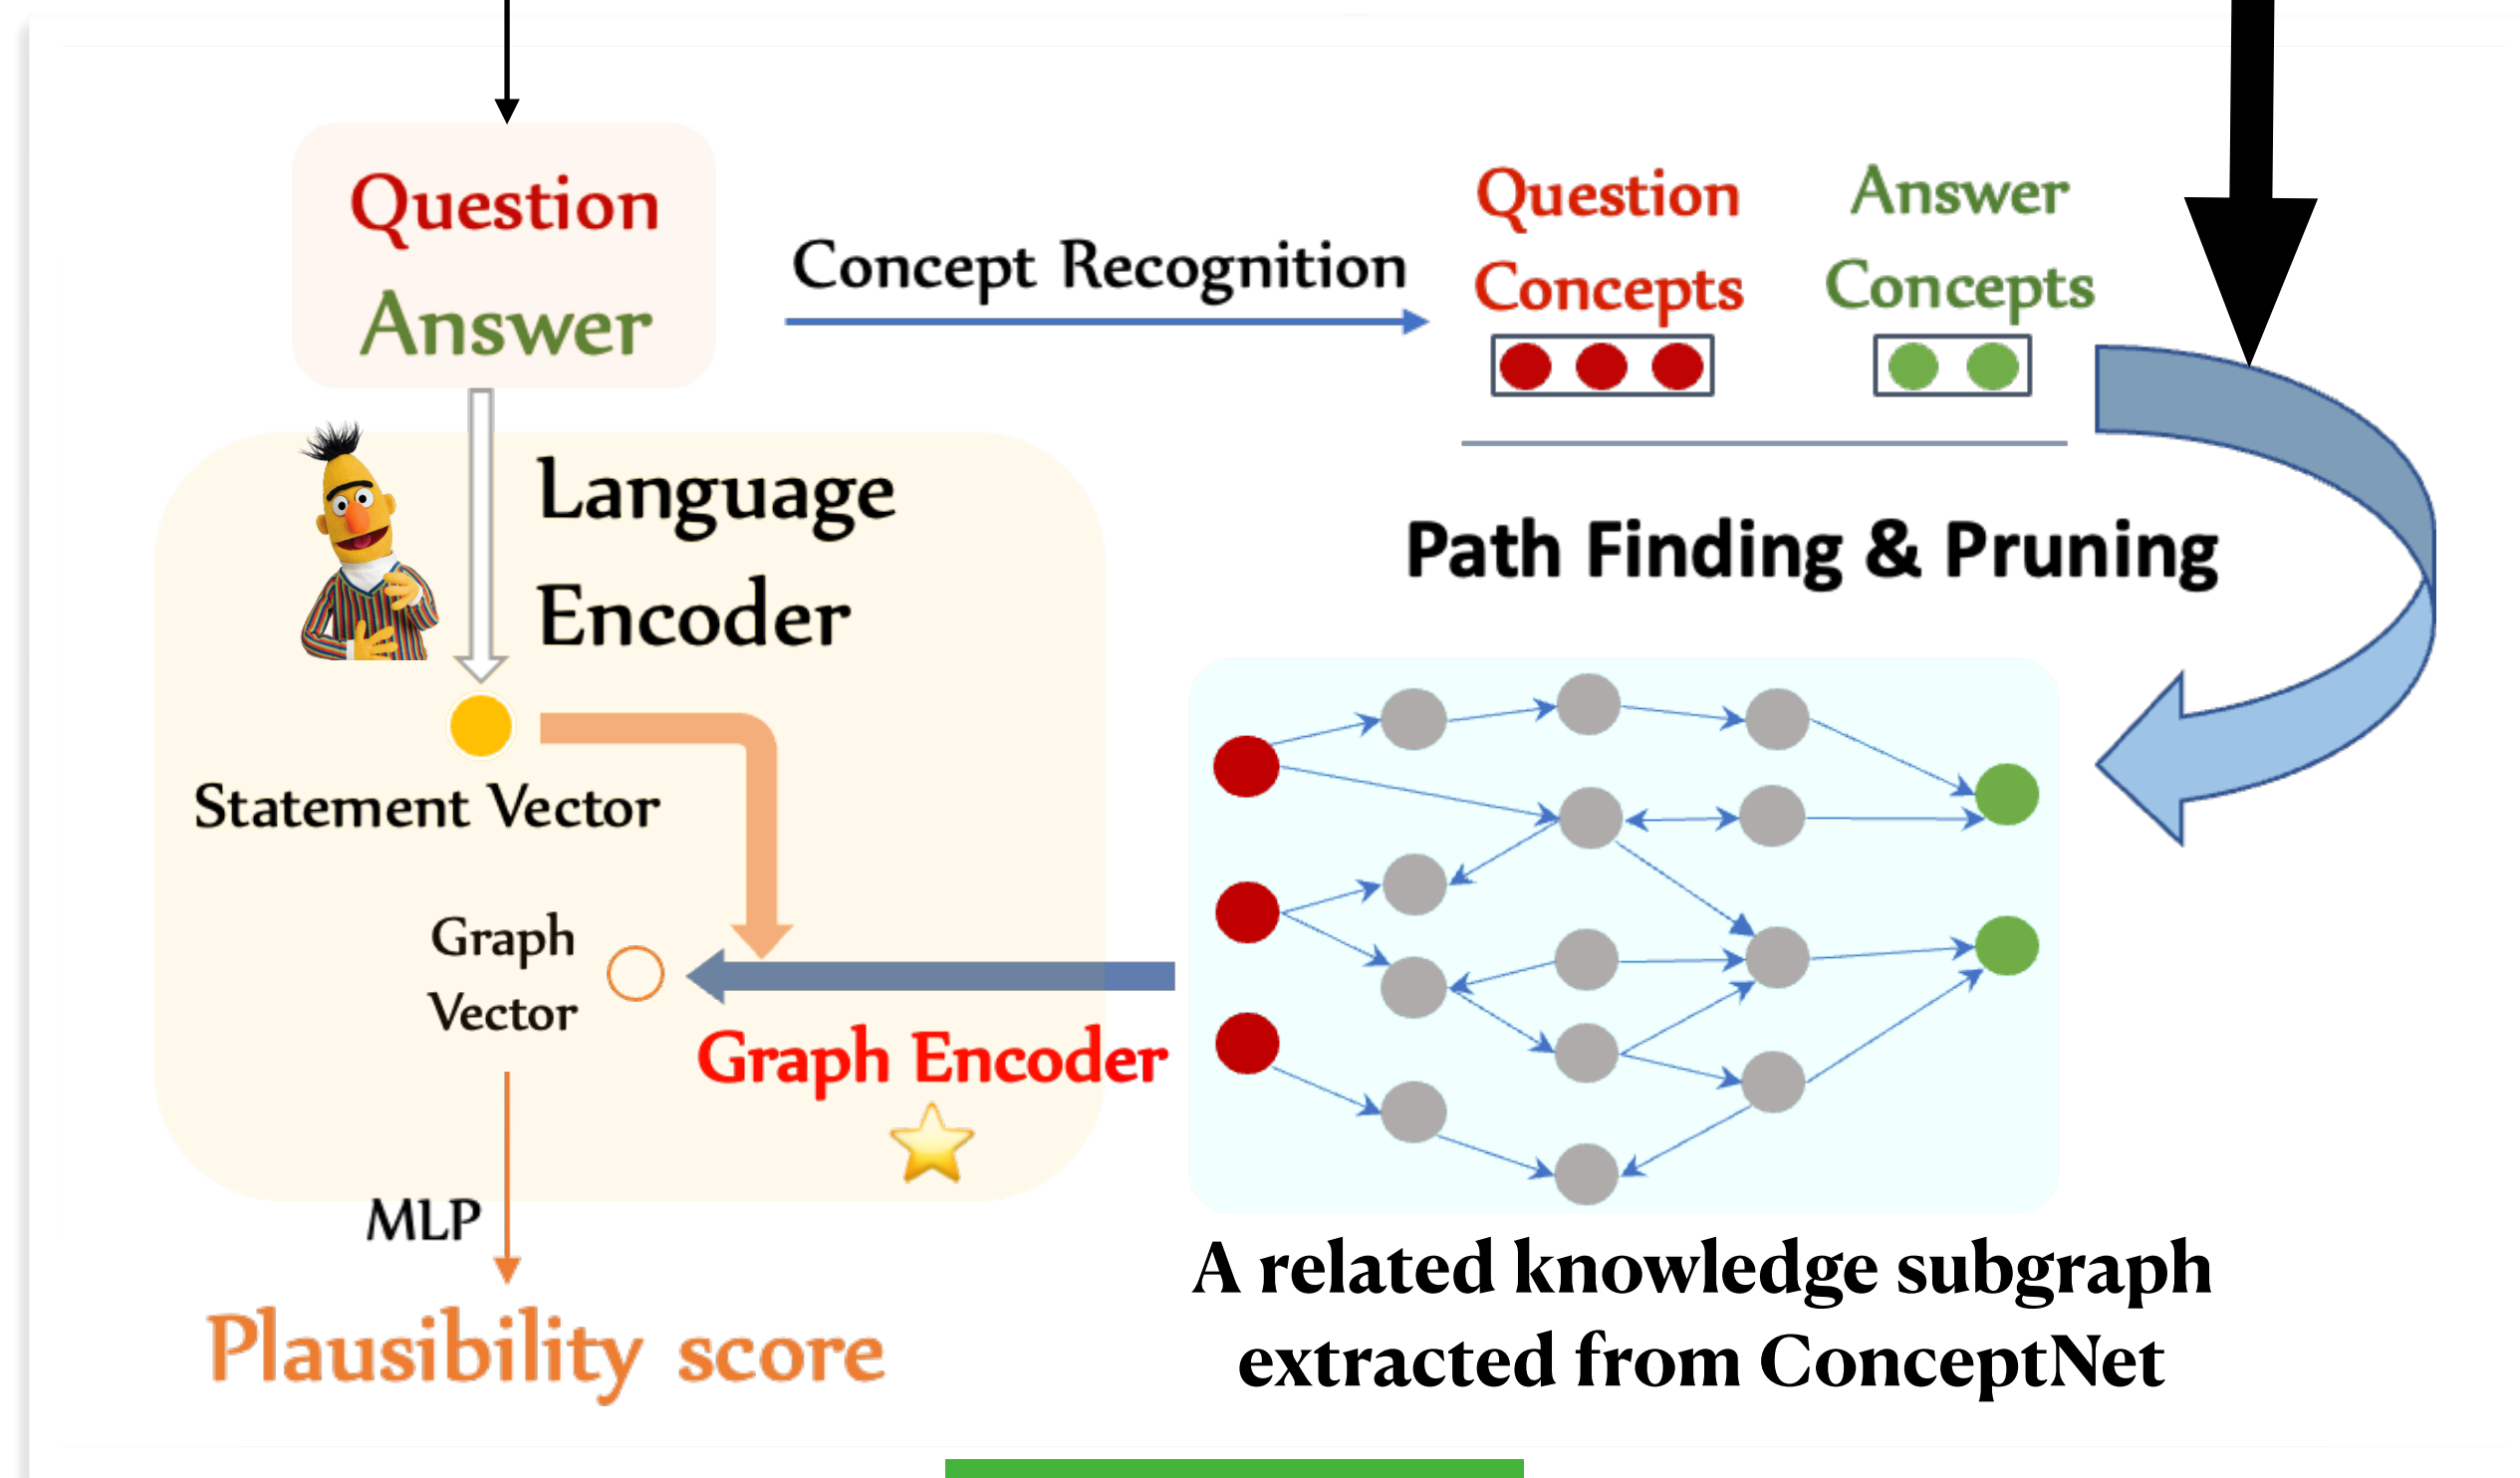
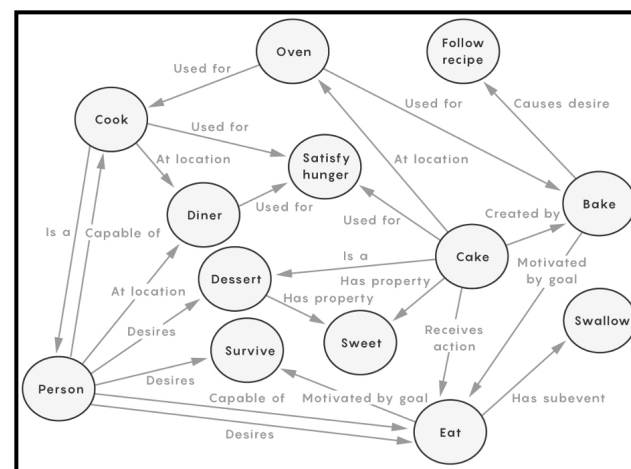
How can we incorporate structured knowledge into neural language models?

Multiple-Choice QA

Question: xxxxx?

Candidates: A) a_1, B) a_2, C) a_3

Knowledge Graph



KagNet [10]

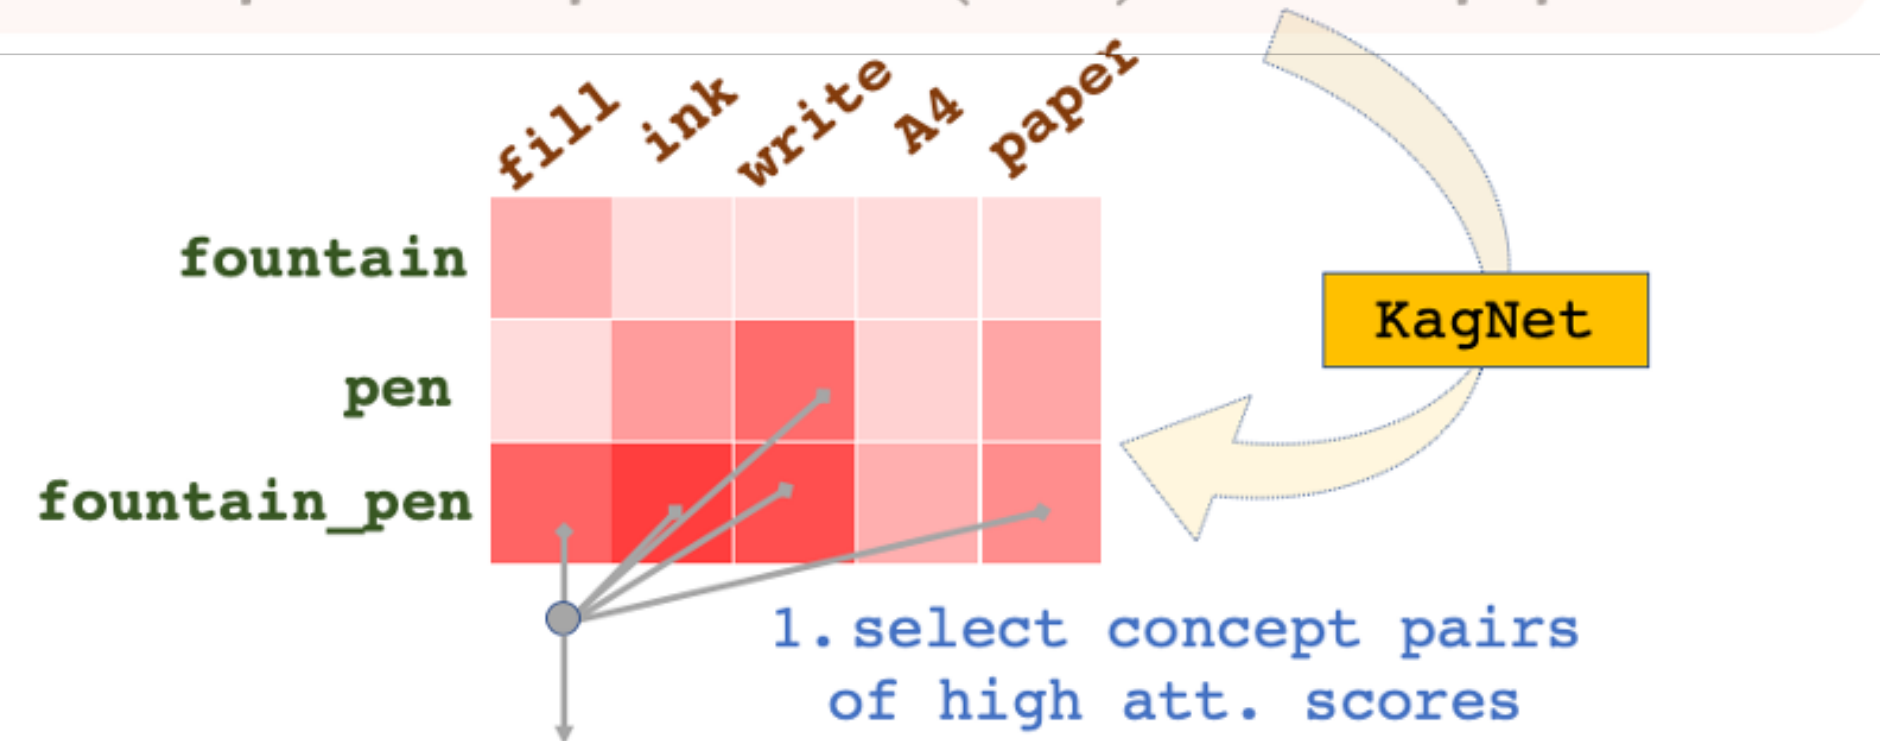
KagNet: Knowledge-Aware Graph Networks for Commonsense Reasoning

Features

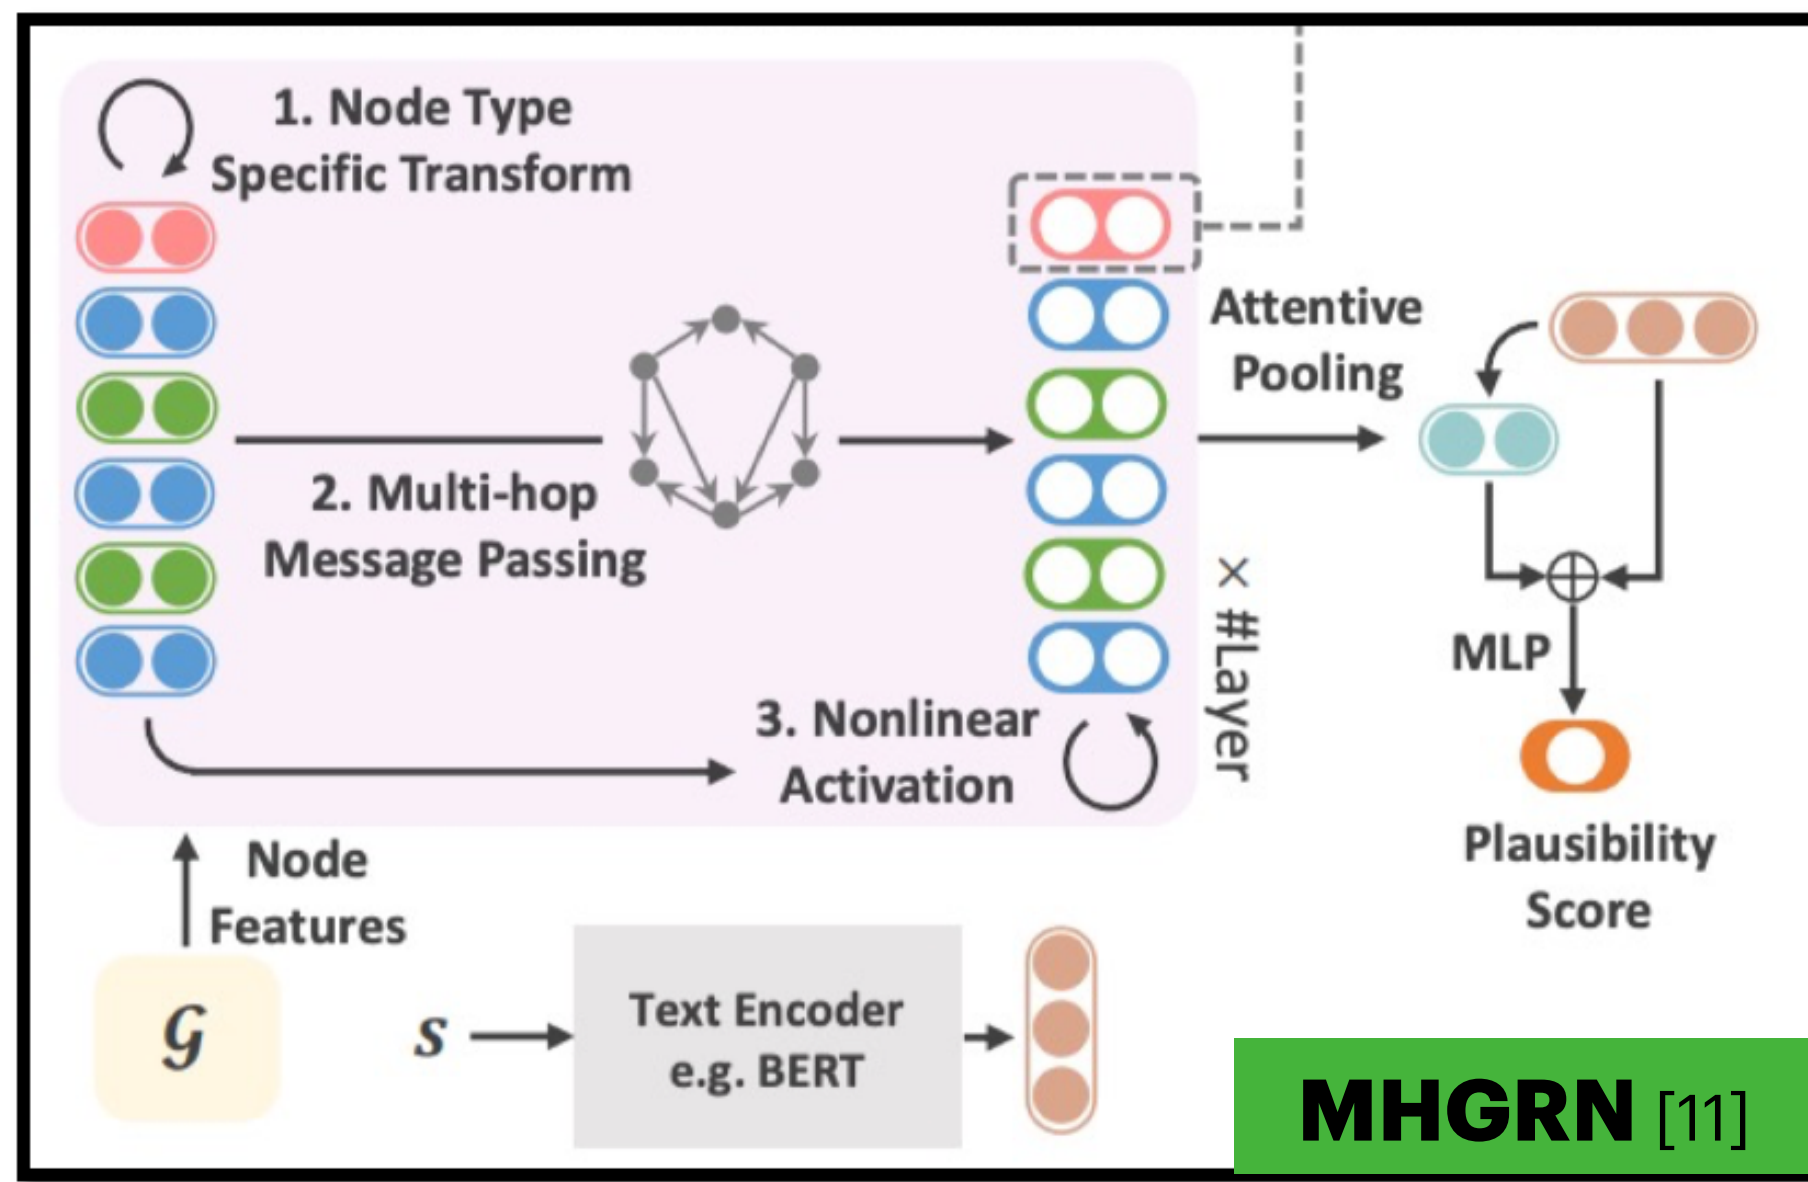
- LM encoder to encode the text of question + candidate
- Graph encoder to encode knowledge subgraphs
- Late fusion of Text Embeddings + Graph Embeddings
- Interpretable ↓

What do you **fill** with **ink** to write on an **A4** paper?

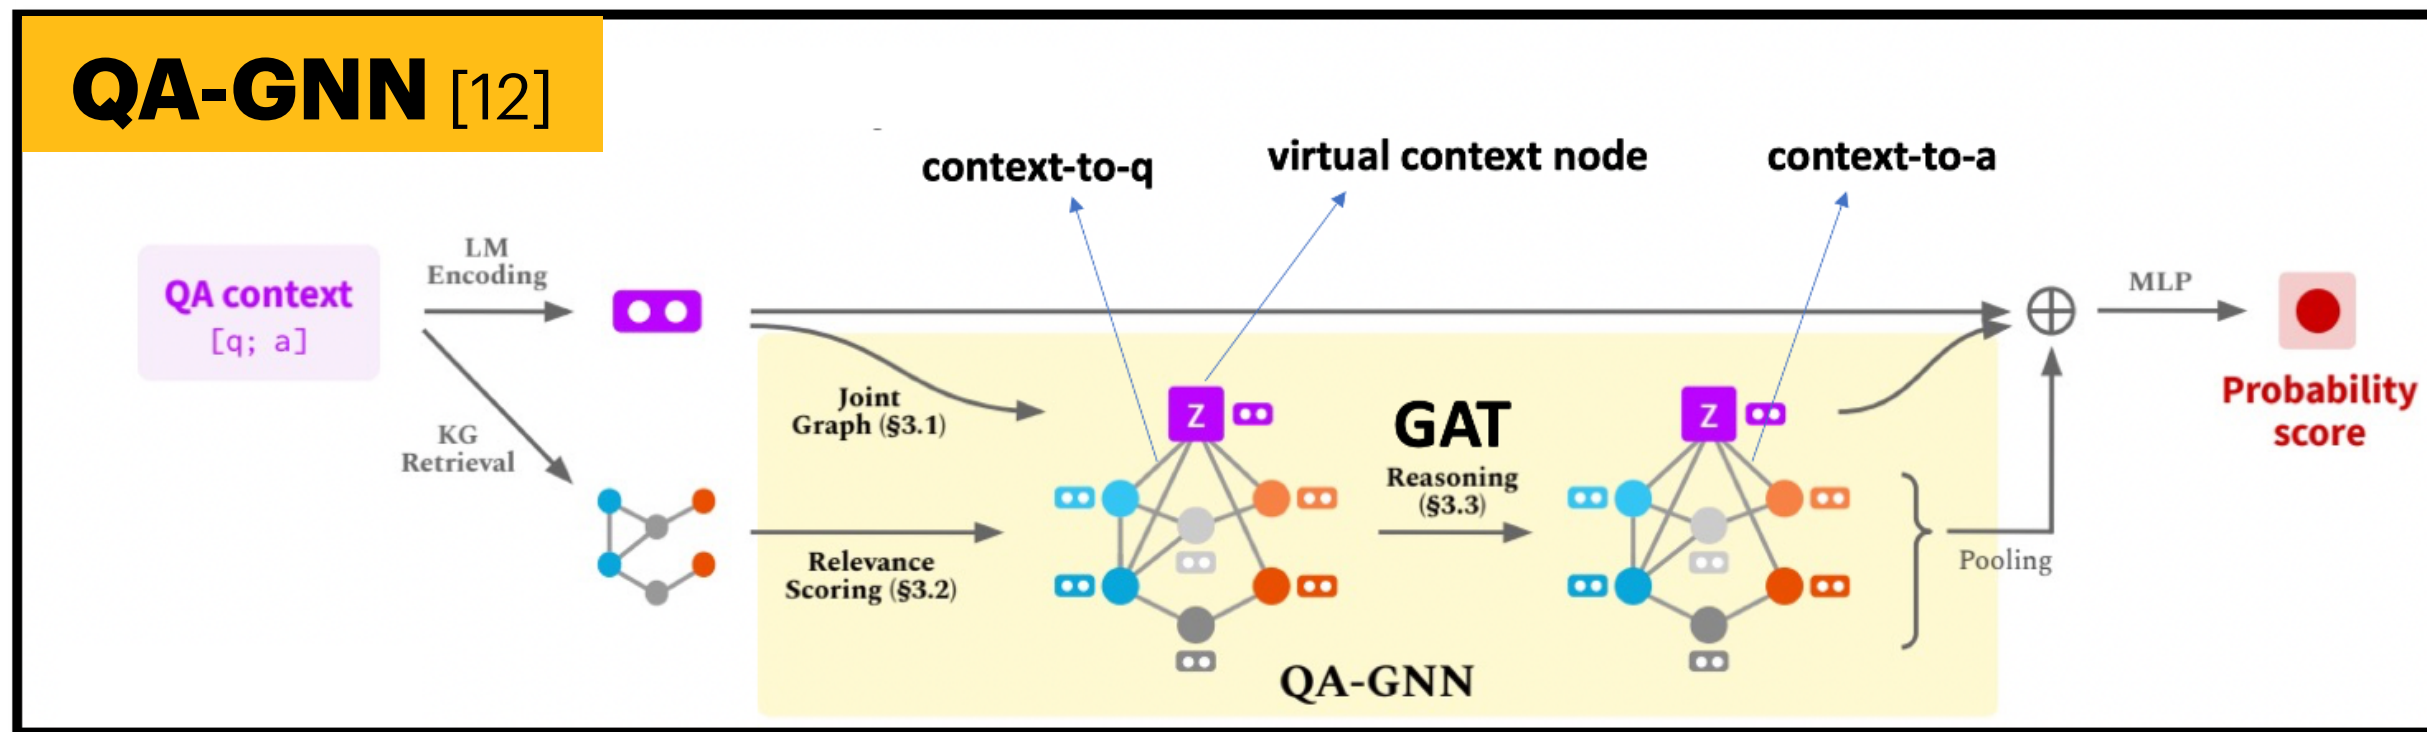
A: fountain pen ✓ (KagNet); B: printer (BERT);
C: squid D: pencil case (GPT); E: newspaper



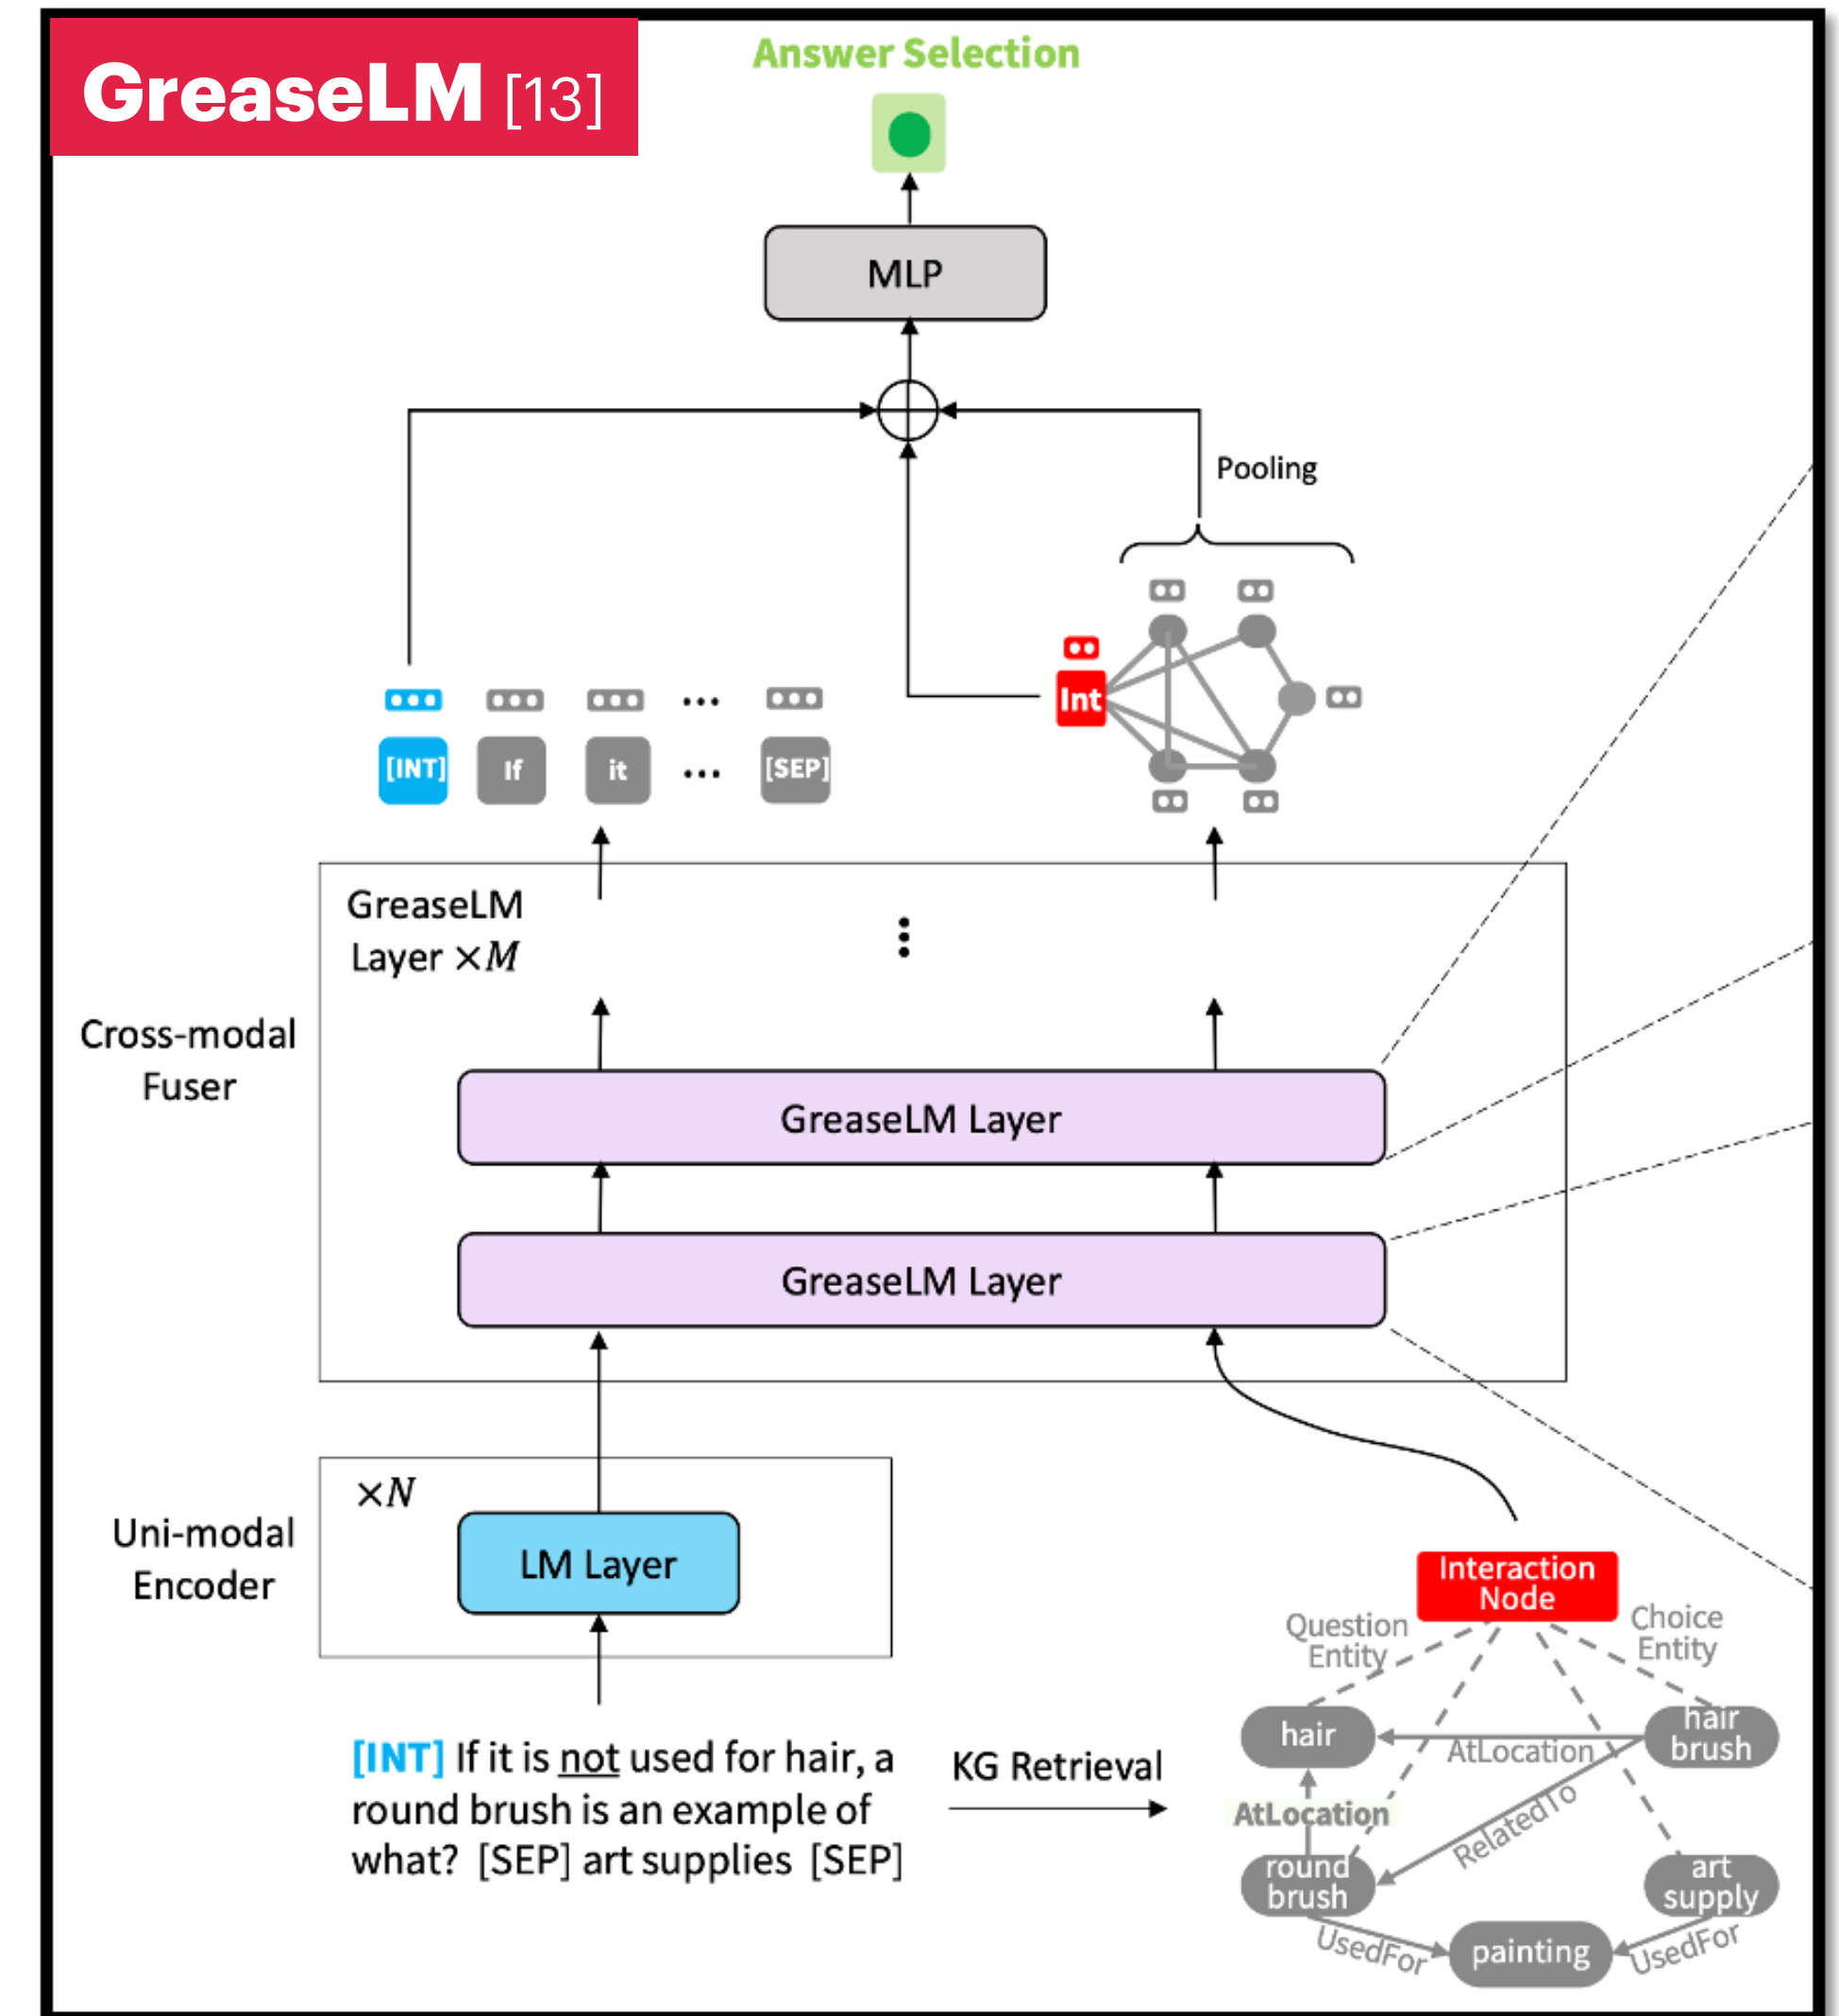
```
ink -PartOf-> fountain_pen
ink -RelatedTo-> container <-IsA- fountain_pen
fill <-HasSubEvent- ink <-AtLocation- fountain_pen
fill -RelatedTo-> container <-IsA- fountain_pen
write <-UsedFor- pen
write <-UsedFor- pen <-IsA- fountain_pen
paper <-RelatedTo- write <-UsedFor- fountain_pen
..... 2. Ranking via path-level attn.
```



○ Multi-hop graph encoder → Scalable

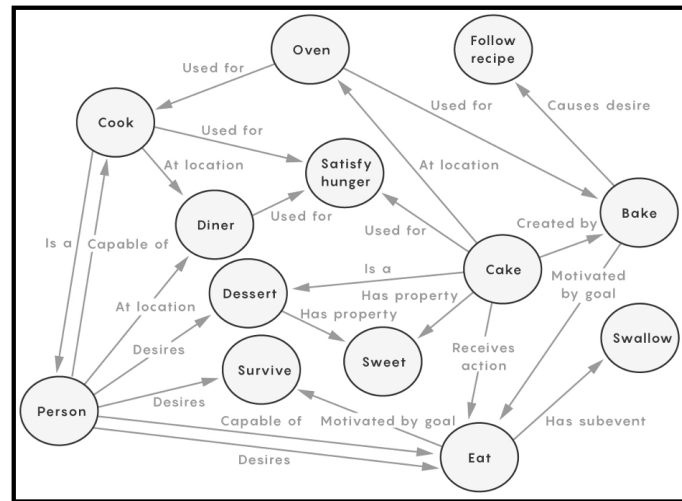


○ Text embeddings as virtual nodes in graph encoding



○ Early fusion of graphs and text encoders

Un/Semi-structured Knowledge




KGs can be limited:

1. Incompleteness
2. Only for binary relations

Tree

- Trees are *perennial plants* that have long woody trunks.
- Most trees add one new ring for each year of growth.
- Trees grow using *photosynthesis*, absorbing carbon dioxide and releasing oxygen.



...




Software system

From Wikipedia, the free encyclopedia

Not to be confused with System software.

A **software system** is a **system** of intercommunicating **components** based on **software** forming part of a **computer system** (a combination of **hardware** and software). It "consists of a number of separate **programs**, **configuration files**, which are used to set up these programs, **system documentation**, which describes the structure of the system, and **user documentation** which explains how to use the system".^[1]

The term "software system" should be distinguished from the terms **"computer program"** and **"software"**. The term computer program generally refers to a set of instructions **(source, or object code)** that perform a specific task. However, a software system generally refers to a more encompassing concept with many more components such as specification, **test results**, end-user documentation, maintenance records, etc.^[2]

Task: Open-Ended QA

Question: xxxxx?
(A target knowledge corpus.)

Complex questions need multiple reasoning steps.

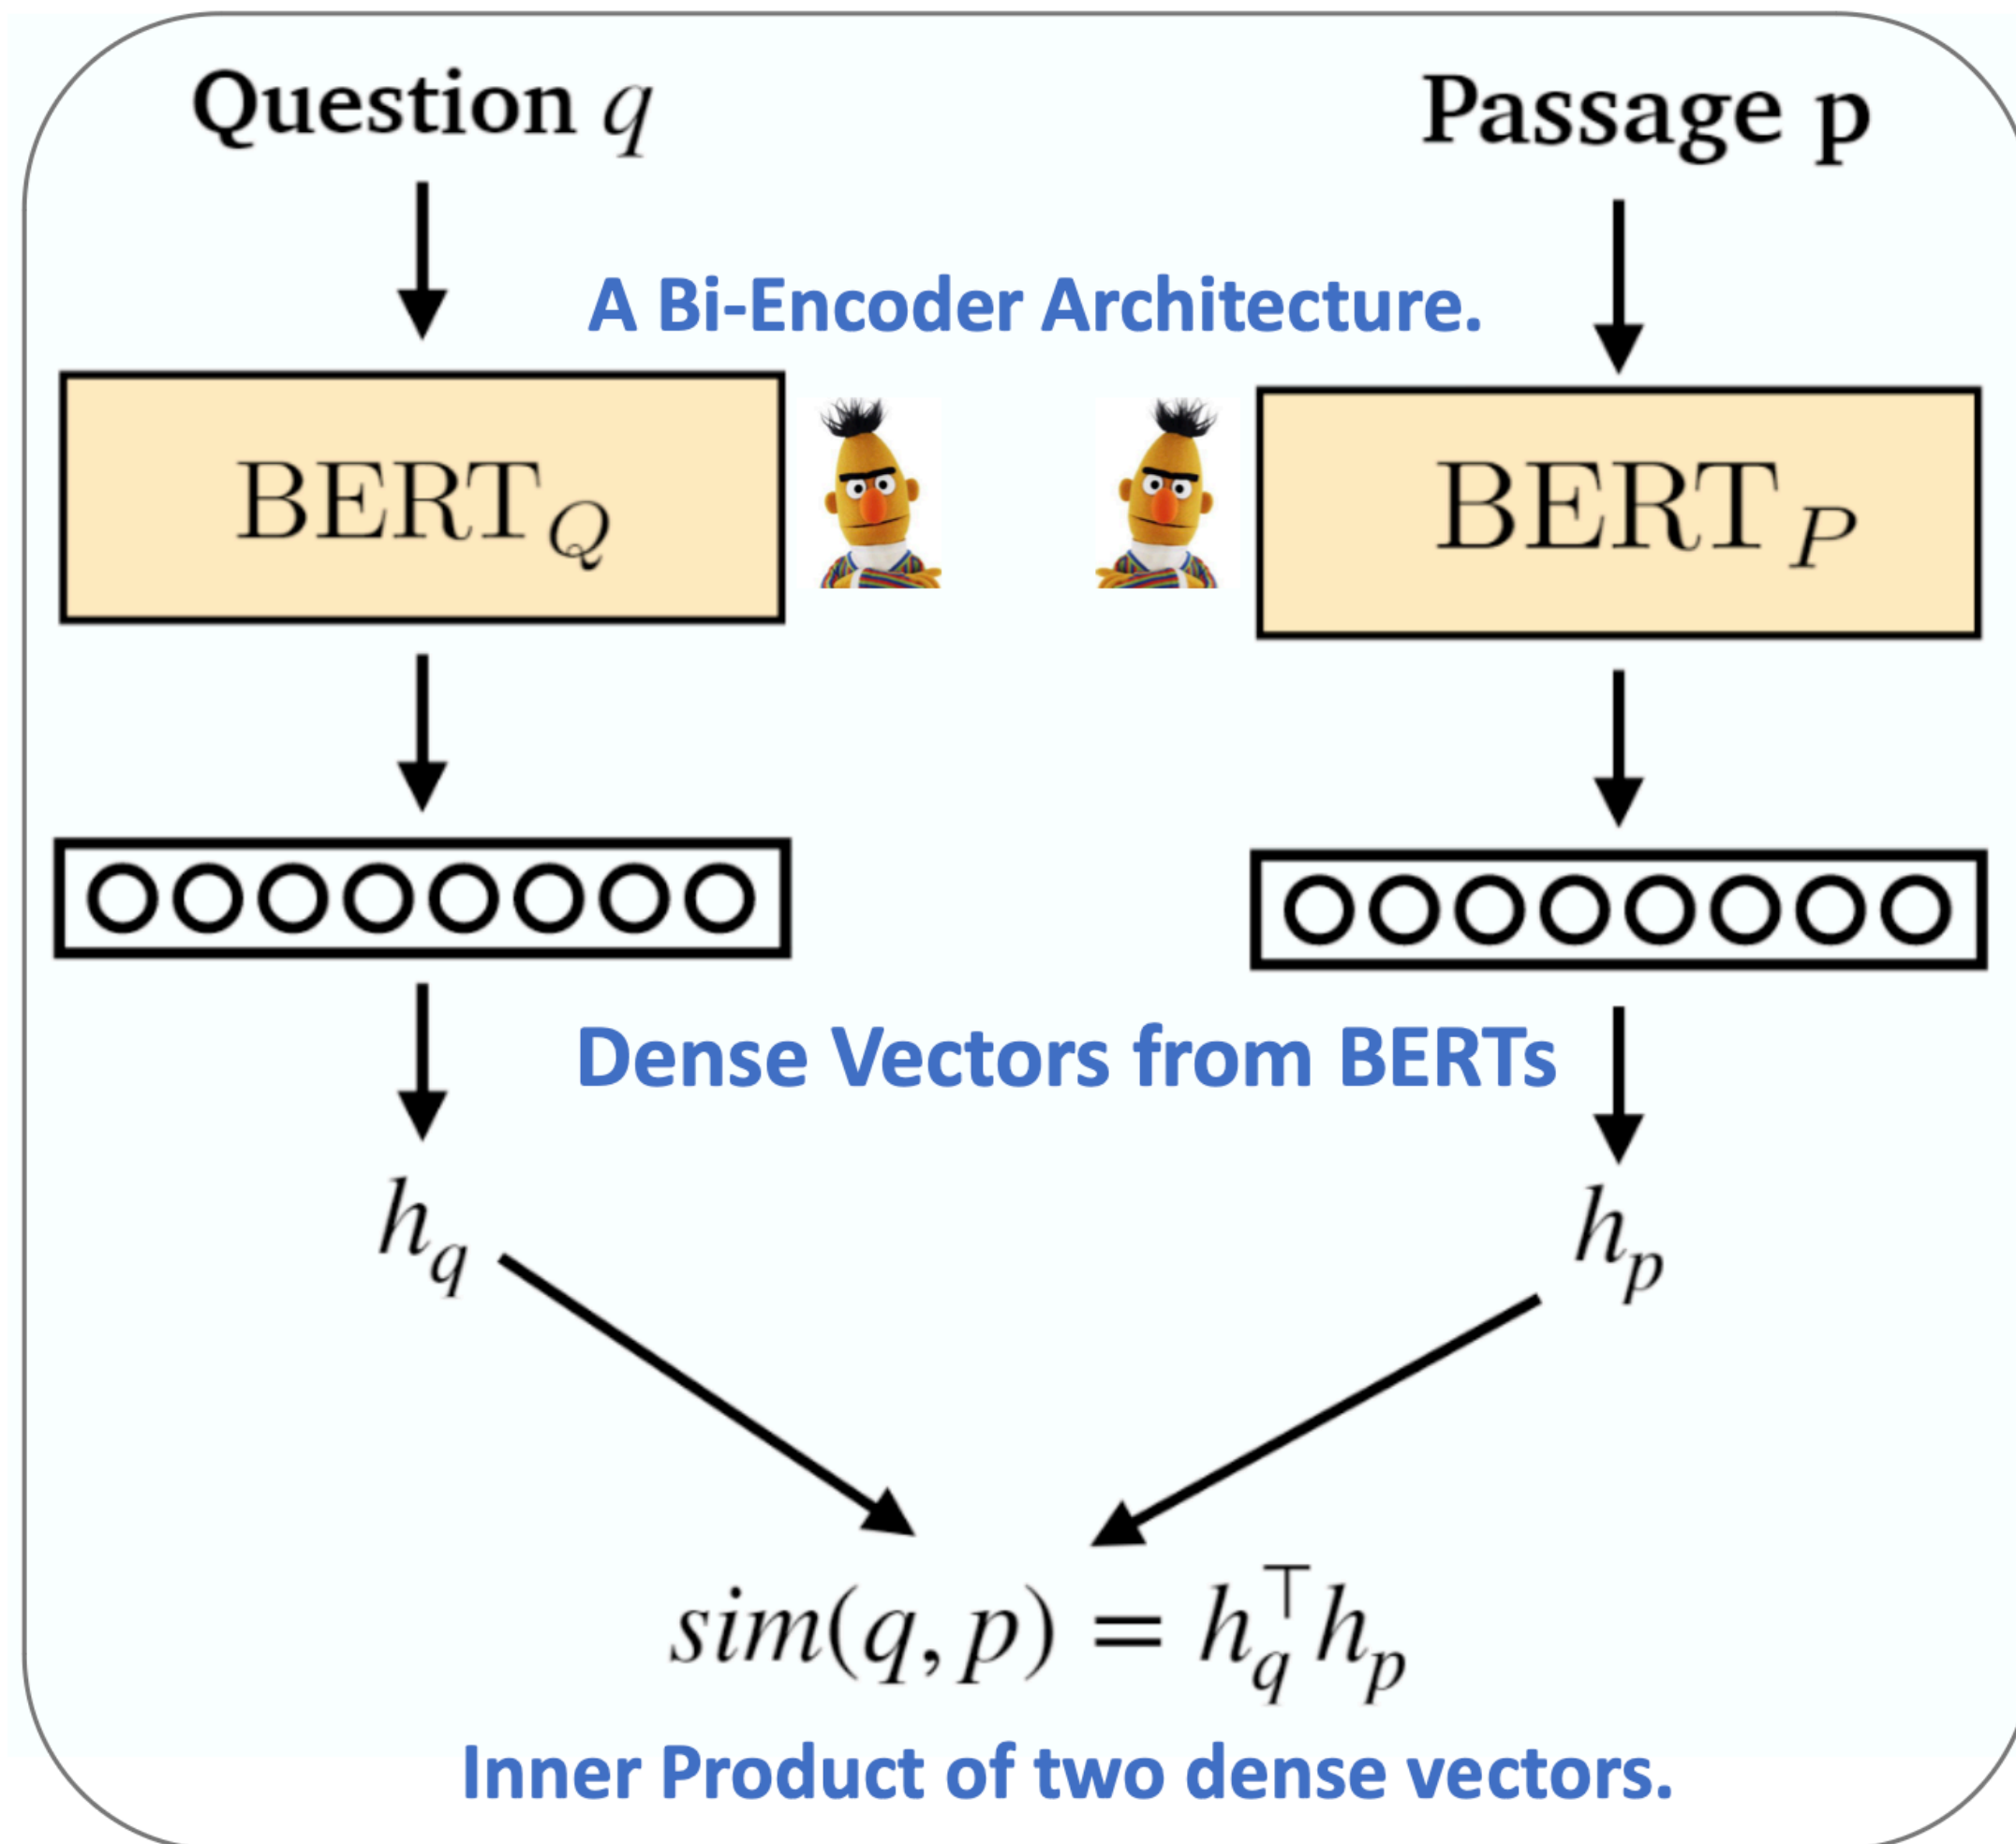
Who voices the dog in the TV show Family Guy ?

What can help alleviate global warming?

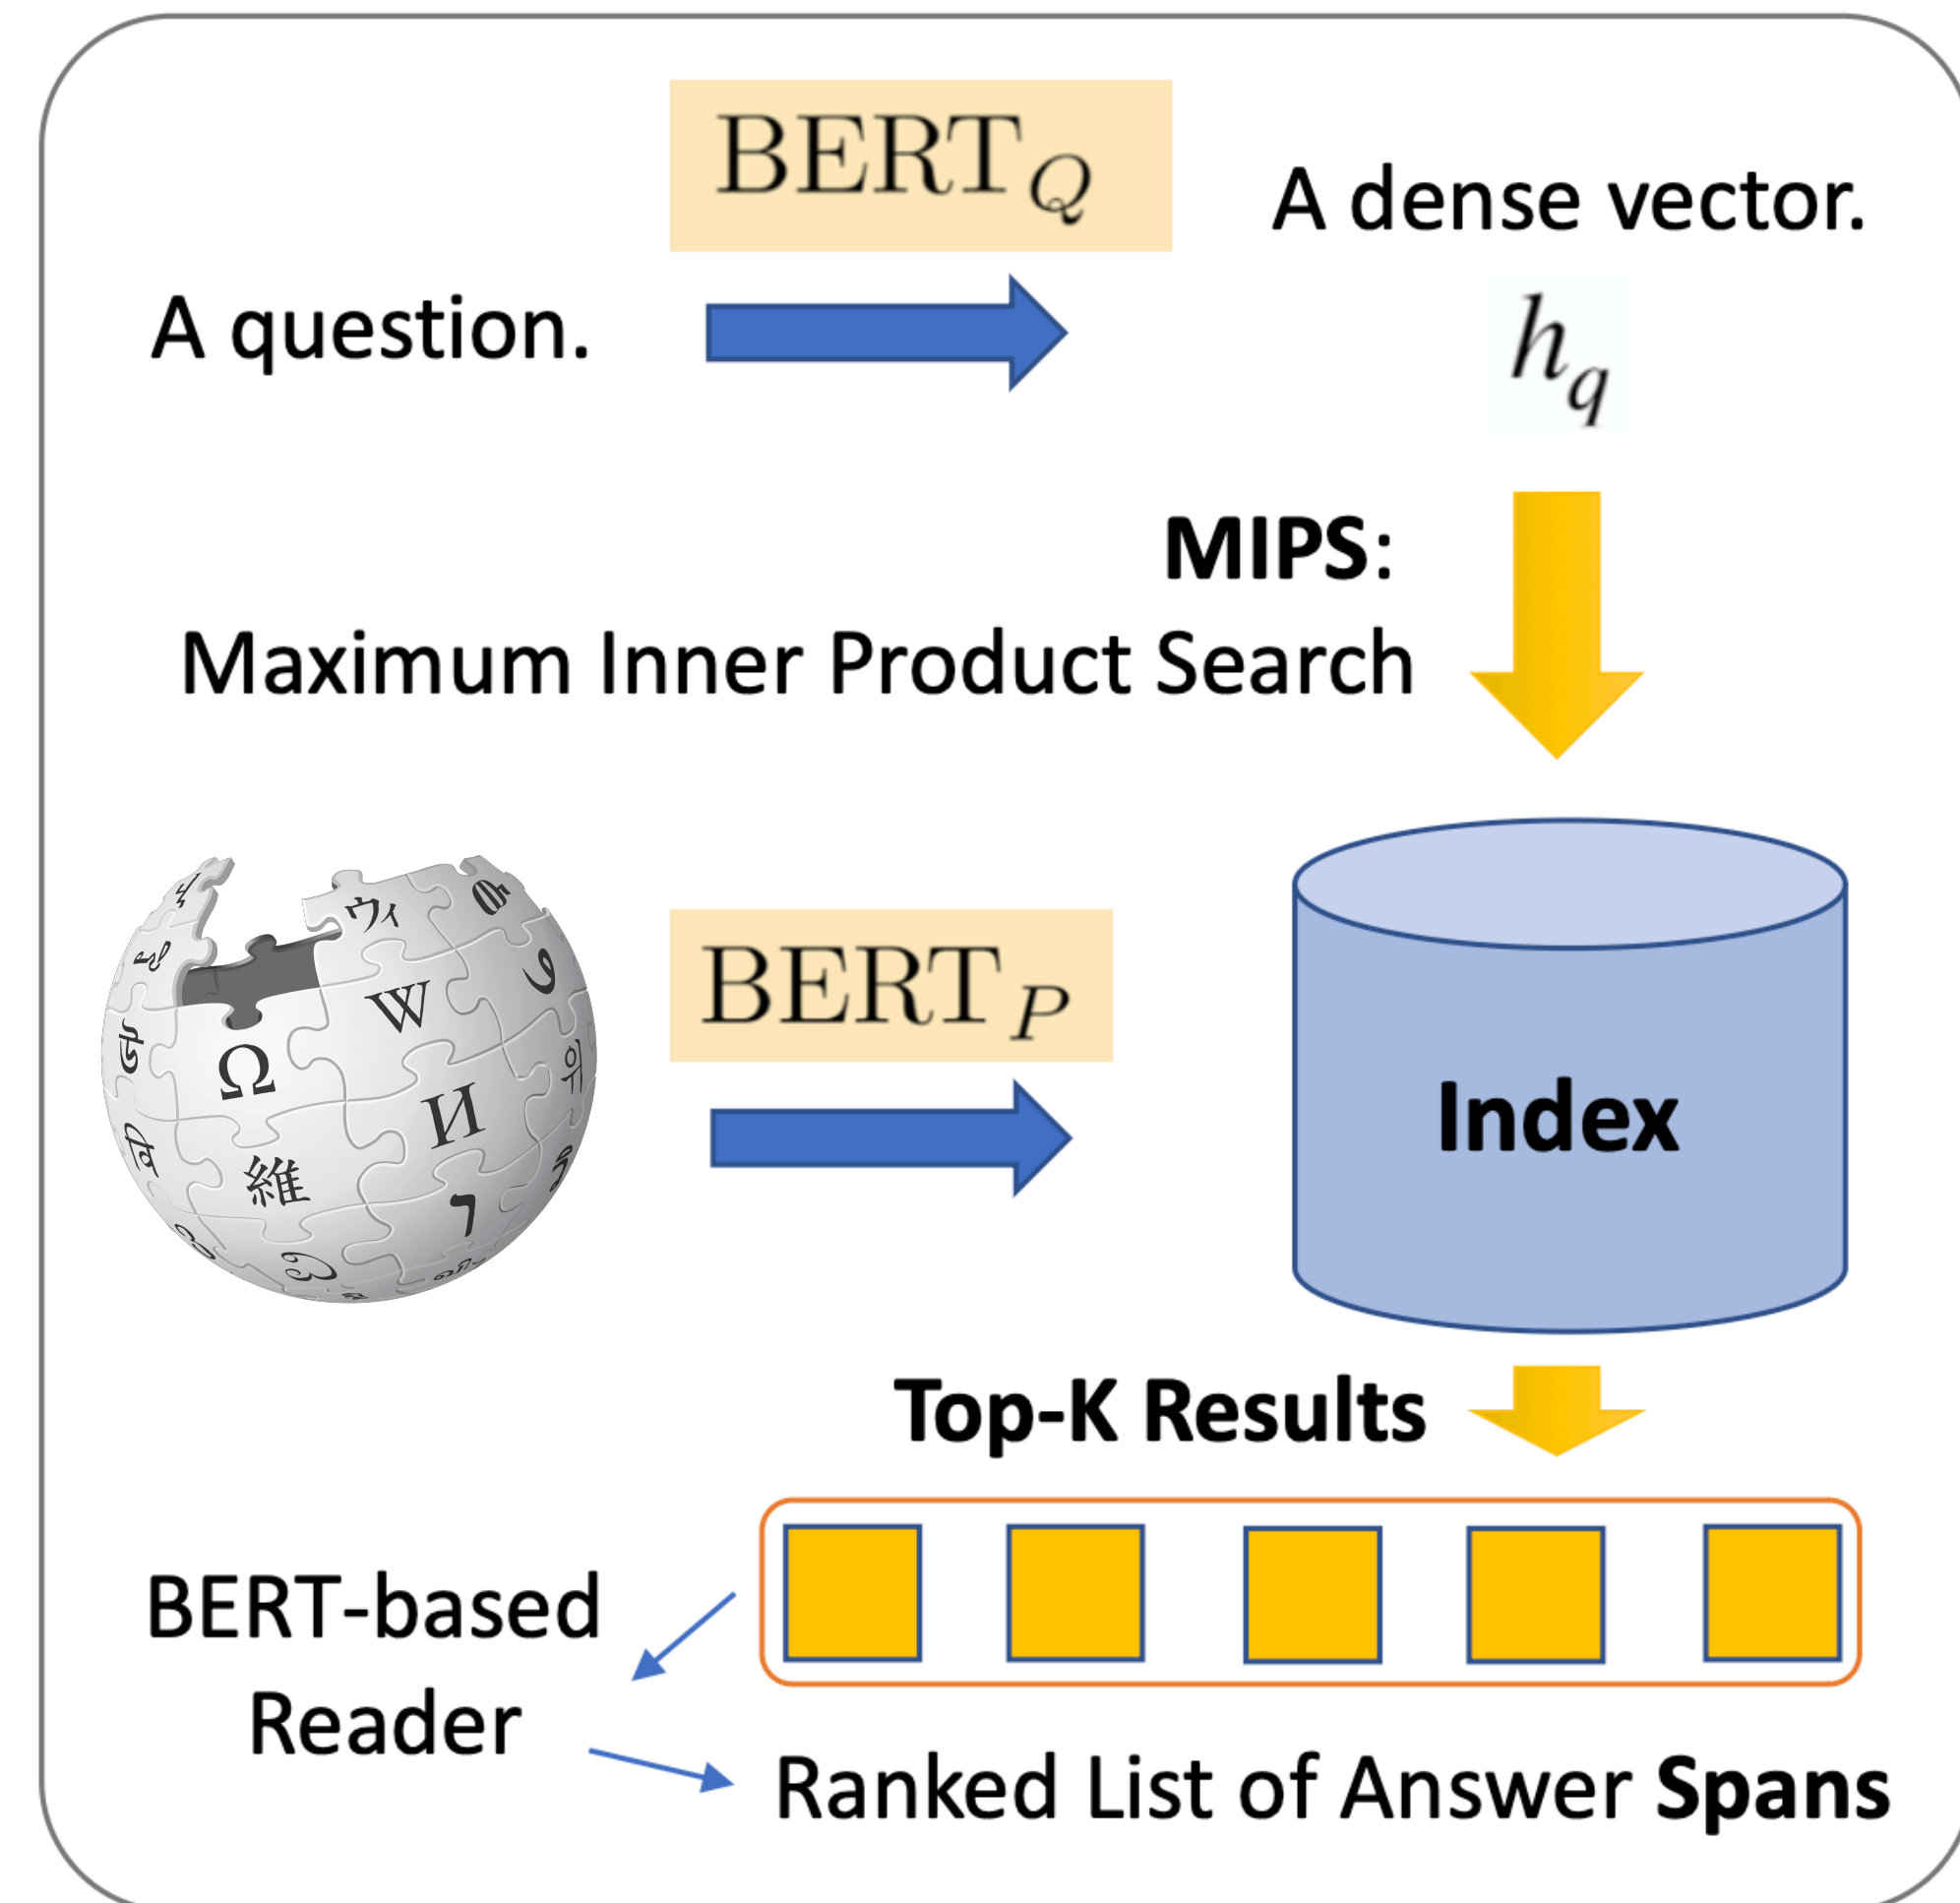
What will separate iron filings from sand?

Dense Passage Retrieval (DPR) [14]

A Trainable Method for Passage Retrieval



MIPS-based Inference Pipeline

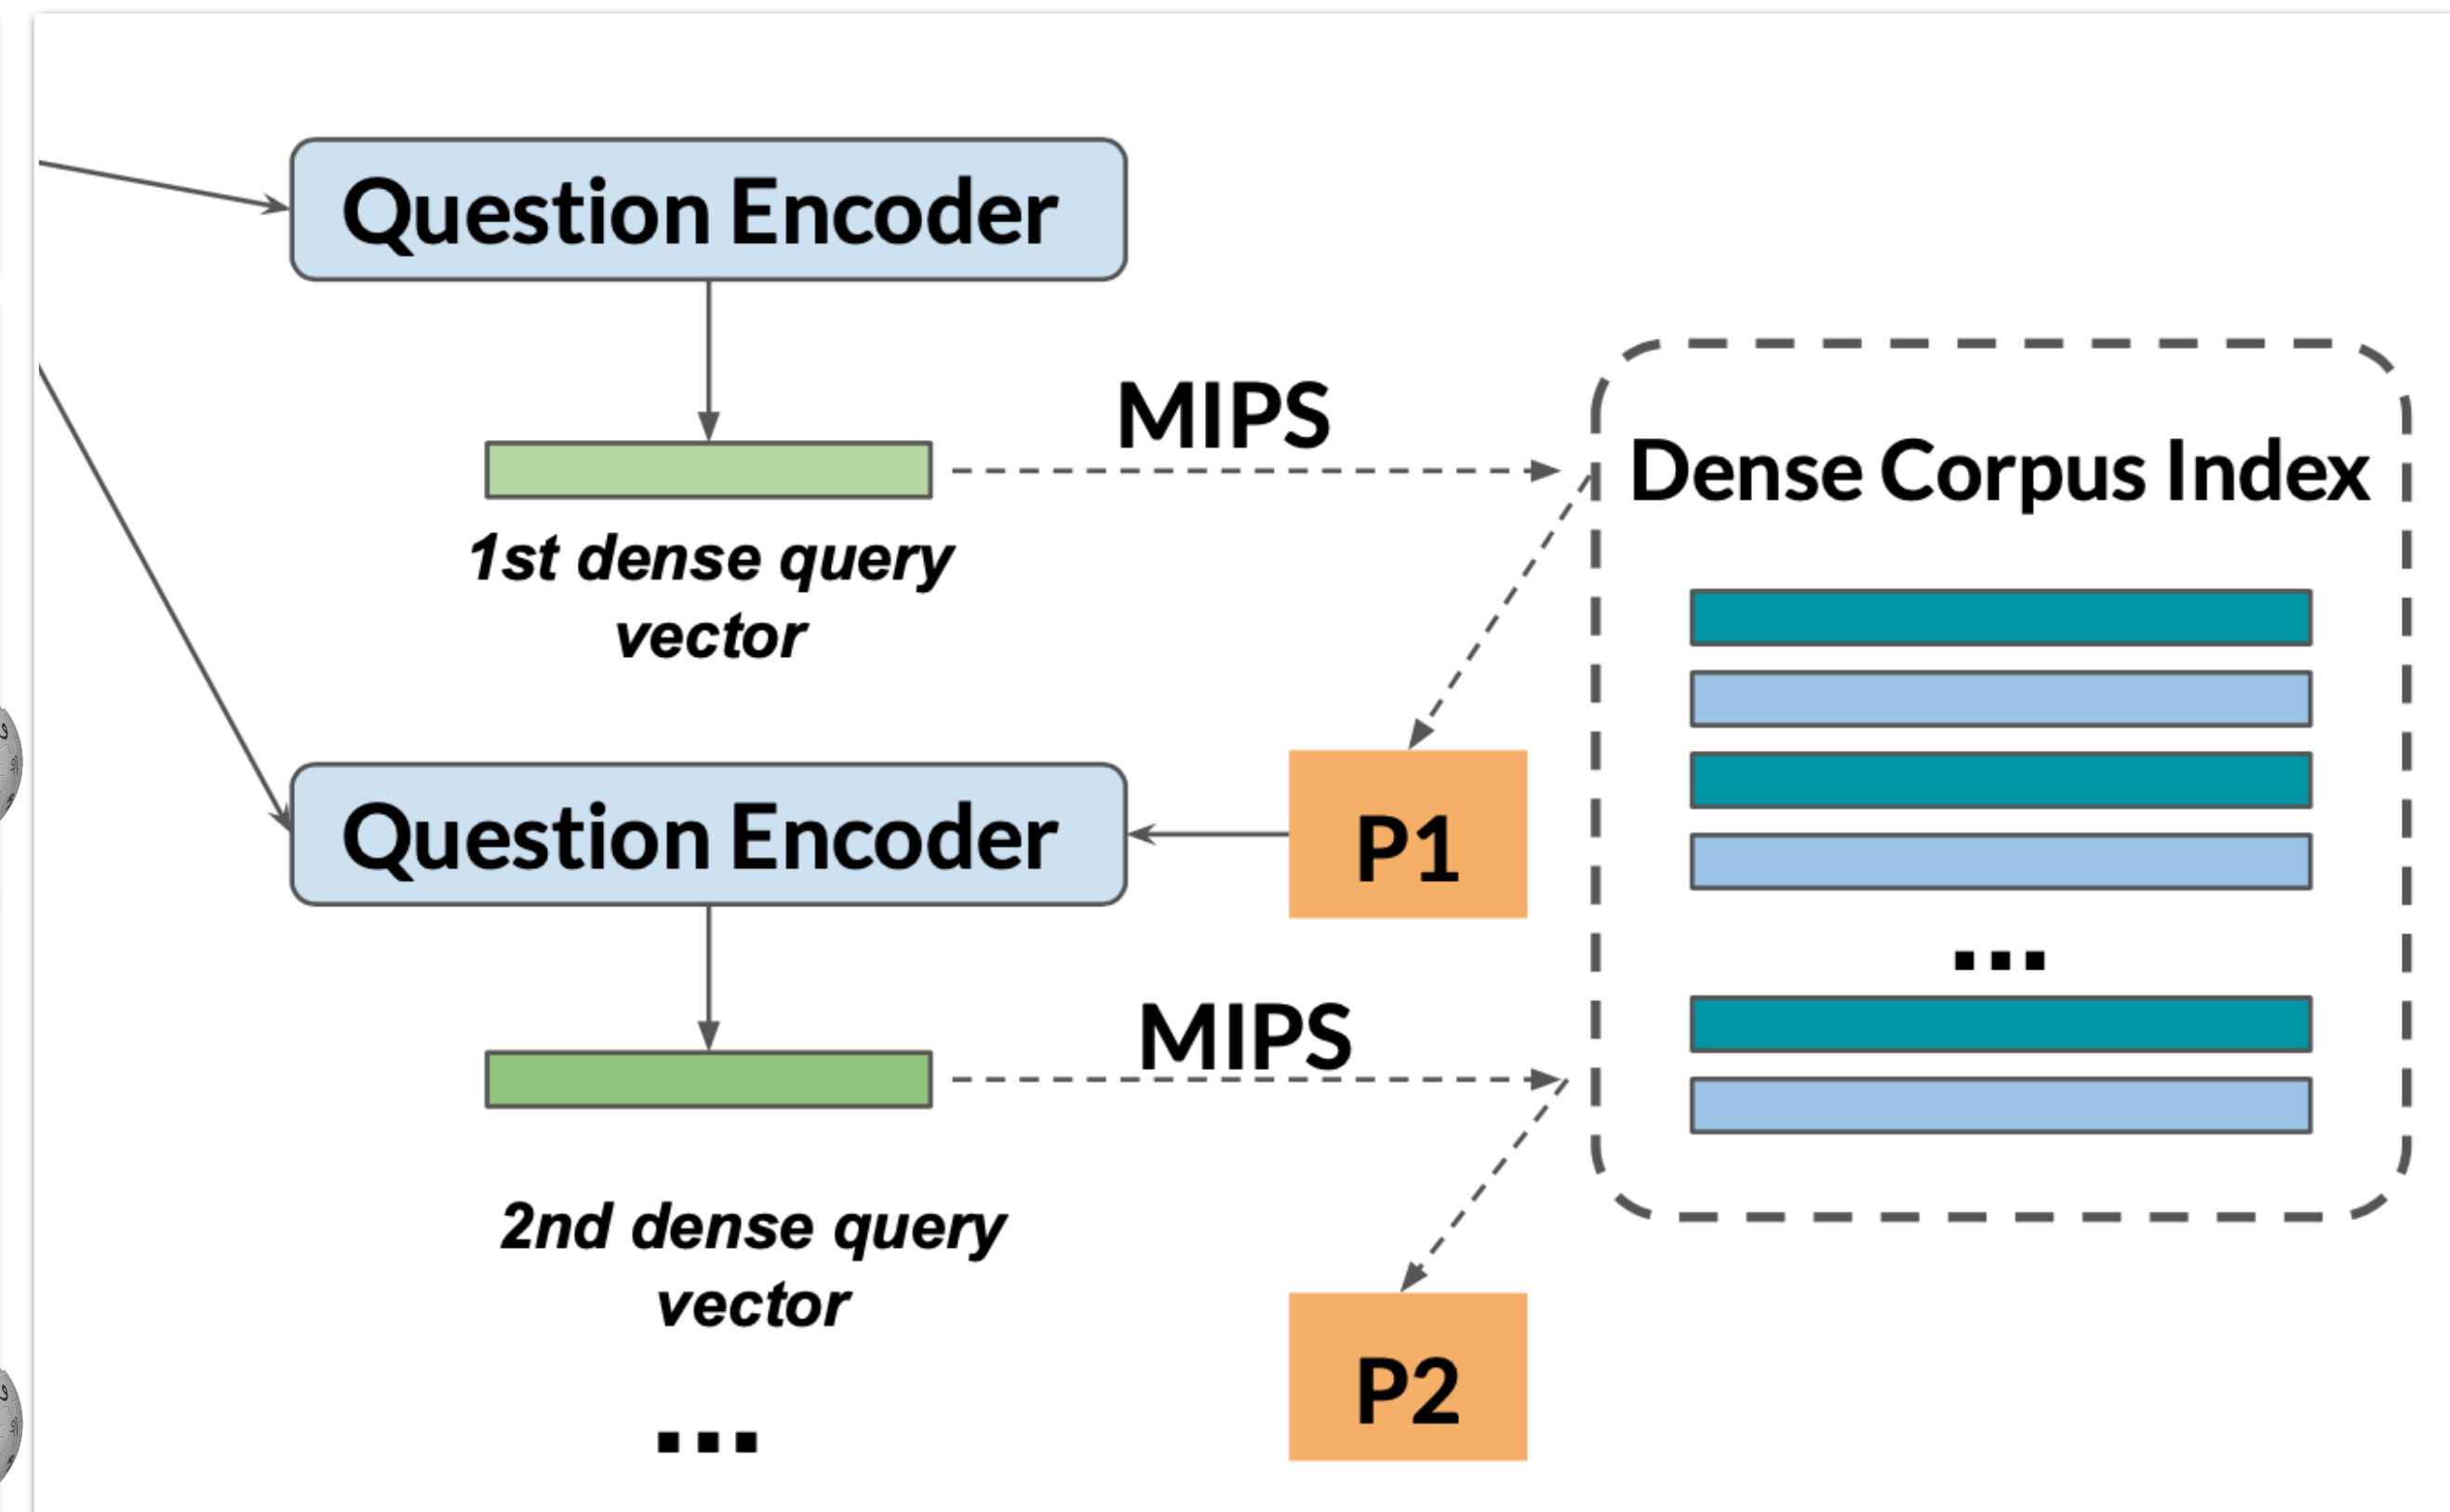


Multi-hop Dense Retrieval (MDR) [15]

Q: What was the nickname of **Judy Lewis**'s father ?

P1: Judy Lewis (born Judith Young; November 6, 1935 – November 25, 2011) was an American actress, writer, producer, and therapist. She was the secret biological daughter of actor **Clark Gable** and actress Loretta Young.

P2: William Clark Gable (February 1, 1901 – November 16, 1960) was an American film actor, often referred to as "**The King of Hollywood**". He had roles in more than 60 motion pictures in a wide variety of genres during a career that lasted 37 years...



Differentiable Fact-Following Operations (DrFact) [17]

Ai2GenericsKB

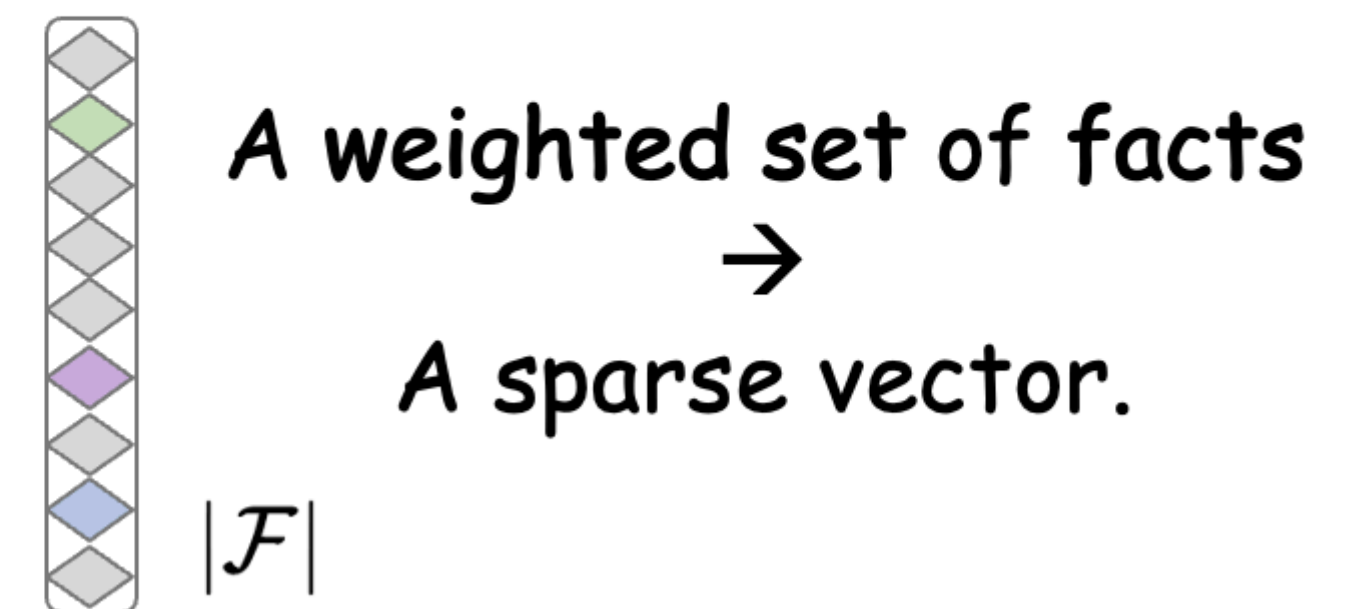
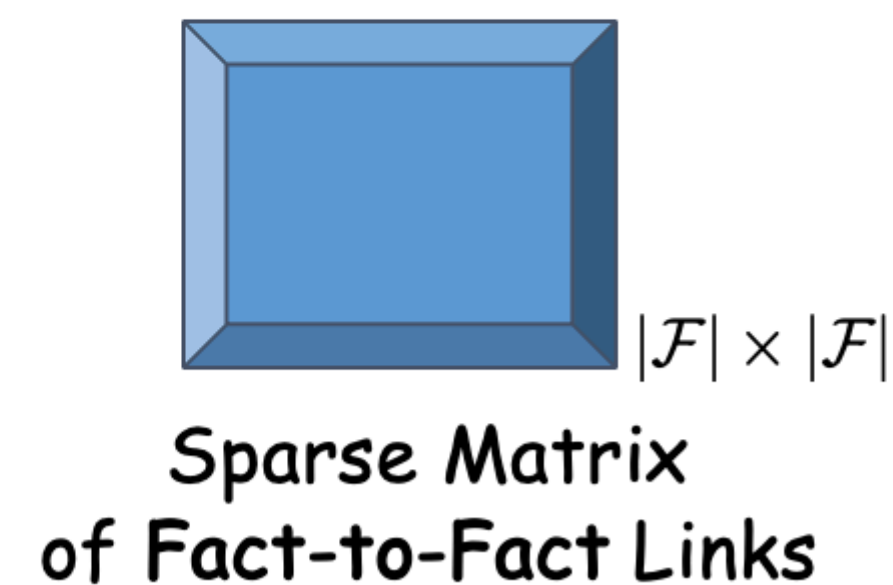
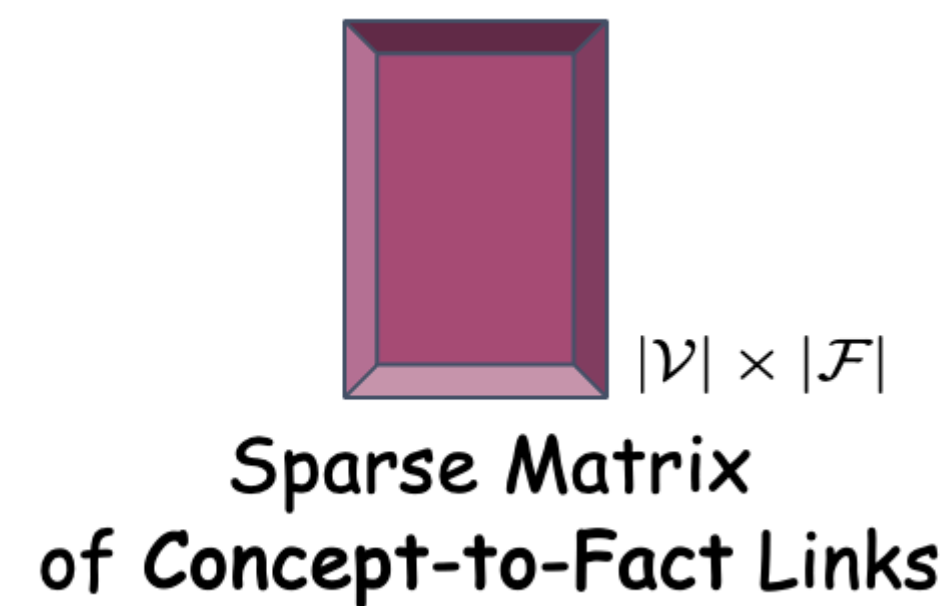
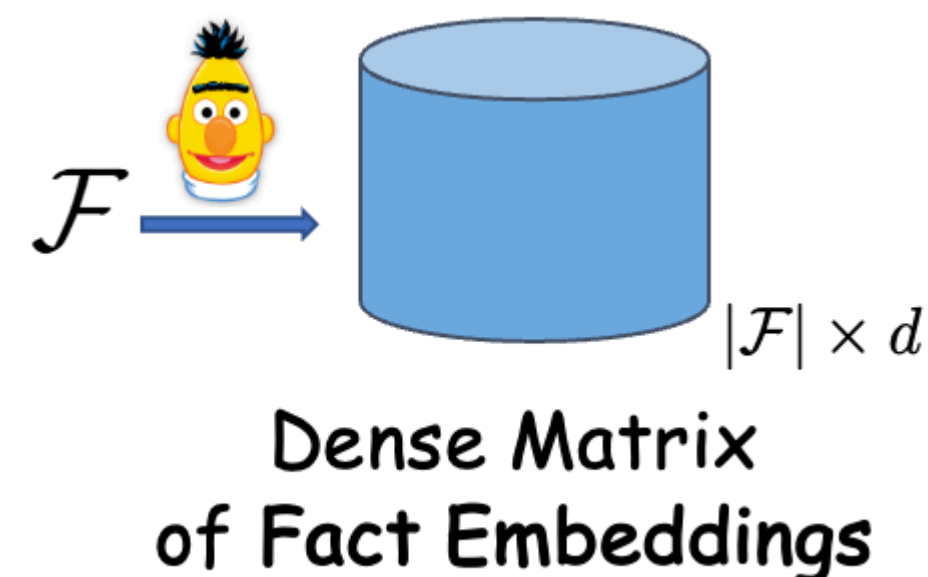
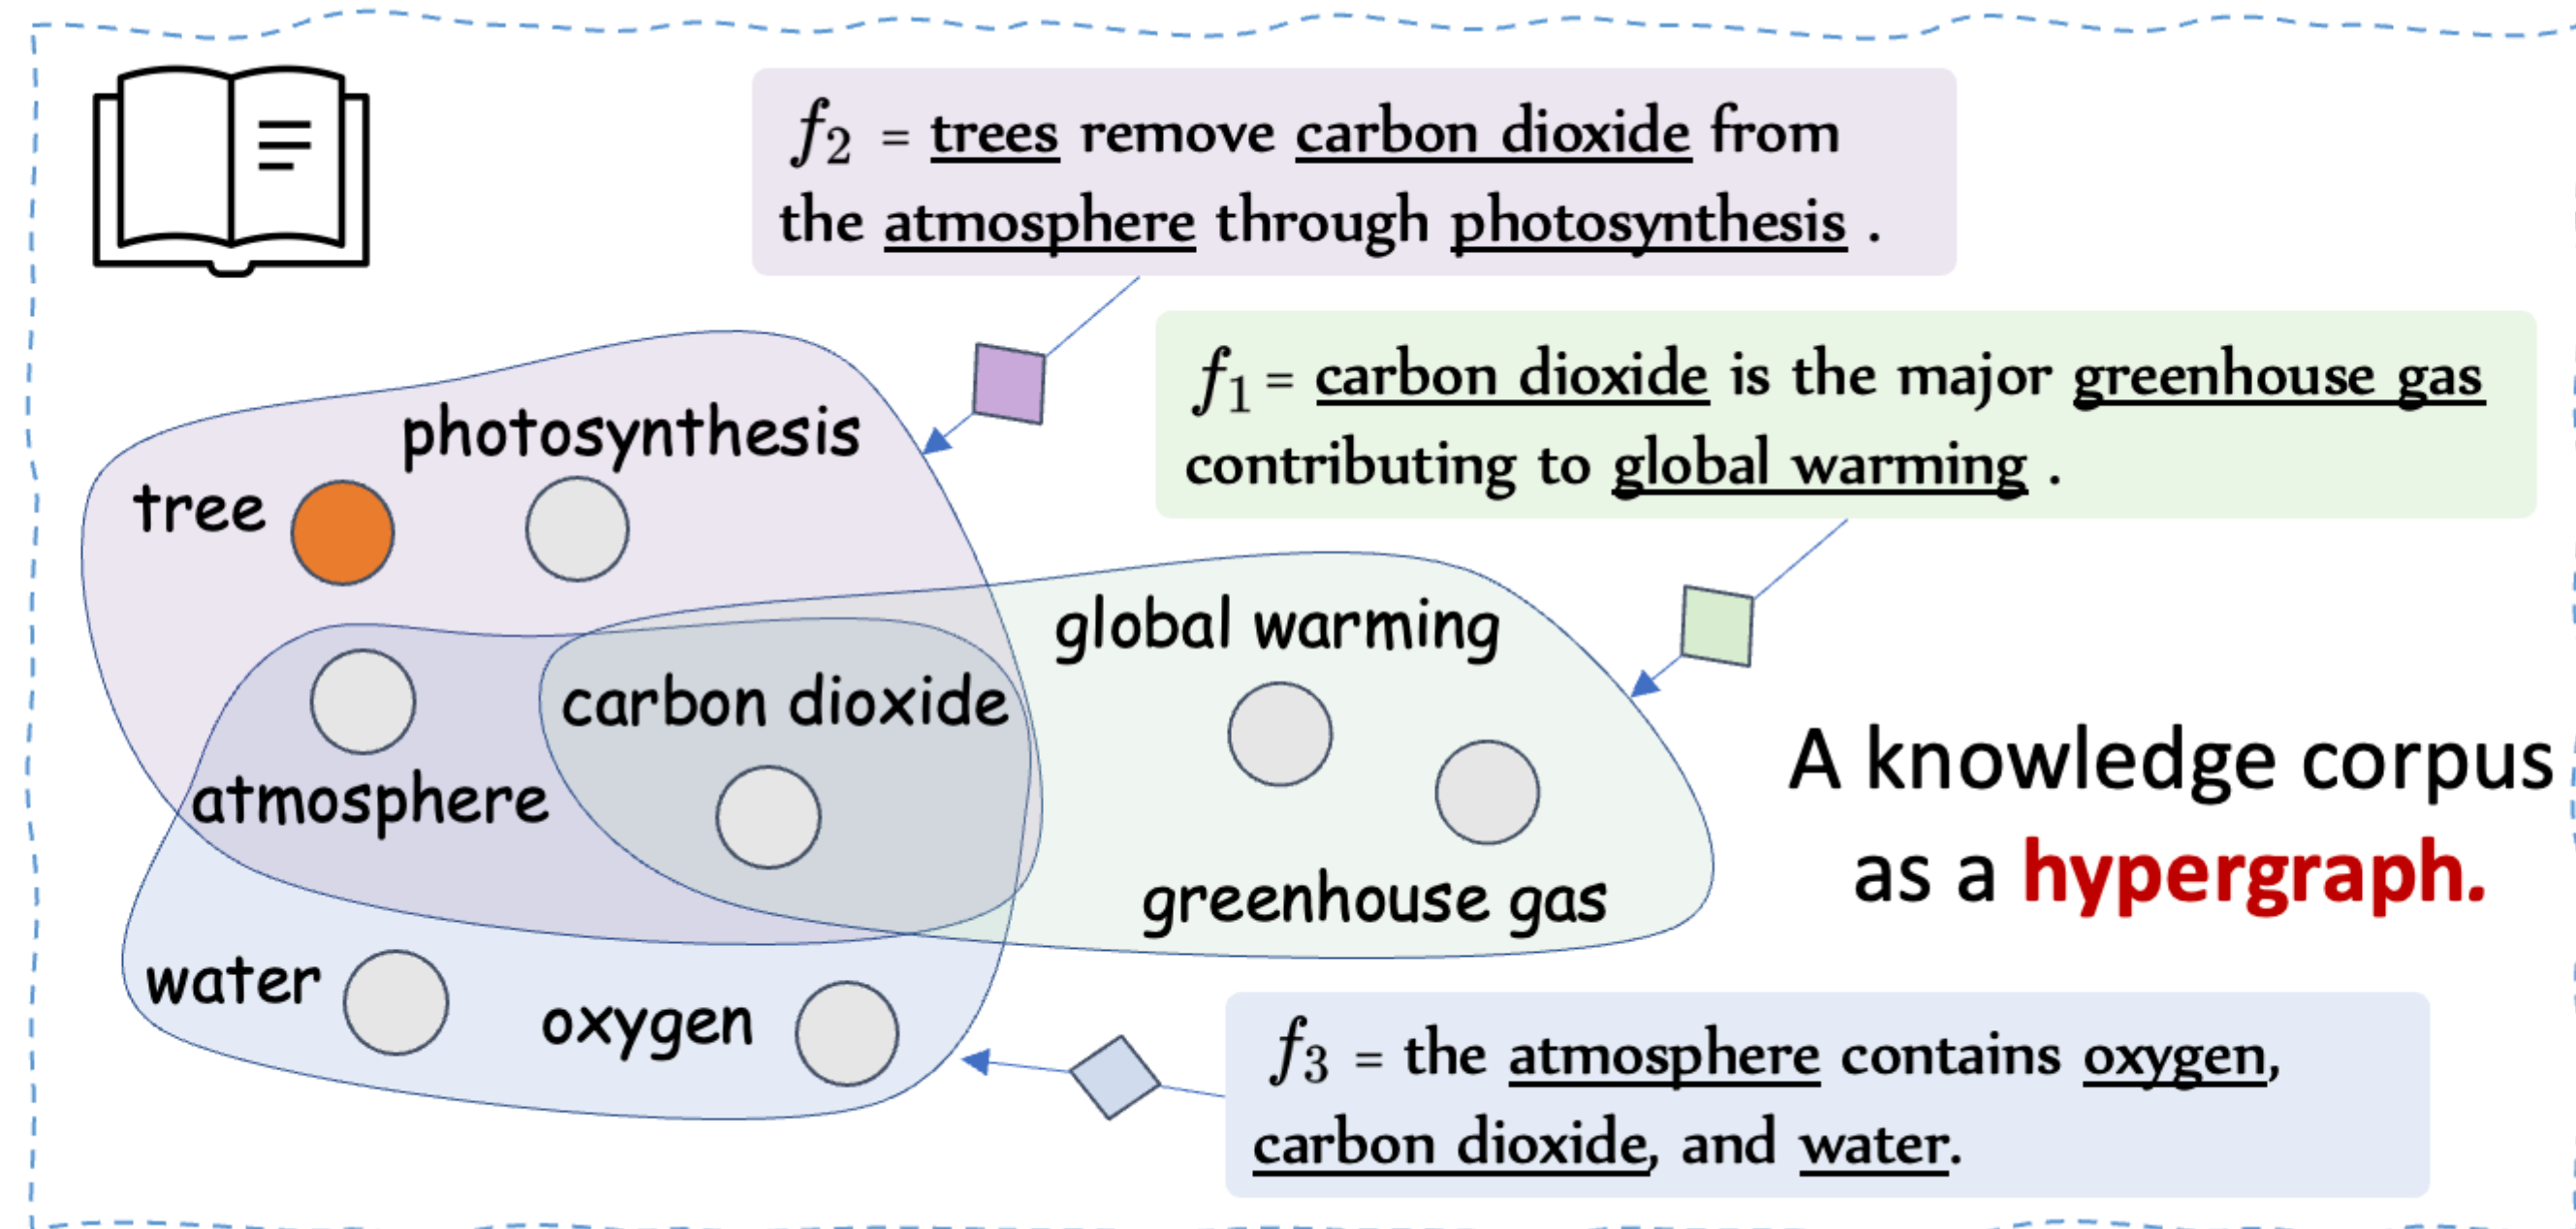
a **corpus** of common-sense facts, e.g., **GenericsKB**. \mathcal{F}

$$f_i \in \mathcal{F}$$

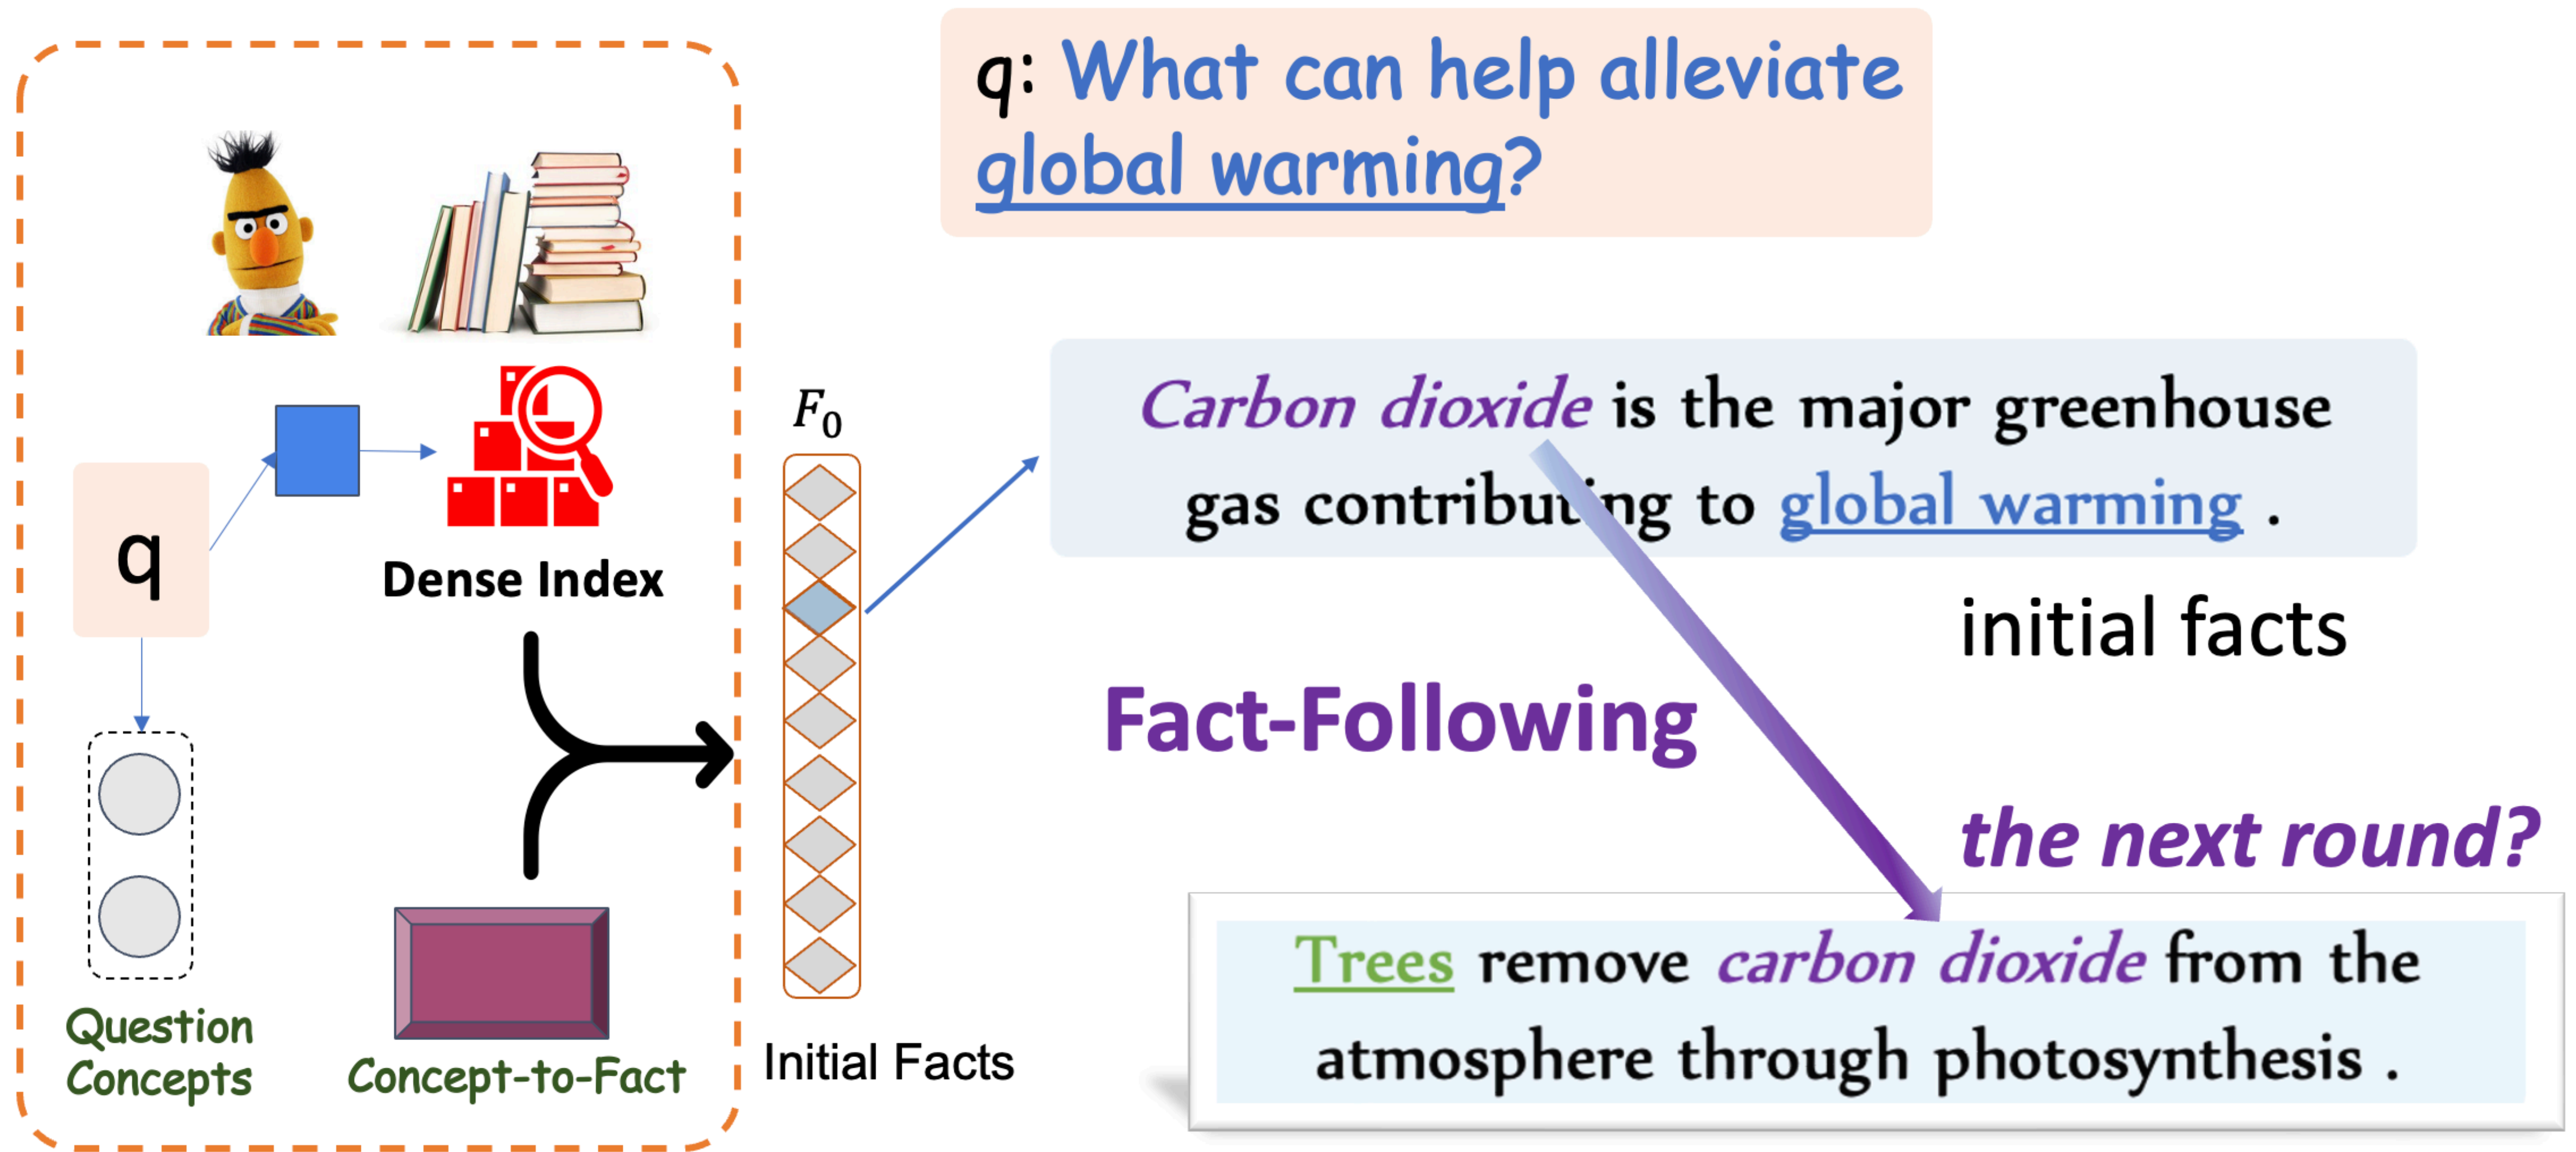
A **fact** is a sentence of generic commonsense knowledge

$$c_j \in \mathcal{V}$$

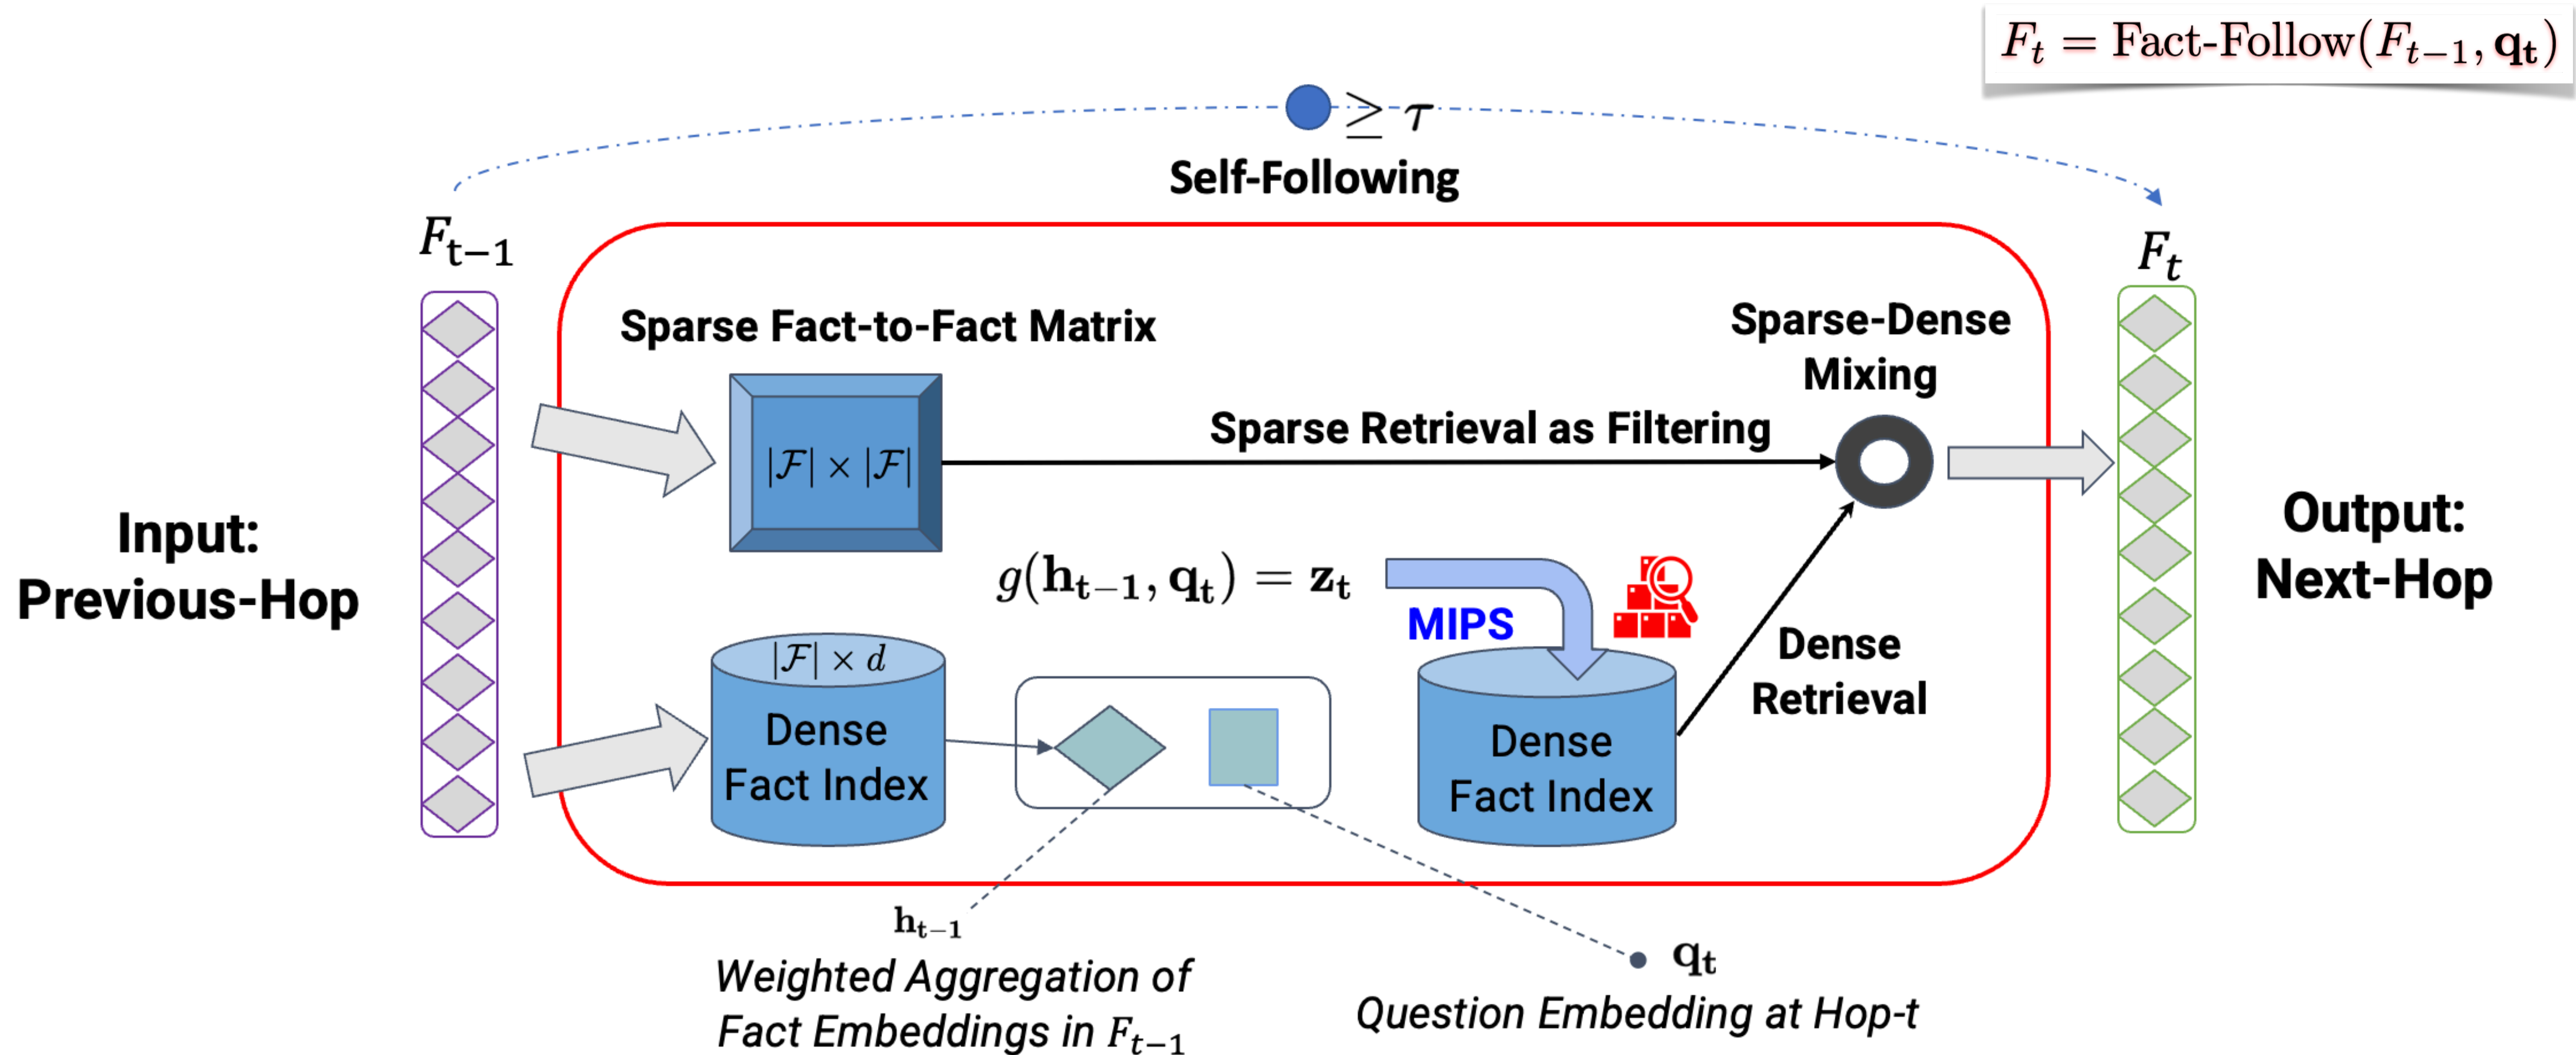
A **concept** is a noun or noun-chunk that are mentioned in \mathcal{F}



Differentiable Fact-Following Operations (DrFact) [17]



Differentiable Fact-Following Operations (DrFact) [17]



A quick comparison

| Methods | BM25 | DPR [14] / MDR [15] | DrKIT [16] | DrFact [17] |
|---------------------------------|--|---------------------|--------------------------------|----------------------------------|
| Knowledge Structure | A set of documents | A set of documents | Mention-Entity Bipartite Graph | Concept-Fact Hypergraph |
| Multi-hop Reasoning Formulation | - | - / Multiple-Round | Entity-Following | Fact-Following |
| Index for Dense Retrieval | - | Passage Embeddings | Mention Embedding | Fact Embeddings |
| Sparse Retrieval Method | TF-IDF based Index+ BM25 Ranking Func. | - | Entity-Mention Cooccurrence | Fact-to-Fact Matrix |
| Multi-Hop Questions | - | - / Single model | Aggregating Multiple Models | A single model w/ Self-Following |
| Intermediate Supervision | - | - | N/A | Distant Supervision |

Textual Knowledge

 **Structured Knowledge**

 **Un/Semi-structured Knowledge**


We are very **accurate**, **easy to modify**, **trustworthy**, and **verifiable**; However, we are **incomplete & hard to query!**

We are super **generalizable** and **easy to query!** But, we often **hallucinate** and make mistakes.

Parametric Knowledge

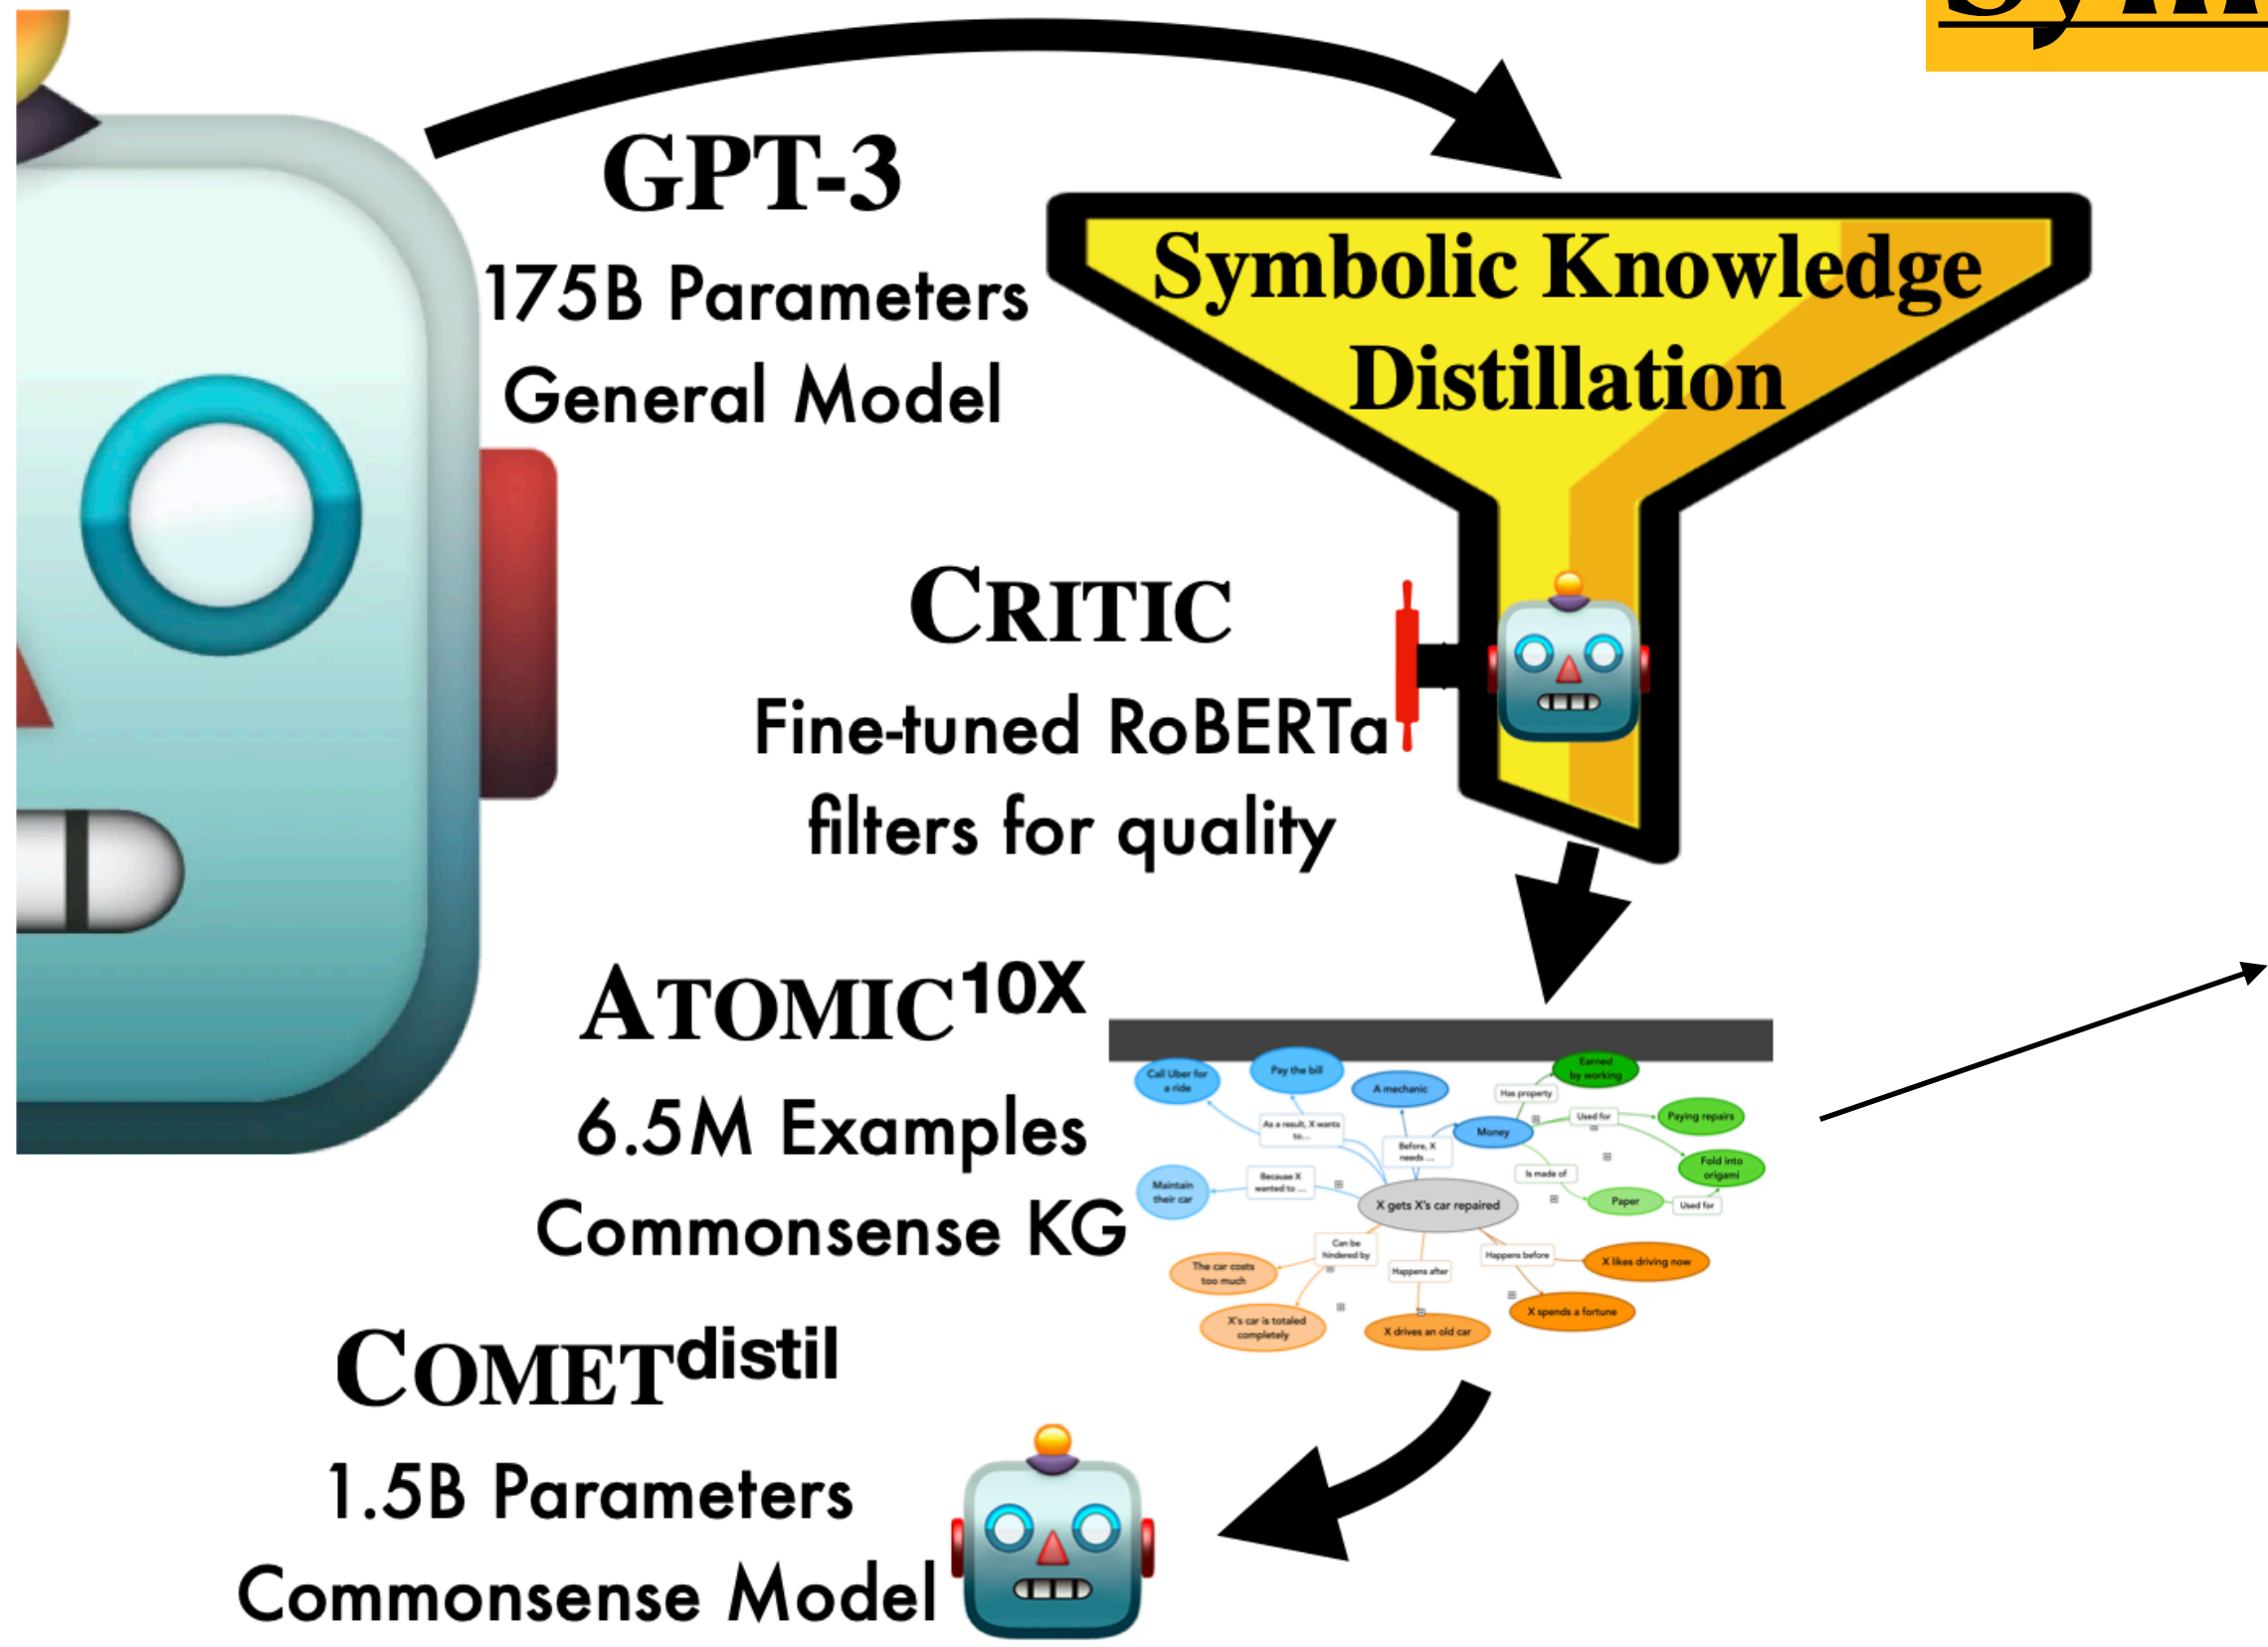


GPT-X



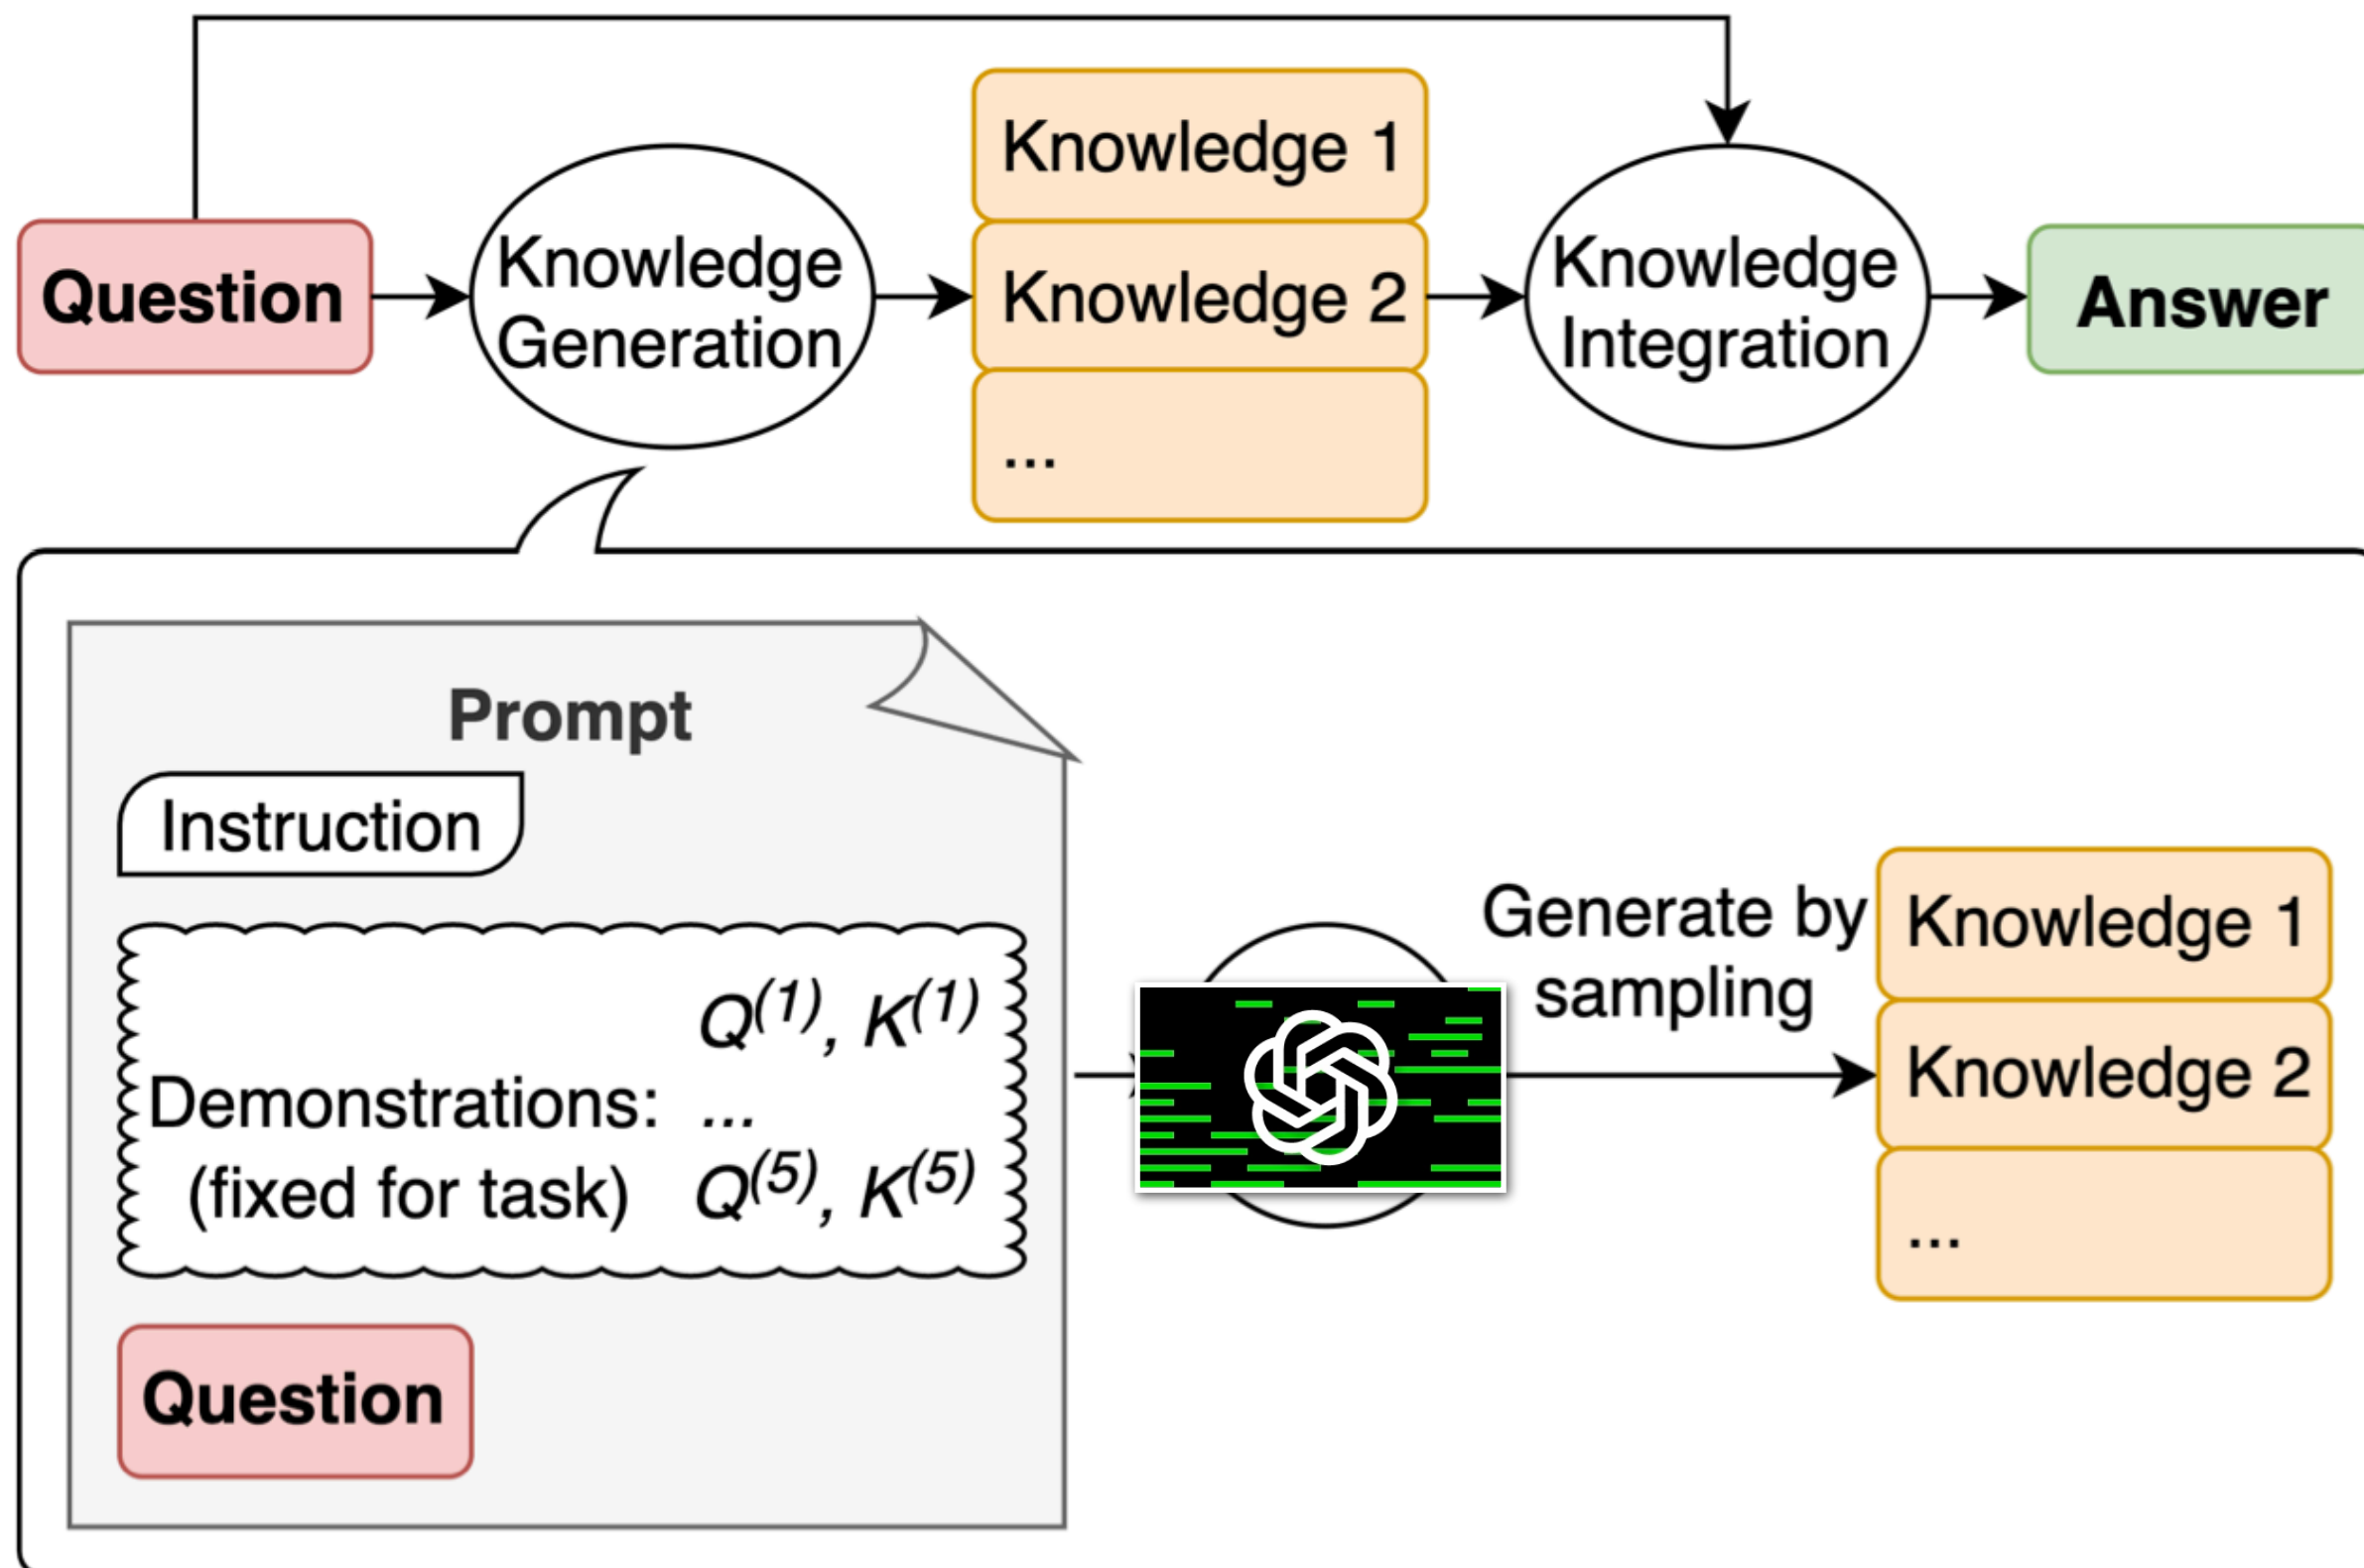
- We can be very expensive.
- You have no idea about our data/weights.

Symbolic Knowledge Distillation [19]



| | ATOMIC^{10X} | |
|----------------------------------|--|--------------------------------|
| X starts running | xEffect <i>so, X</i> | gets in shape |
| X and Y engage in an argument | xWant <i>so, X wants</i> | to avoid Y |
| X learns to type fast | xNeed <i>X needed</i> | to have taken typing lessons |
| X steals his grandfather's sword | xEffect <i>so, X</i> | is punished by his grandfather |
| X takes up new employment | xIntent <i>because X wants</i> | to be self sufficient |

Generated Knowledge Prompting (GKP) [20]



| Task | NumerSense [22] |
|---------------|--|
| Prompt | <p>Generate some numerical facts about objects. Examples:</p> <p>Input: penguins have <mask> wings. Knowledge: <i>Birds have two wings. Penguin is a kind of bird.</i></p> <p>...</p> <p>Input: a typical human being has <mask> limbs. Knowledge: <i>Human has two arms and two legs.</i></p> <p>Input: {question} Knowledge:</p> |

Input: The word children means <mask> or more kids.
Knowledge: The word child means one kid. The word children is the plural form of the word child.
Prediction: two

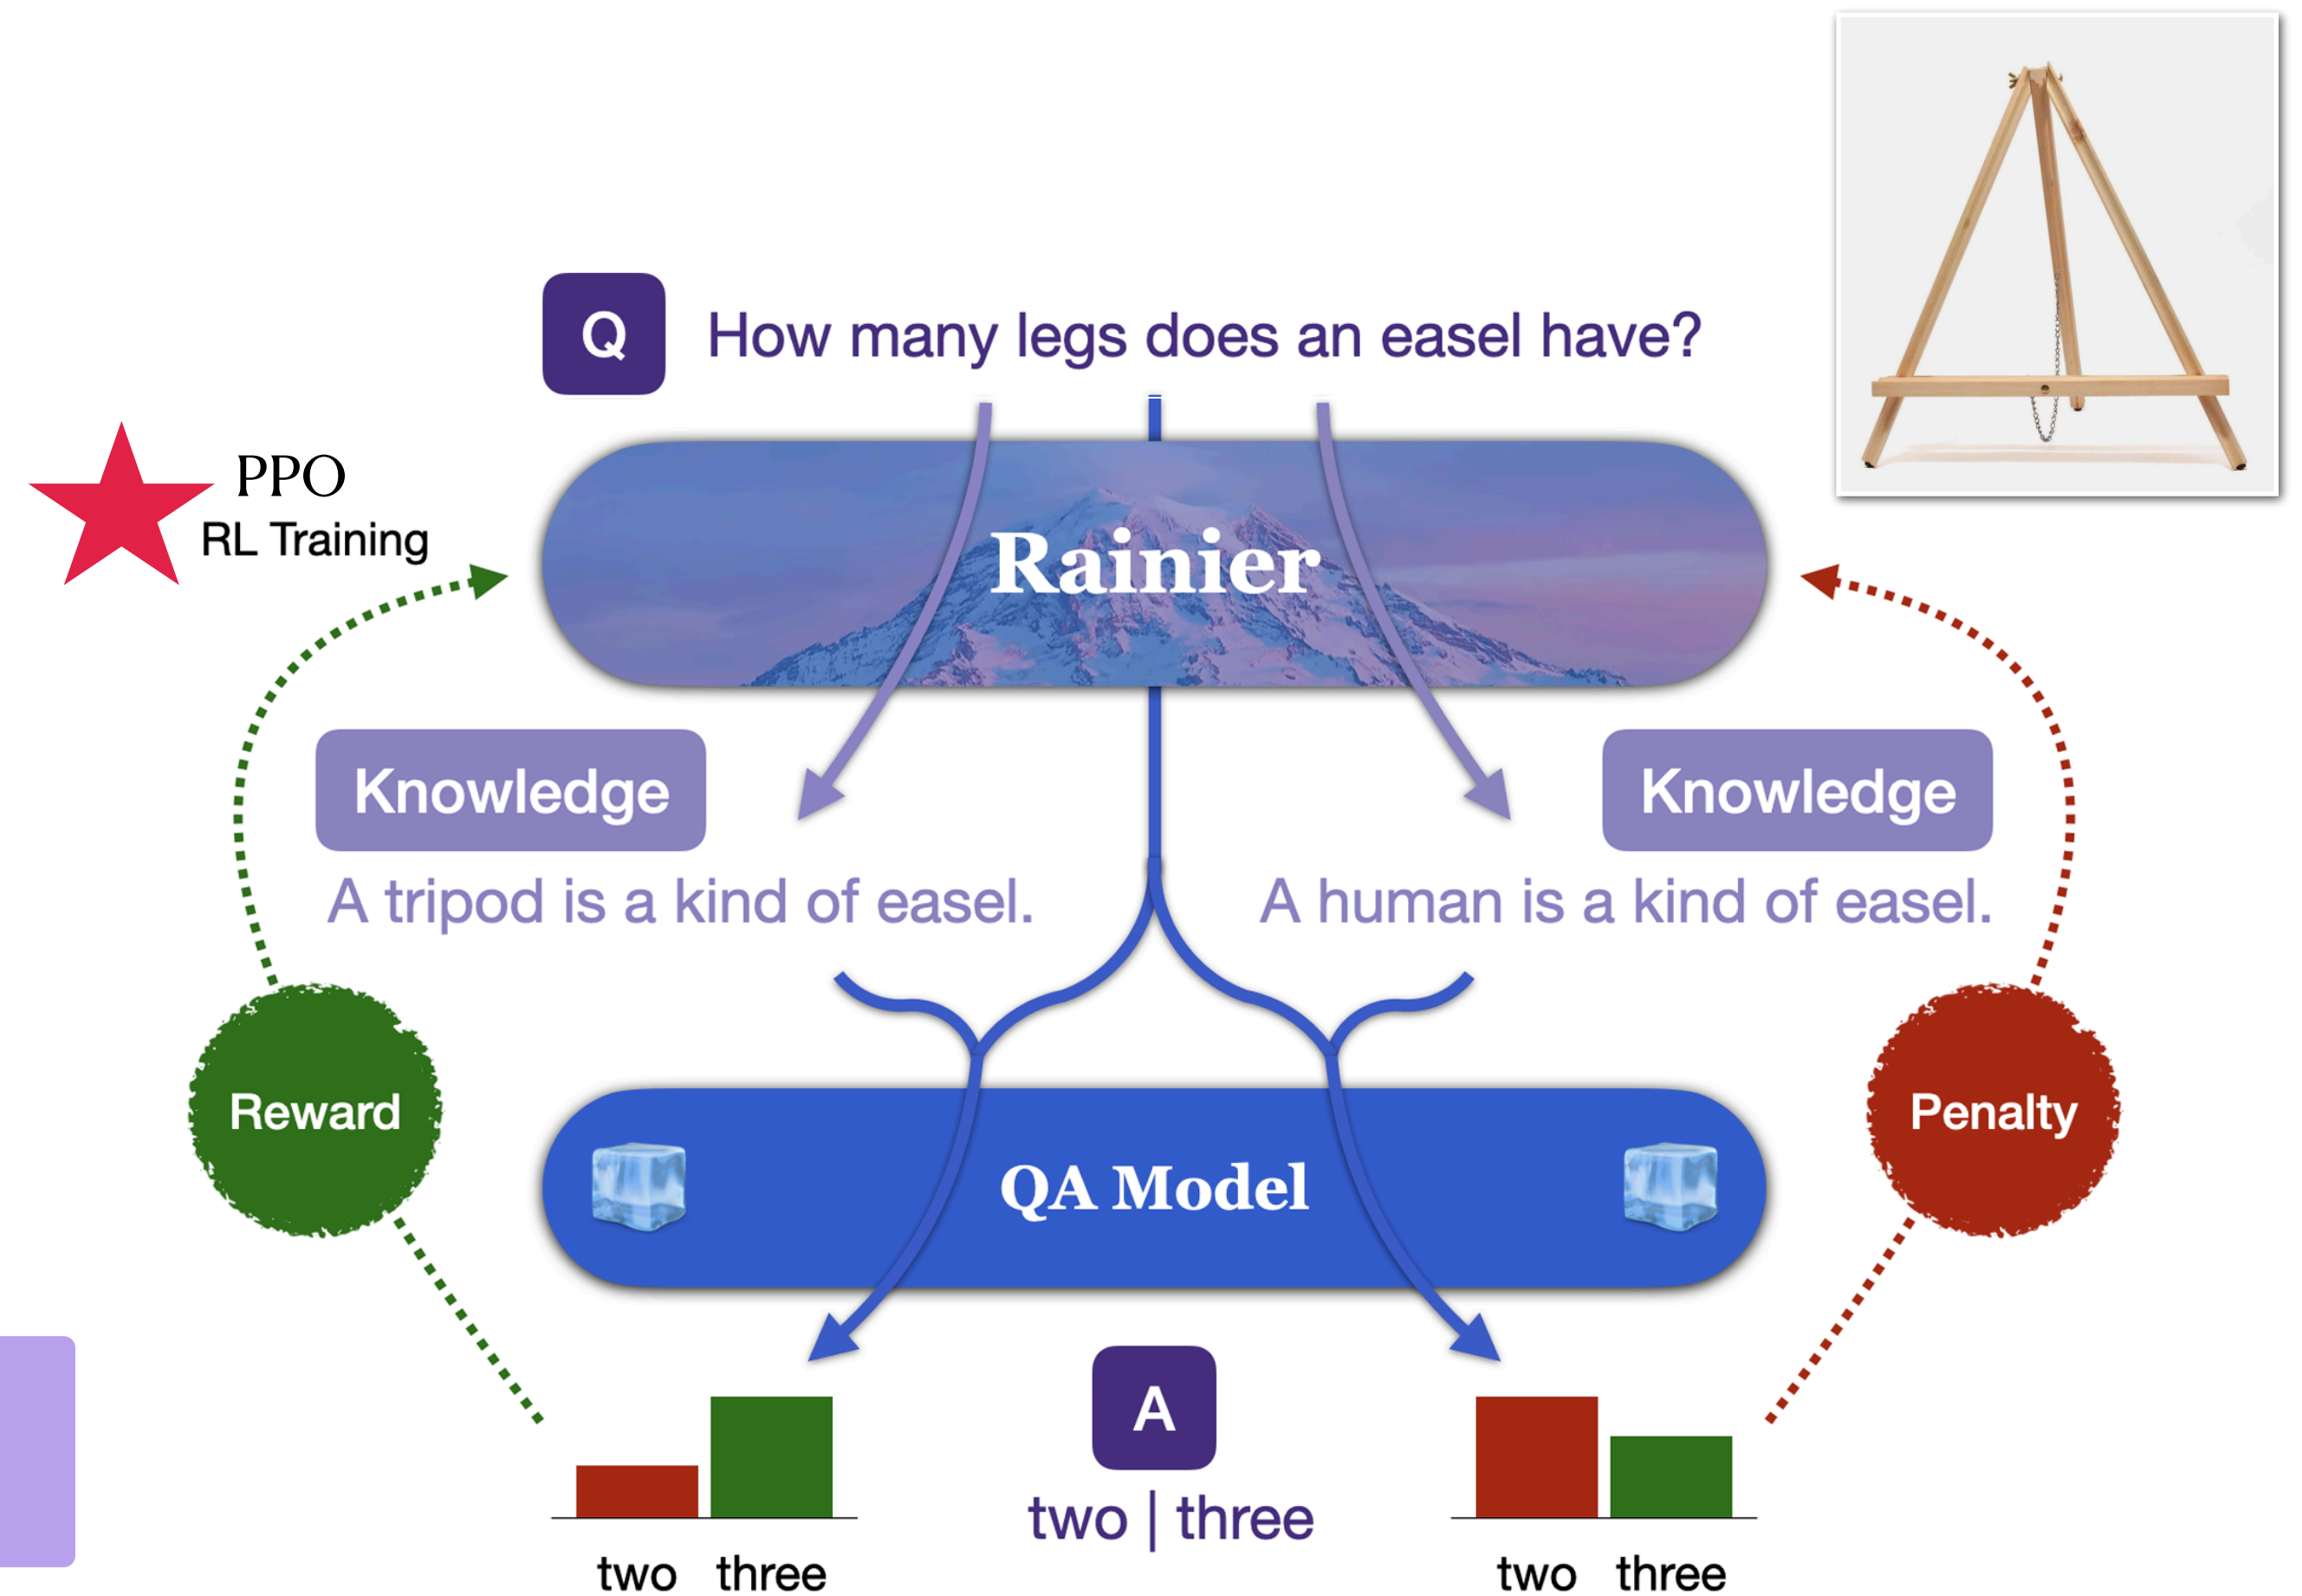
Reinforced Knowledge Introspector (Rainier) [21]

| Task | NumerSense |
|---------------|--|
| Prompt | Generate some numerical facts about objects. Examples: Input: penguins have <mask> wings. Knowledge: <i>Birds have two wings. Penguin is a kind of bird.</i> ... Input: a typical human being has <mask> limbs. Knowledge: <i>Human has two arms and two legs.</i> Input: How many legs does an easel have? Knowledge: |

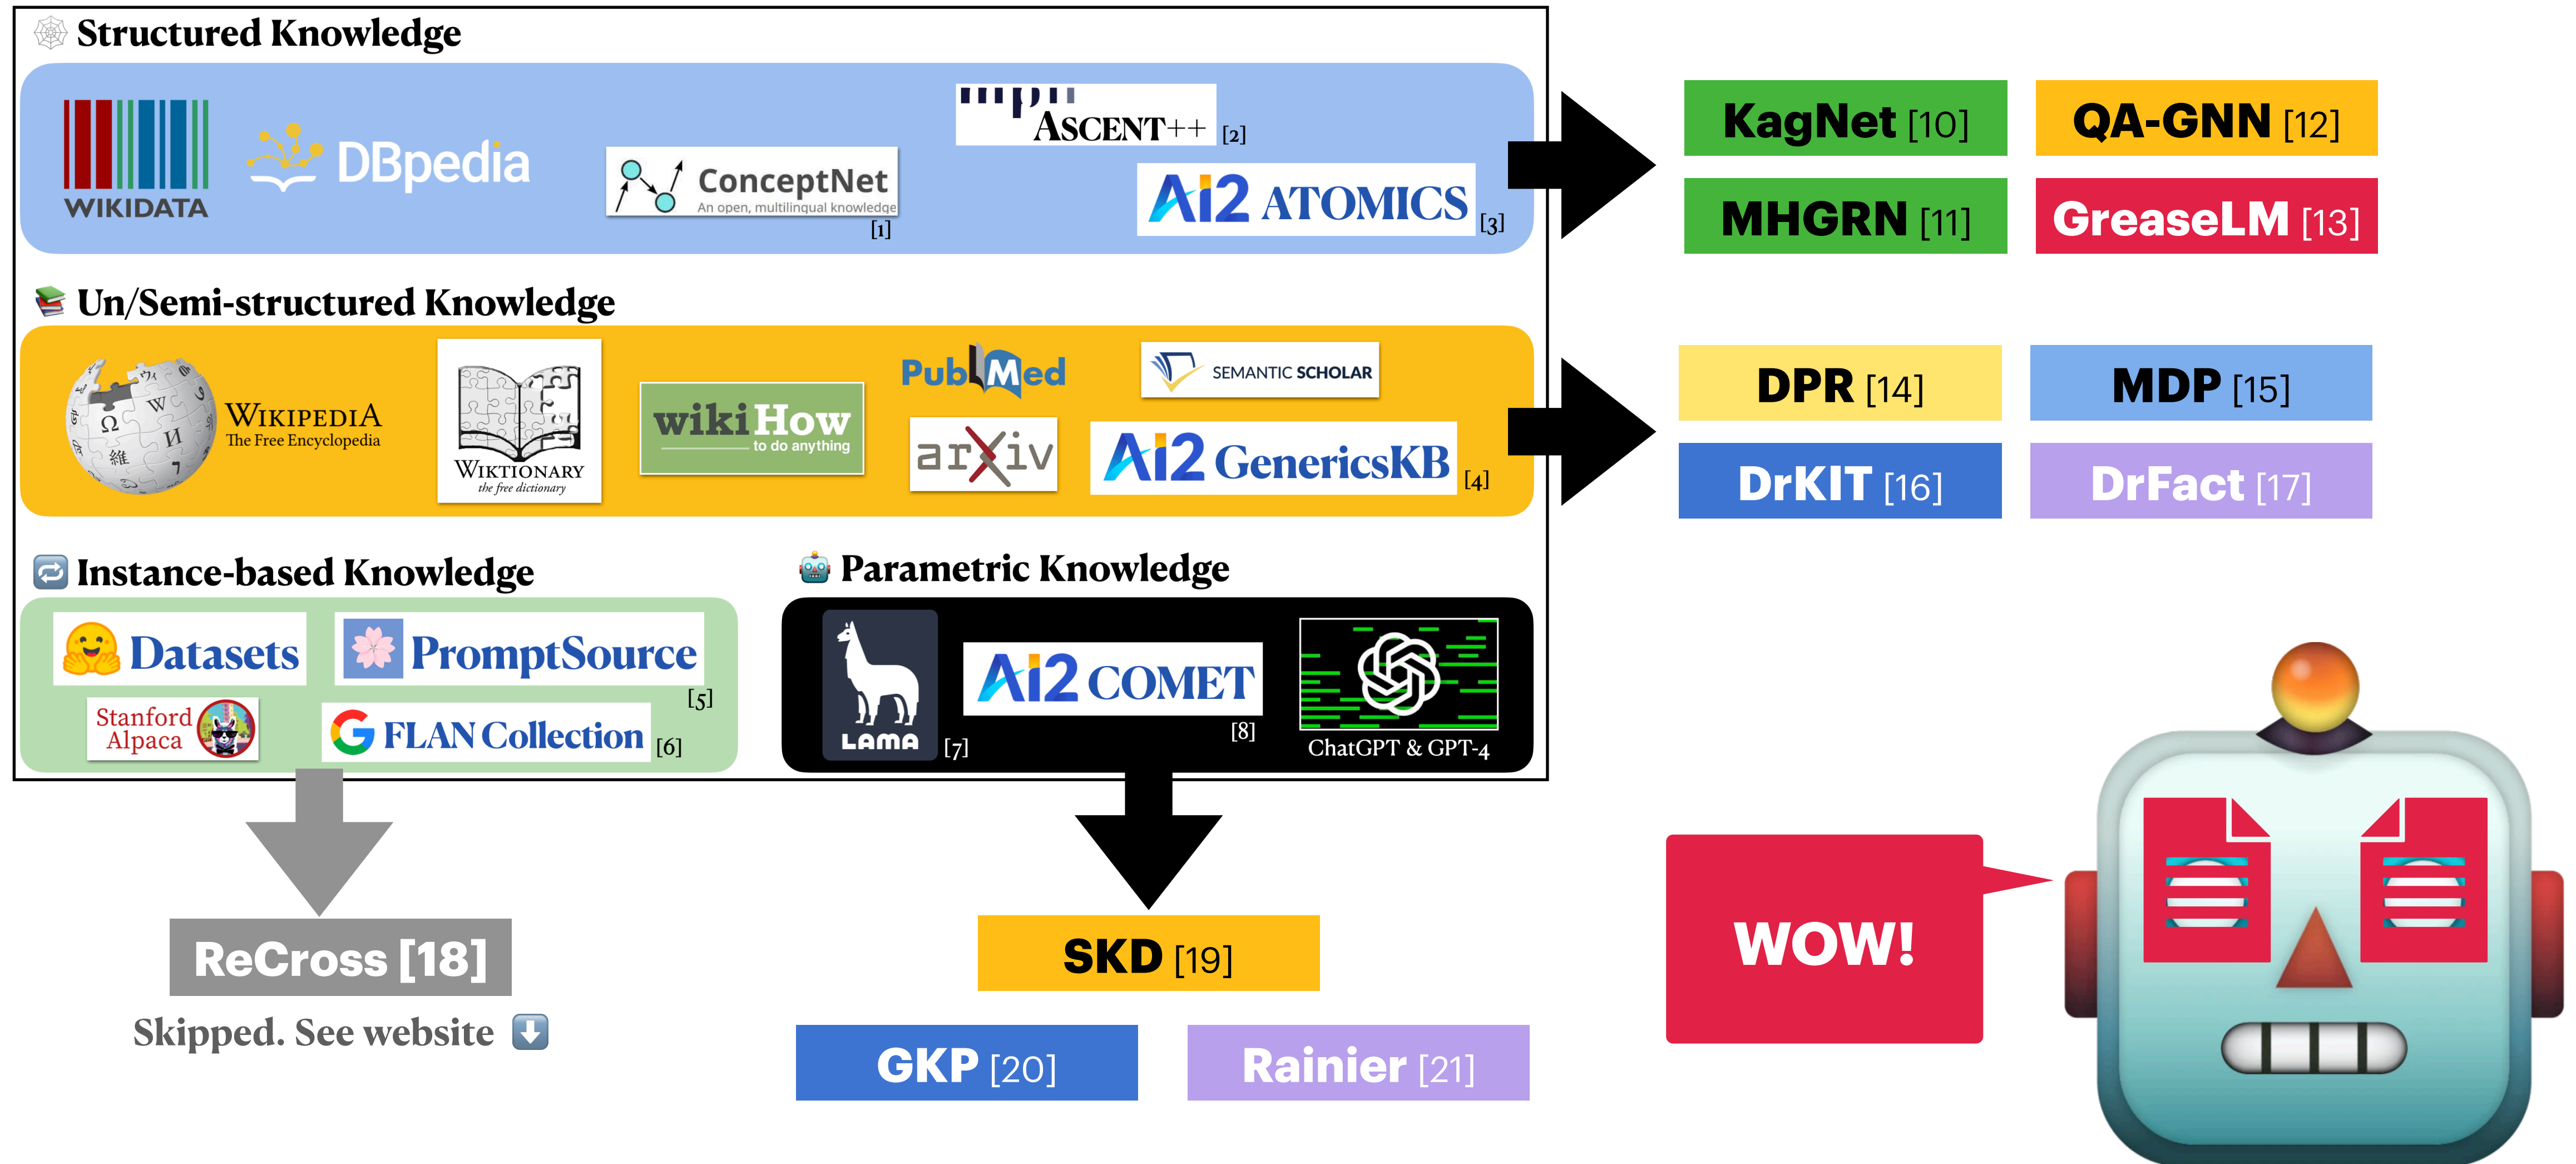
Generated Knowledge Prompting (GKP)



- A tripod is a kind of easel.
- A human is a kind of easel.




Conclusion



Reference is on the website: <https://yuchenlin.xyz/>

Future Directions



- 1 How can we merge all these different sources of knowledge into a **unified knowledge model**?
- 2 How can we deliver a more **faithful** and **interpretable** reasoning models with **low cost**?
- 3 How do we **collect** and **ground** knowledge of **real world** and **social interactions** between humans?

