

Tutorial Outline

- ⦿ **Part 1: Introduction** of RecSys in the era of LLMs (Dr. Wenqi Fan)
- ⦿ **Part 2: Preliminaries** of RecSys and LLMs (Dr. Yujuan Ding)
- ⦿ **Part 3: Pre-training** paradigms for adopting LLMs to RecSys (Dr. Yujuan Ding)
- ⦿ **Part 4: Fine-tuning paradigms for adopting LLMs to RecSys (Liangbo Ning)**
- **Part 5: Prompting** paradigms for adopting LLMs to RecSys (Shijie Wang)
- **Part 6: Future directions** of LLM-empowered RecSys (Dr. Wenqi Fan)

Website of this tutorial
Check out the slides and more information!



PART 4: RecSys Fine-tuning



Presenter
Liangbo NING
HK PolyU

- **Fine-tuning in NLP**
 - Pre-training then fine-tuning
 - Parameter Efficient Fine-tuning (PEFT)
 - Fine-tuning with reinforcement learning
- Fine-tuning LLM-based RecSys
- Parameter efficient fine-tuning for LLM-based RecSys

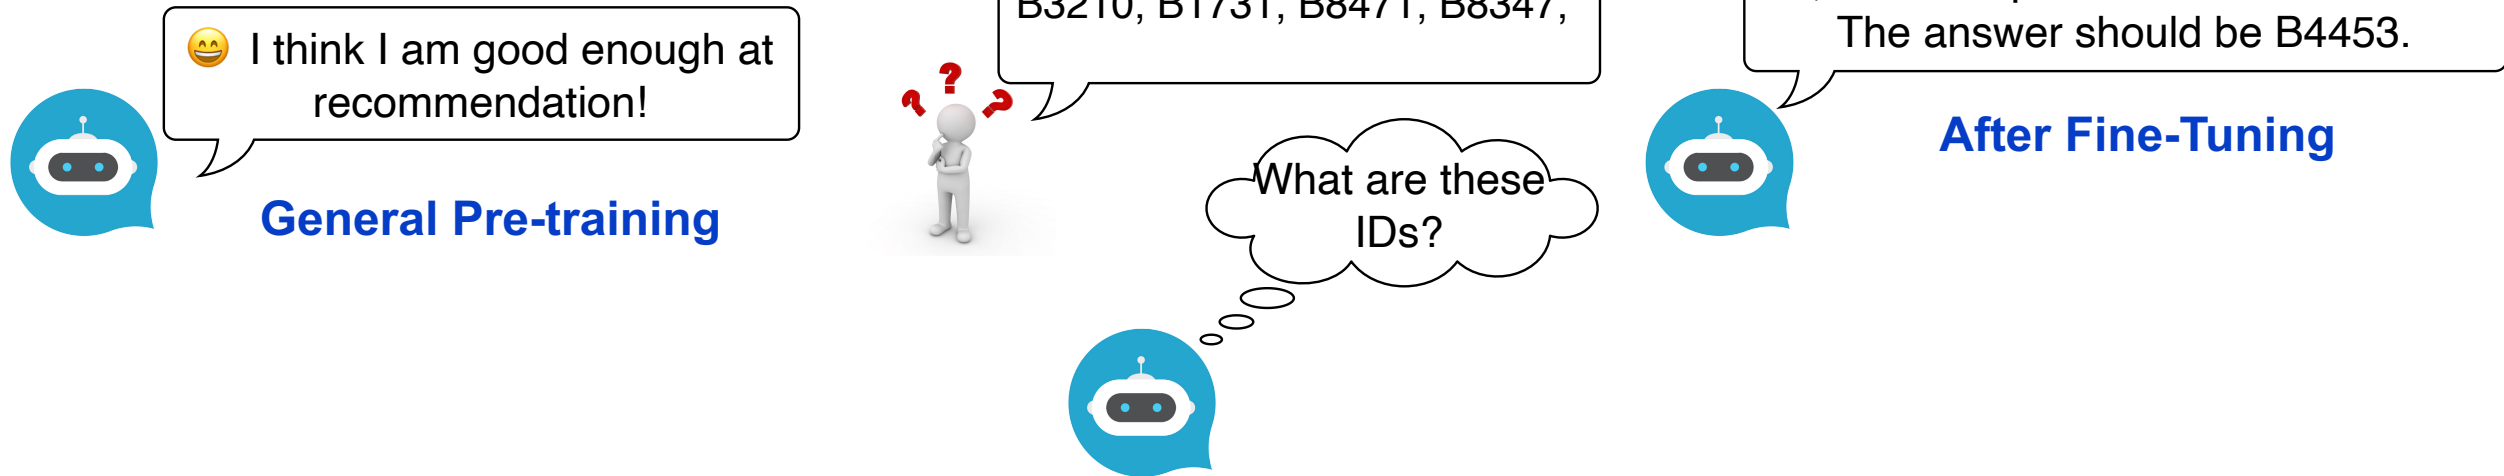


Fine-tuning in NLP



□ What is Fine-tuning and Why Fine-tuning?

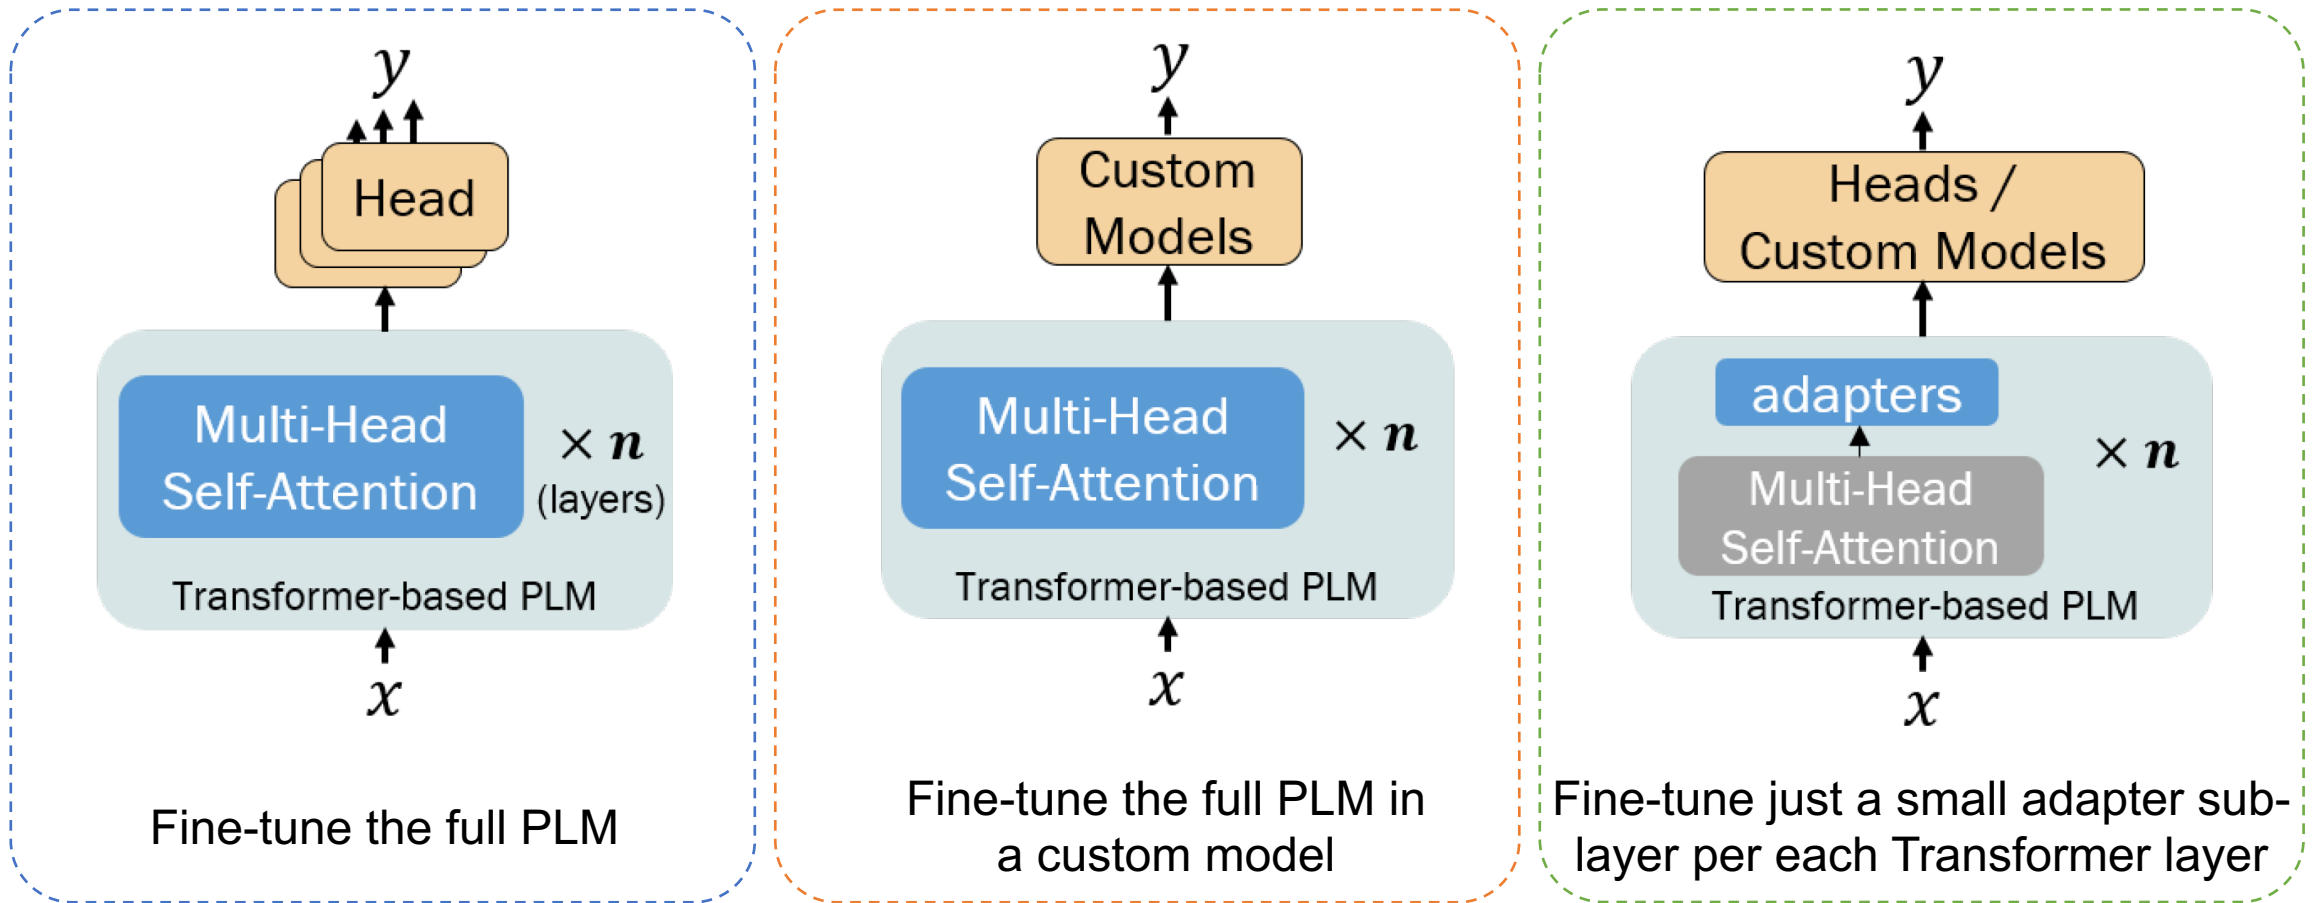
- ❖ **Gaps** between the pre-training tasks and downstream tasks still exist
- ❖ Fine-tuning means **training pre-trained LLMs on downstream tasks** to fit the requirements



Fine-tuning in NLP: pre-train then fine-tune

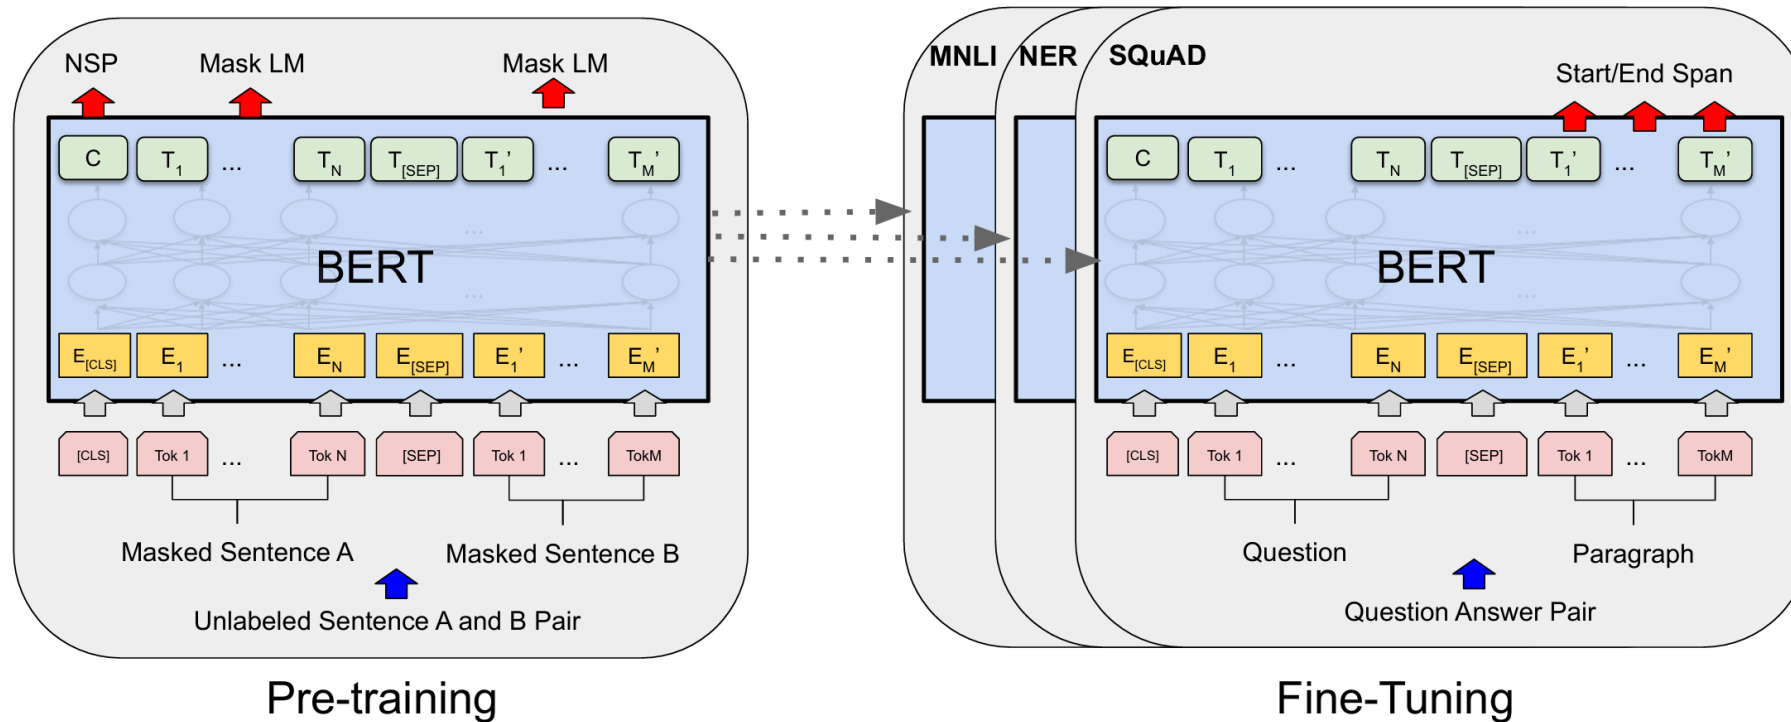


- Typical “pre-train then fine-tune” strategies

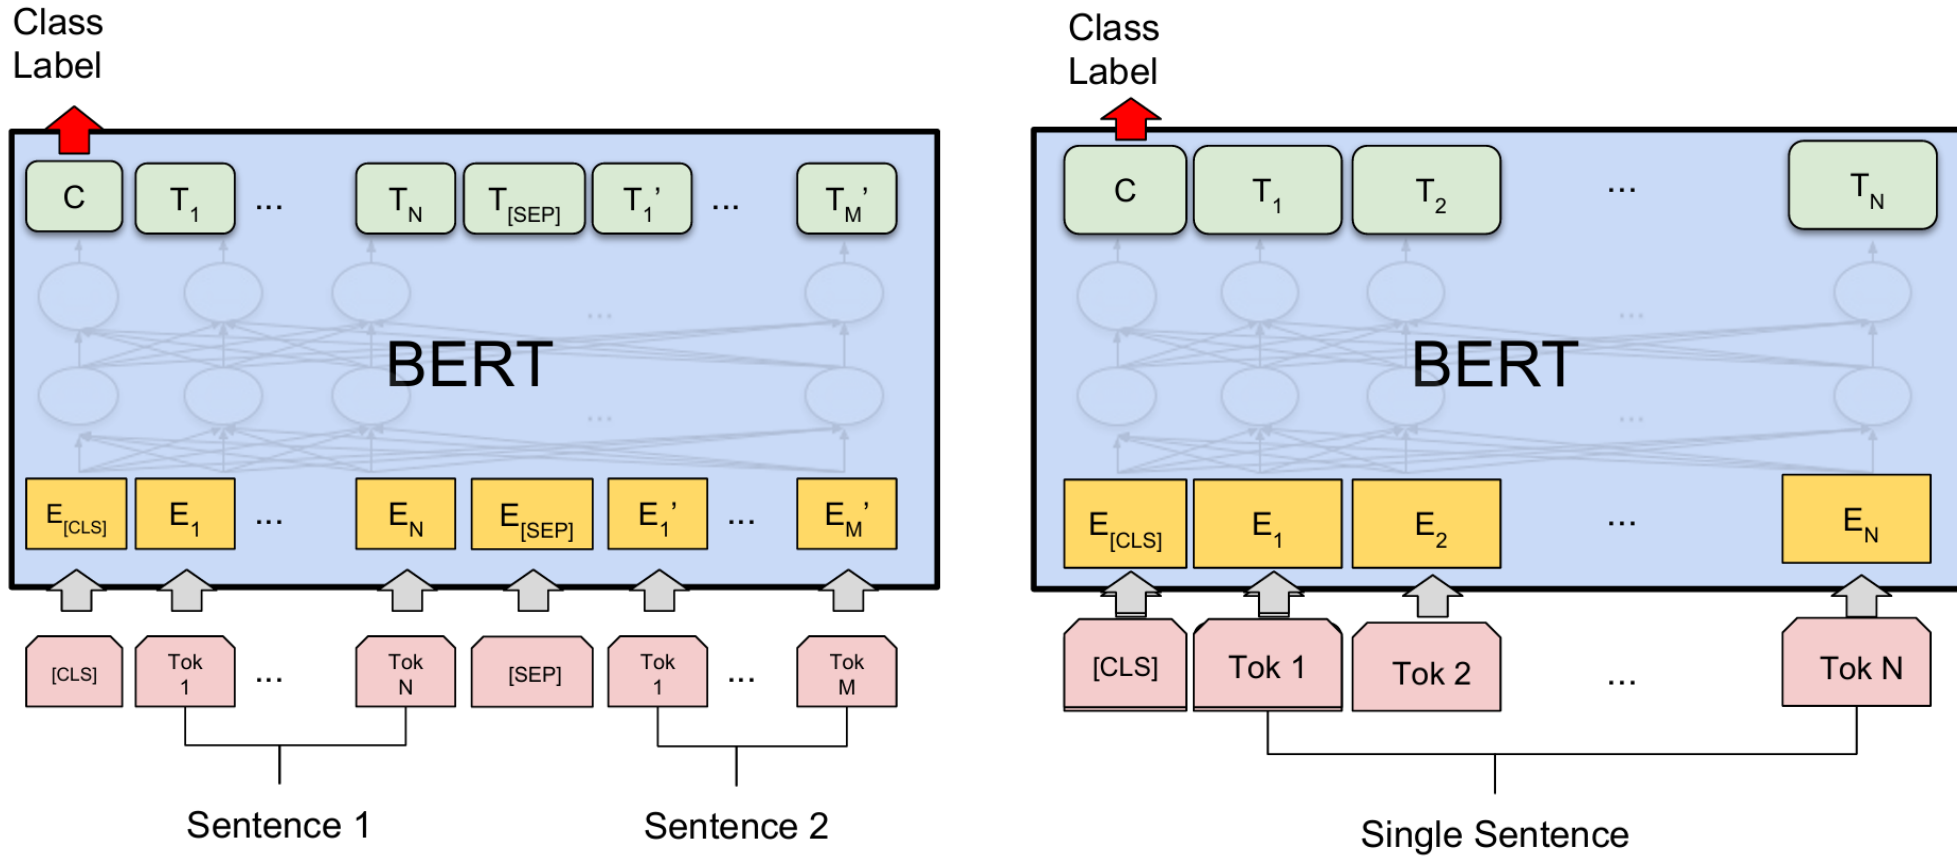


Fine-tuning the pre-trained language model (PLM)

- This approach **fine-tunes some or all the layers of the PLM** and then adds one or two simple output layers (known as **prediction heads**).



Fine-tuning in NLP: Fine-tuning the PLM

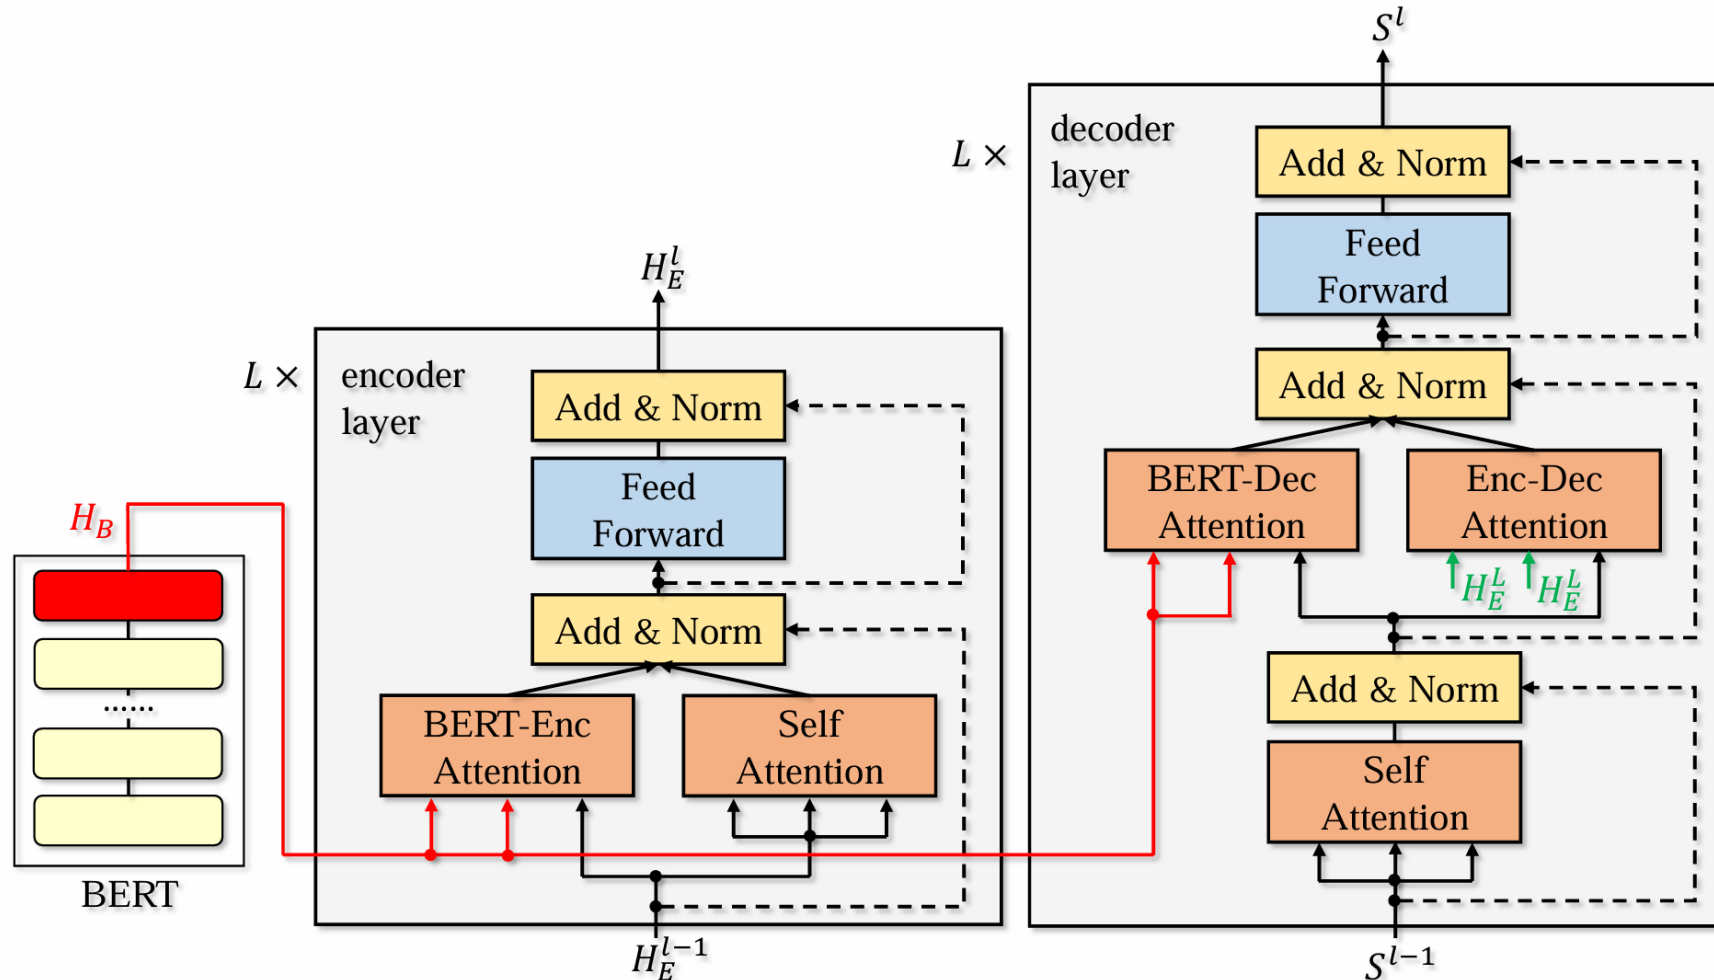


Fine-tuning in NLP: Fine-tuning the PLM



❑ Customized Models

❖ Some tasks require significant **additional architecture on top of a language model.**



PART 4: RecSys Fine-tuning



Website of this tutorial

- ⦿ **Fine-tuning in NLP**
 - ⦿ Pre-training then fine-tuning
 - **Parameter Efficient Fine-tuning (PEFT)**
 - Fine-tuning with reinforcement learning
- Fine-tuning LLM-based RecSys
- Parameter efficient fine-tuning for LLM-based RecSys



Fine-tuning in NLP: Parameter Efficient Fine-tuning (PEFT)



- ❑ What is Parameter Efficient Fine-tuning (PEFT)?
 - ❖ As LLMs scale up to **billion weights**, consumable GPUs like 3090 and 4090 gradually fail to contain all the weights in their memory
 - ❖ Parameter Efficient Fine-tuning aims to save GPU memory and boost training

- ❑ Why PEFT?
 - ❖ Making fine-tuning feasible for consumable GPUs
 - ❖ With major parameters fixed, it might relieve the problem of **catastrophic forgetting**

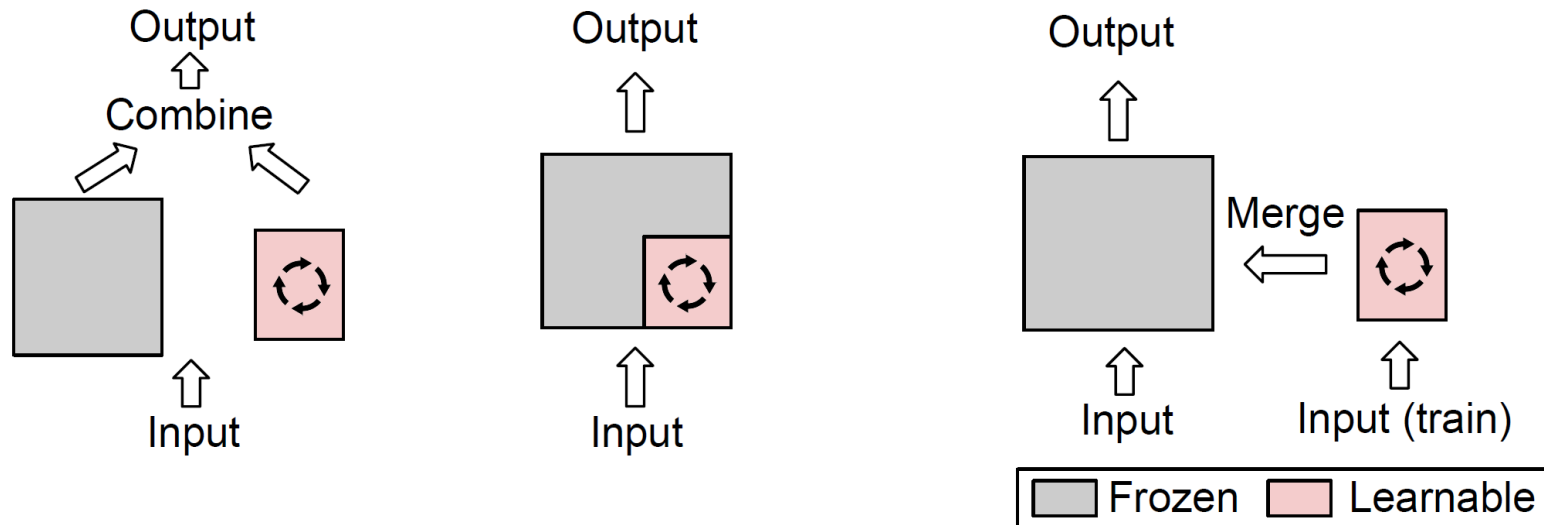


Fine-tuning in NLP: PEFT



- ❑ **Fine-tuning a separate, small network** that is tightly coupled with the PLM.
- ❑ **Selecting only a small number of the PLM's weights** to fine-tune or keep.
- ❑ Introduce additional **low-rank trainable parameters**, which are integrated with the original model.

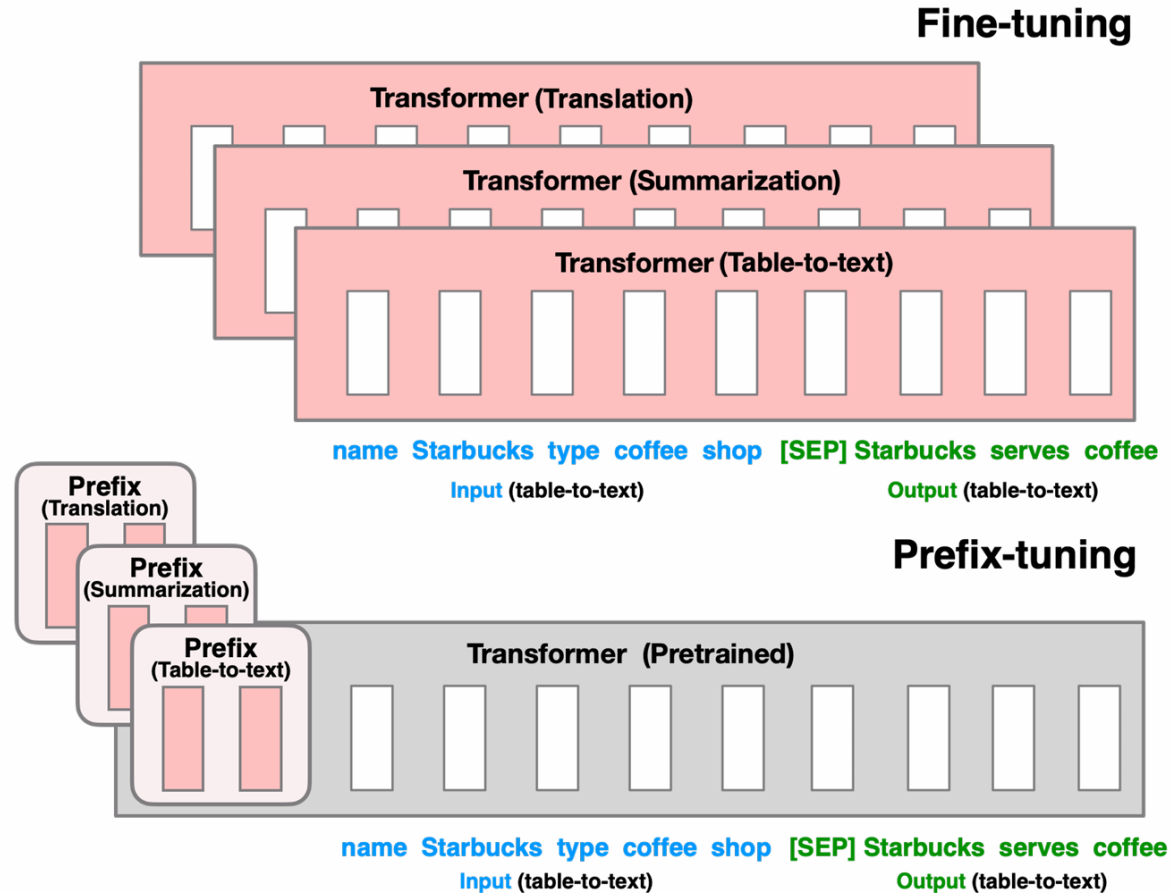
(a) Additive PEFT (b) Selective PEFT (c) Reparameterization PEFT



Fine-tuning in NLP: Prefix-tuning



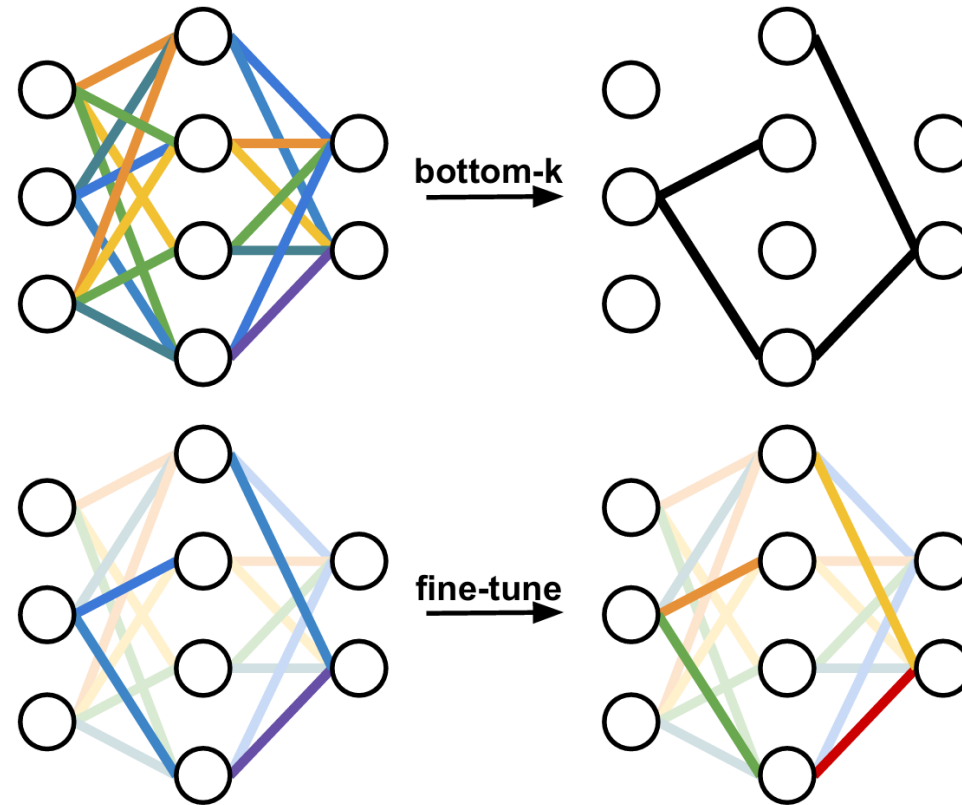
- Freezes the Transformer parameters and **only optimizes the prefix** (the red prefix blocks)



Fine-tuning in NLP: PaFi



- Select model parameters with the **smallest absolute magnitude as trainable**

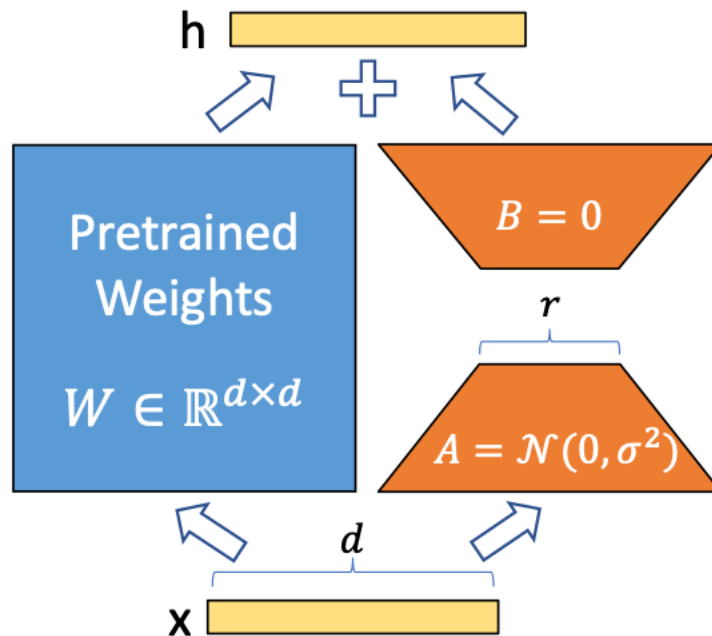




Fine-tuning in NLP:

Low-Rank Adaptation of LLMs (LoRA)

- ❑ Fine-tuning a 7B model needs $7,000,000,000 * 8 / 1024^3 \cong 52GB$ GPU memory
- ❑ LoRA only fine-tunes the **feed-forward networks** (FFNs)
 - ❖ Making it possible for consumable GPUs to train 7B and even 13B LLMs



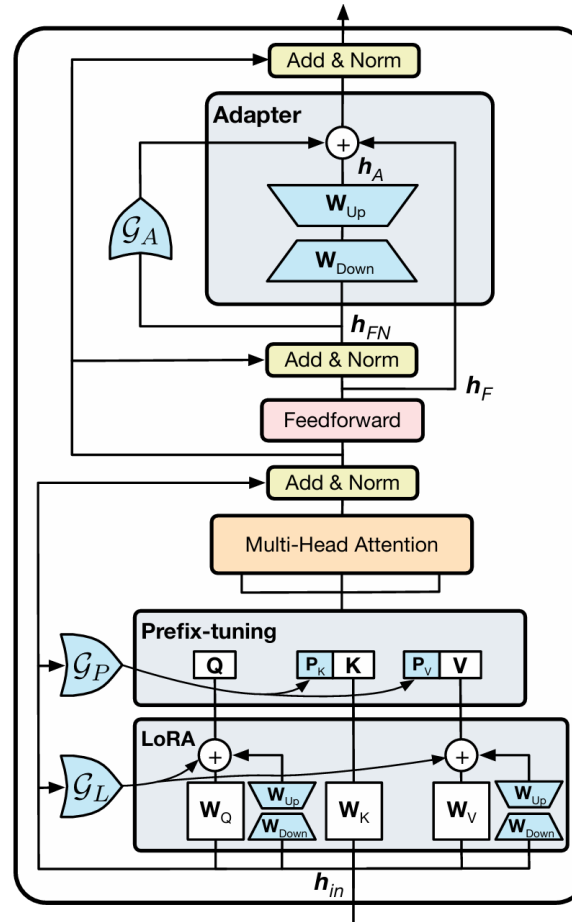
$$h = W_0x + \Delta Wx = W_0x + BAx$$



Fine-tuning in NLP: UNIPELT



- Integrate multiple PELT methods and controls them via a **gating mechanism**



PART 4: RecSys Fine-tuning



Website of this tutorial

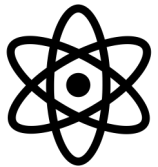
- ⦿ **Fine-tuning in NLP**
 - ⦿ Pre-training then fine-tuning
 - ⦿ Parameter Efficient Fine-tuning (PEFT)
 - **Fine-tuning with reinforcement learning**
- **Fine-tuning LLM-based RecSys**
- **Parameter efficient fine-tuning for LLM-based RecSys**



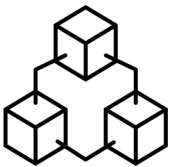
Fine-tuning with Reinforcement Learning



Reinforcement Learning based on Human Feedbacks (RLHF)



Proximal Policy Optimization (PPO)



Direct Preference Optimization (DPO)



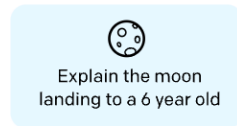
Fine-tuning with Reinforcement Learning: RLHF



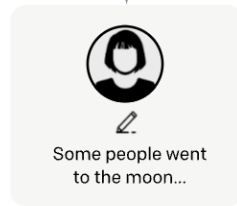
Step 1

Collect demonstration data, and train a supervised policy.

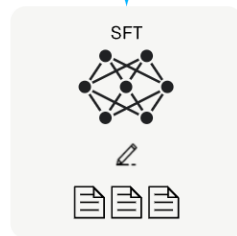
A prompt is sampled from our prompt dataset.



A labeler demonstrates the desired output behavior.



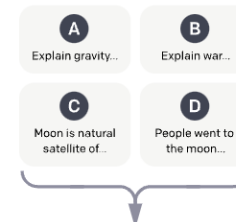
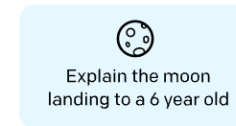
This data is used to fine-tune GPT-3 with supervised learning.



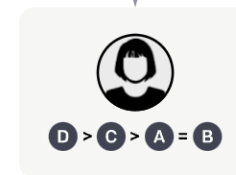
Step 2

Collect comparison data, and train a reward model.

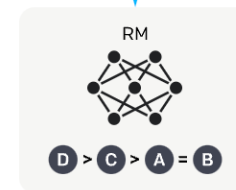
A prompt and several model outputs are sampled.



A labeler ranks the outputs from best to worst.



This data is used to train our reward model.



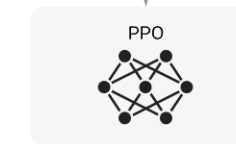
Step 3

Optimize a policy against the reward model using reinforcement learning.

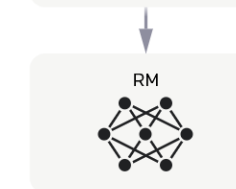
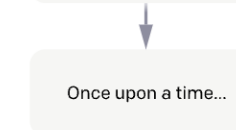
A new prompt is sampled from the dataset.



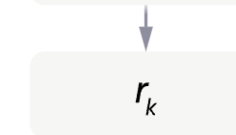
The policy generates an output.



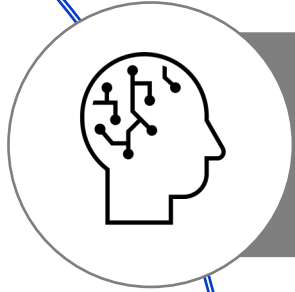
The reward model calculates a reward for the output.



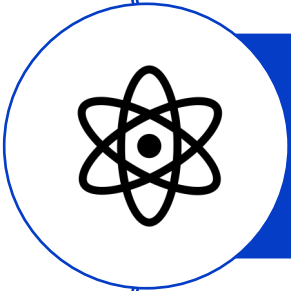
The reward is used to update the policy using PPO.



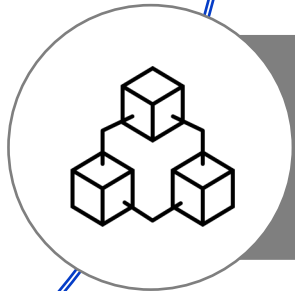
Fine-tuning with Reinforcement Learning



Reinforcement Learning based on Human Feedbacks (RLHF)



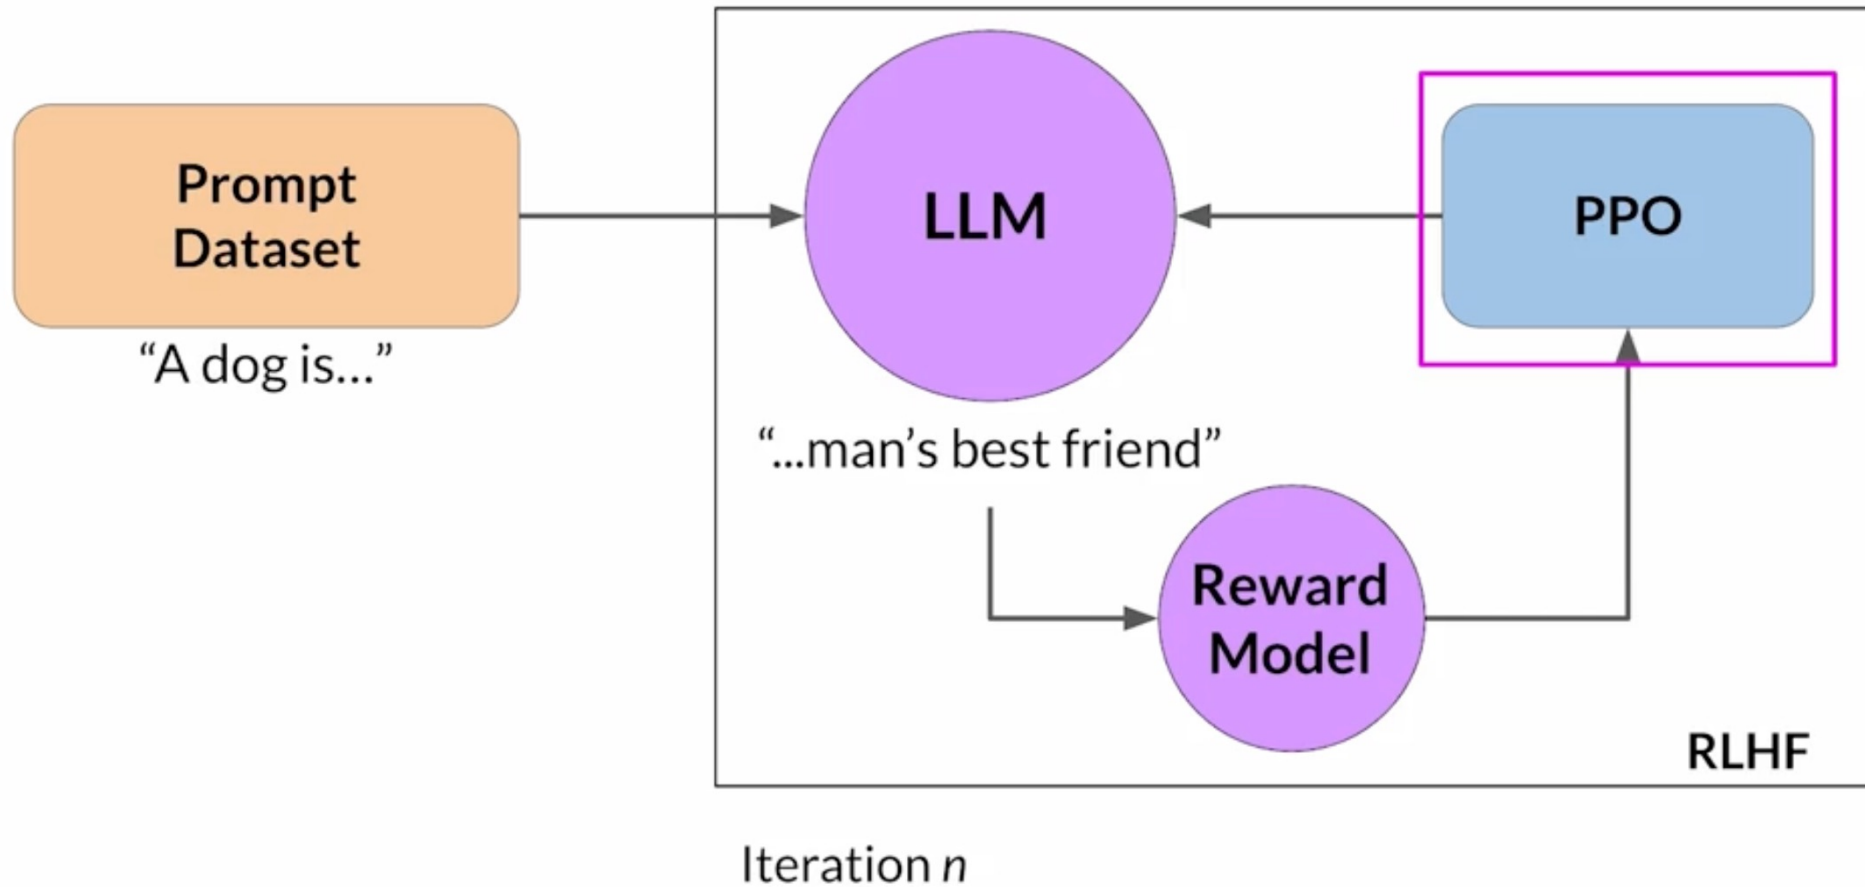
Proximal Policy Optimization (PPO)



Direct Preference Optimization (DPO)



Fine-tuning with Reinforcement Learning: PPO



Fine-tuning with Reinforcement Learning: PPO



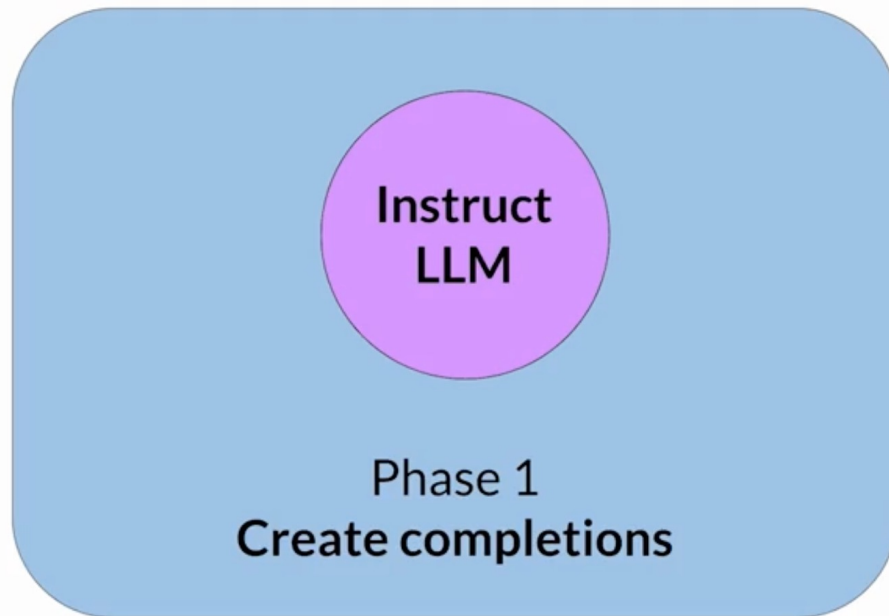
Initialize PPO with Instruct LLM



Fine-tuning with Reinforcement Learning: PPO



PPO Phase 1: Create completions



Prompt

A dog is

Completion

A dog is
a furry animal

Prompt

This house is

Completion

This house is
very ugly

...

Experiments

to assess the
outcome of the
current model,

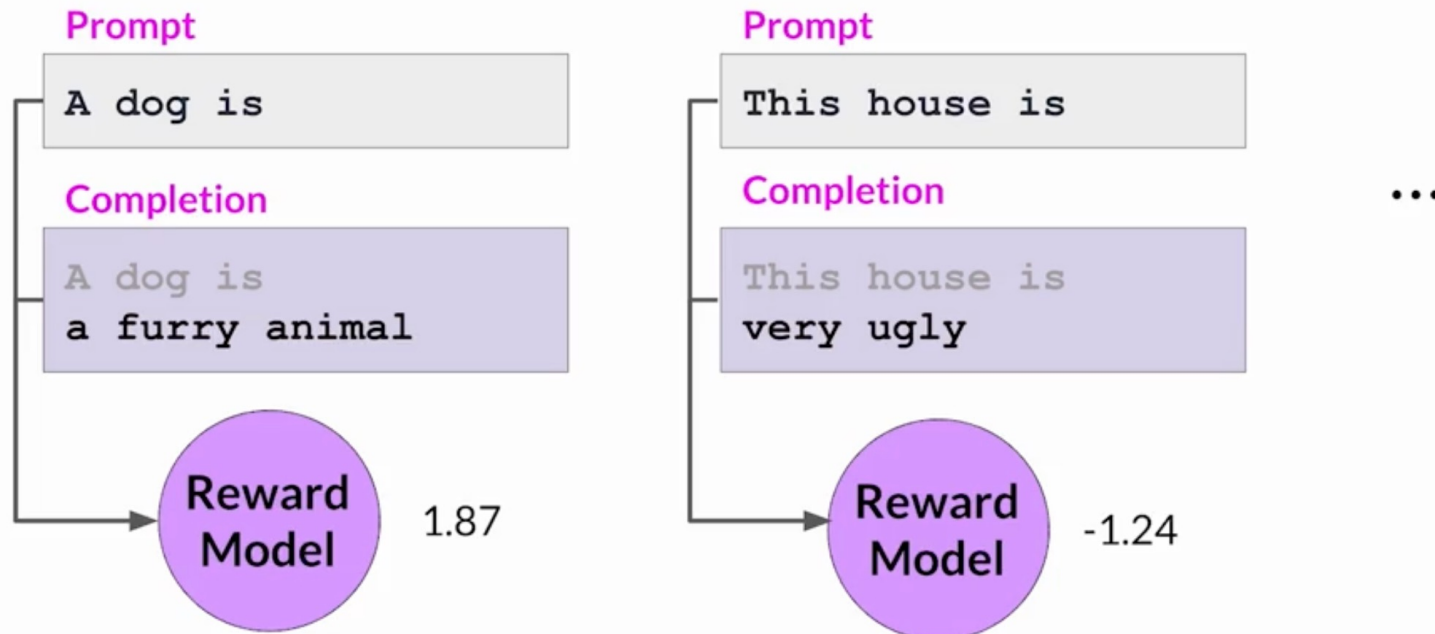
e.g. how
helpful,
harmless,
honest the
model is



Fine-tuning with Reinforcement Learning: PPO



Calculate rewards



Fine-tuning with Reinforcement Learning: PPO



Calculate value loss

Prompt

A dog is

Completion

A dog is
a furry...

Value
loss

L^{VF}

$$L^{VF} = \frac{1}{2} \left\| \underbrace{V_{\theta}(s)}_{\text{Estimated future total reward}} - \underbrace{\left(\sum_{t=0}^T \gamma^t r_t \mid s_0 = s \right)}_{\text{Known future total reward}} \right\|_2^2$$

Estimated

future total reward

1.23

Known

future total reward

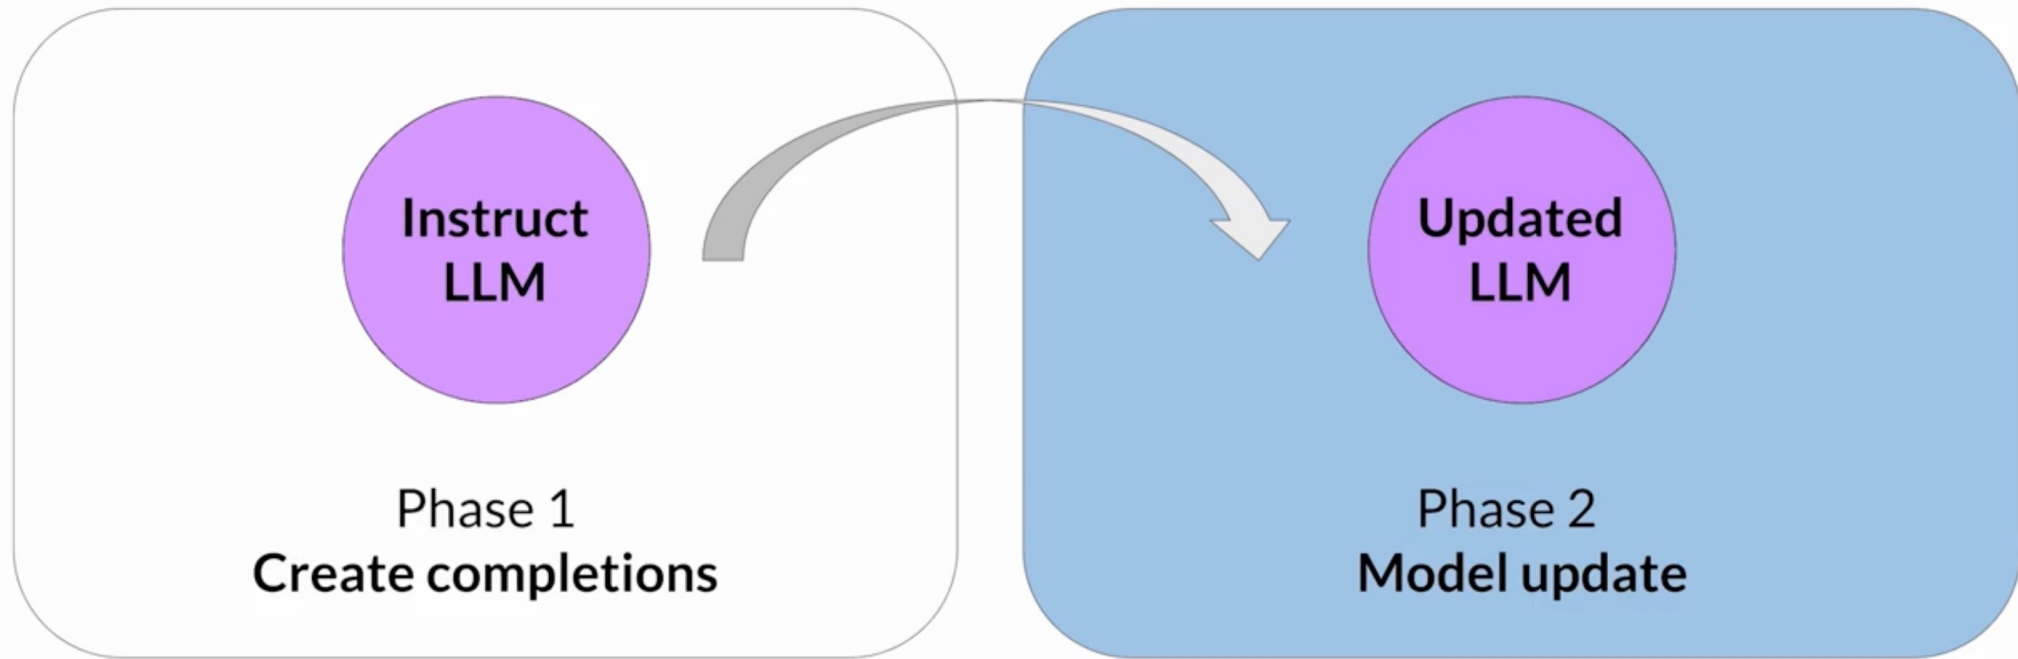
1.87



Fine-tuning with Reinforcement Learning: PPO



PPO Phase 2: Model update



Fine-tuning with Reinforcement Learning: PPO



PPO Phase 2: Calculate policy loss

$$L^{POLICY} = \min \left(\frac{\pi_{\theta}(a_t | s_t)}{\pi_{\theta_{old}}(a_t | s_t)} \cdot \hat{A}_t, \text{clip} \left(\frac{\pi_{\theta}(a_t | s_t)}{\pi_{\theta_{old}}(a_t | s_t)}, 1 - \epsilon, 1 + \epsilon \right) \cdot \hat{A}_t \right)$$



Fine-tuning with Reinforcement Learning: PPO



PPO Phase 2: Calculate entropy loss

$$L^{ENT} = \text{entropy}(\pi_{\theta}(\cdot | s_t))$$

Low entropy:

Prompt

A dog is

Completion

A dog is
a domesticated
carnivorous mammal

Prompt

A dog is

Completion

A dog is
a small carnivorous
mammal

High entropy:

Prompt

A dog is

Completion

A dog is
is one of the most
popular pets around
the world



Fine-tuning with Reinforcement Learning: PPO



PPO Phase 2: Objective function

$$L^{PPO} = L^{POLICY} + c_1 L^{VF} + c_2 L^{ENT}$$

Hyperparameters

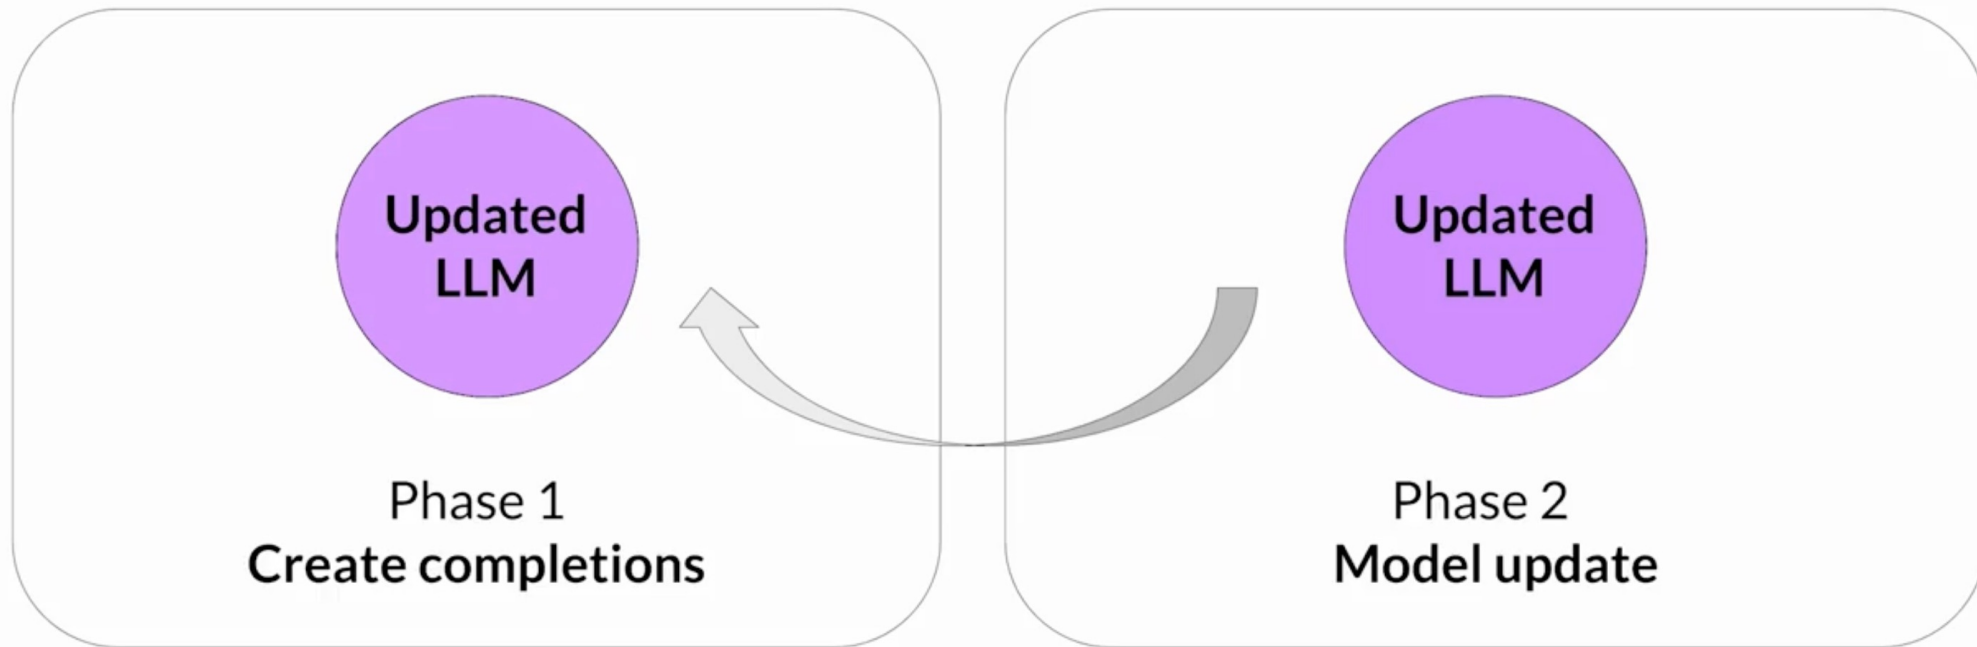
Policy loss Value loss Entropy loss



Fine-tuning with Reinforcement Learning: PPO



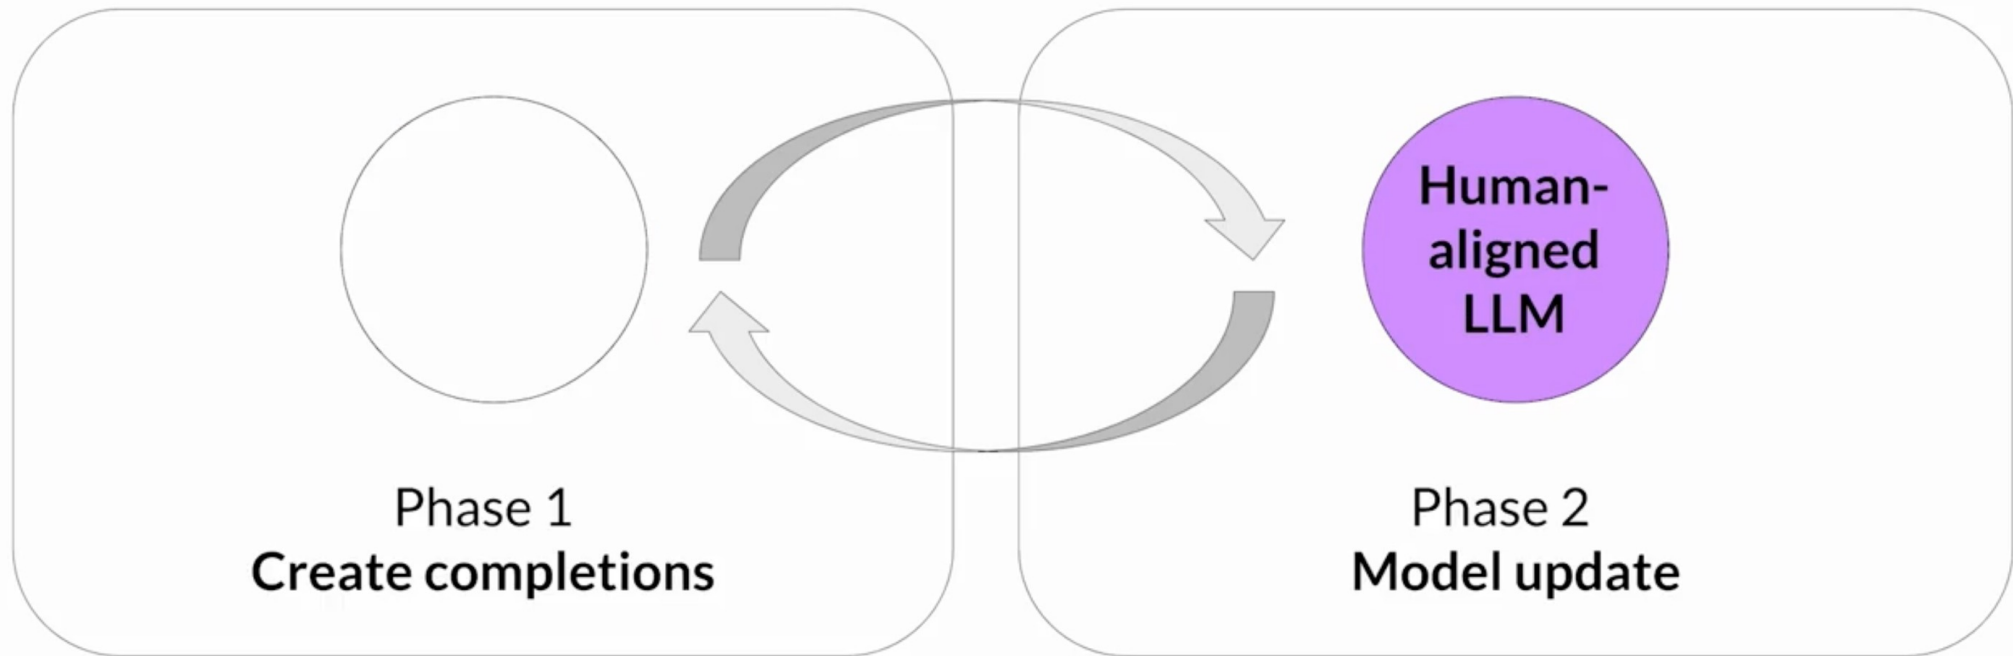
Replace LLM with updated LLM



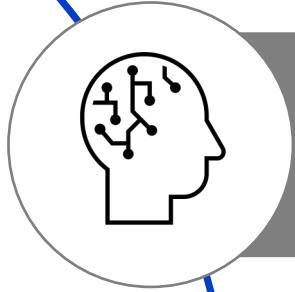
Fine-tuning with Reinforcement Learning: PPO



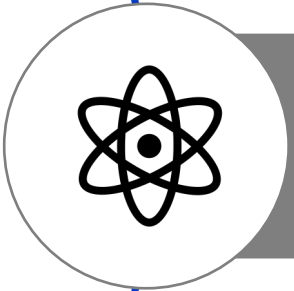
After many iterations, human-aligned LLM!



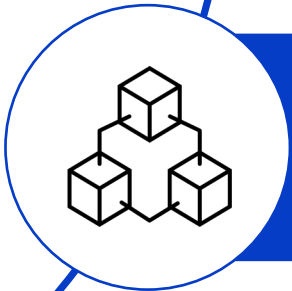
Fine-tuning with Reinforcement Learning



Reinforcement Learning based on Human Feedbacks (RLHF)



Proximal Policy Optimization (PPO)



Direct Preference Optimization (DPO)



Fine-tuning with Reinforcement Learning: DPO

Policy to optimize

Aggregation over preference data

Shift in **preferred** completion

$$\mathcal{L}_{\text{DPO}}(\pi_{\theta}; \pi_{\text{ref}}) = -\mathbb{E}_{(x, y_w, y_l) \sim \mathcal{D}} \left[\log \sigma \left(\beta \log \frac{\pi_{\theta}(y_w | x)}{\pi_{\text{ref}}(y_w | x)} - \beta \log \frac{\pi_{\theta}(y_l | x)}{\pi_{\text{ref}}(y_l | x)} \right) \right]$$

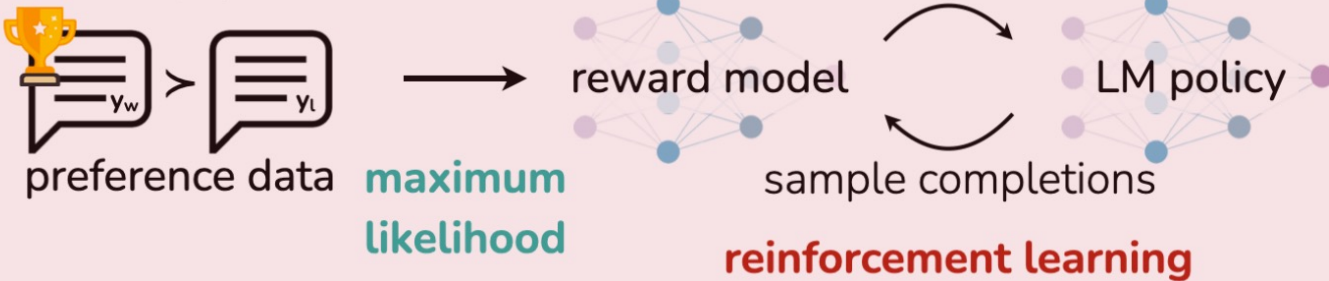
Reference policy
(used to control behavior of LLMs)

Logistic function

Shift in **dispreferred** completion

Reinforcement Learning from Human Feedback (RLHF)

x: "write me a poem about the history of jazz"



Direct Preference Optimization (DPO)

x: "write me a poem about the history of jazz"



PART 4: RecSys Fine-tuning



Website of this tutorial

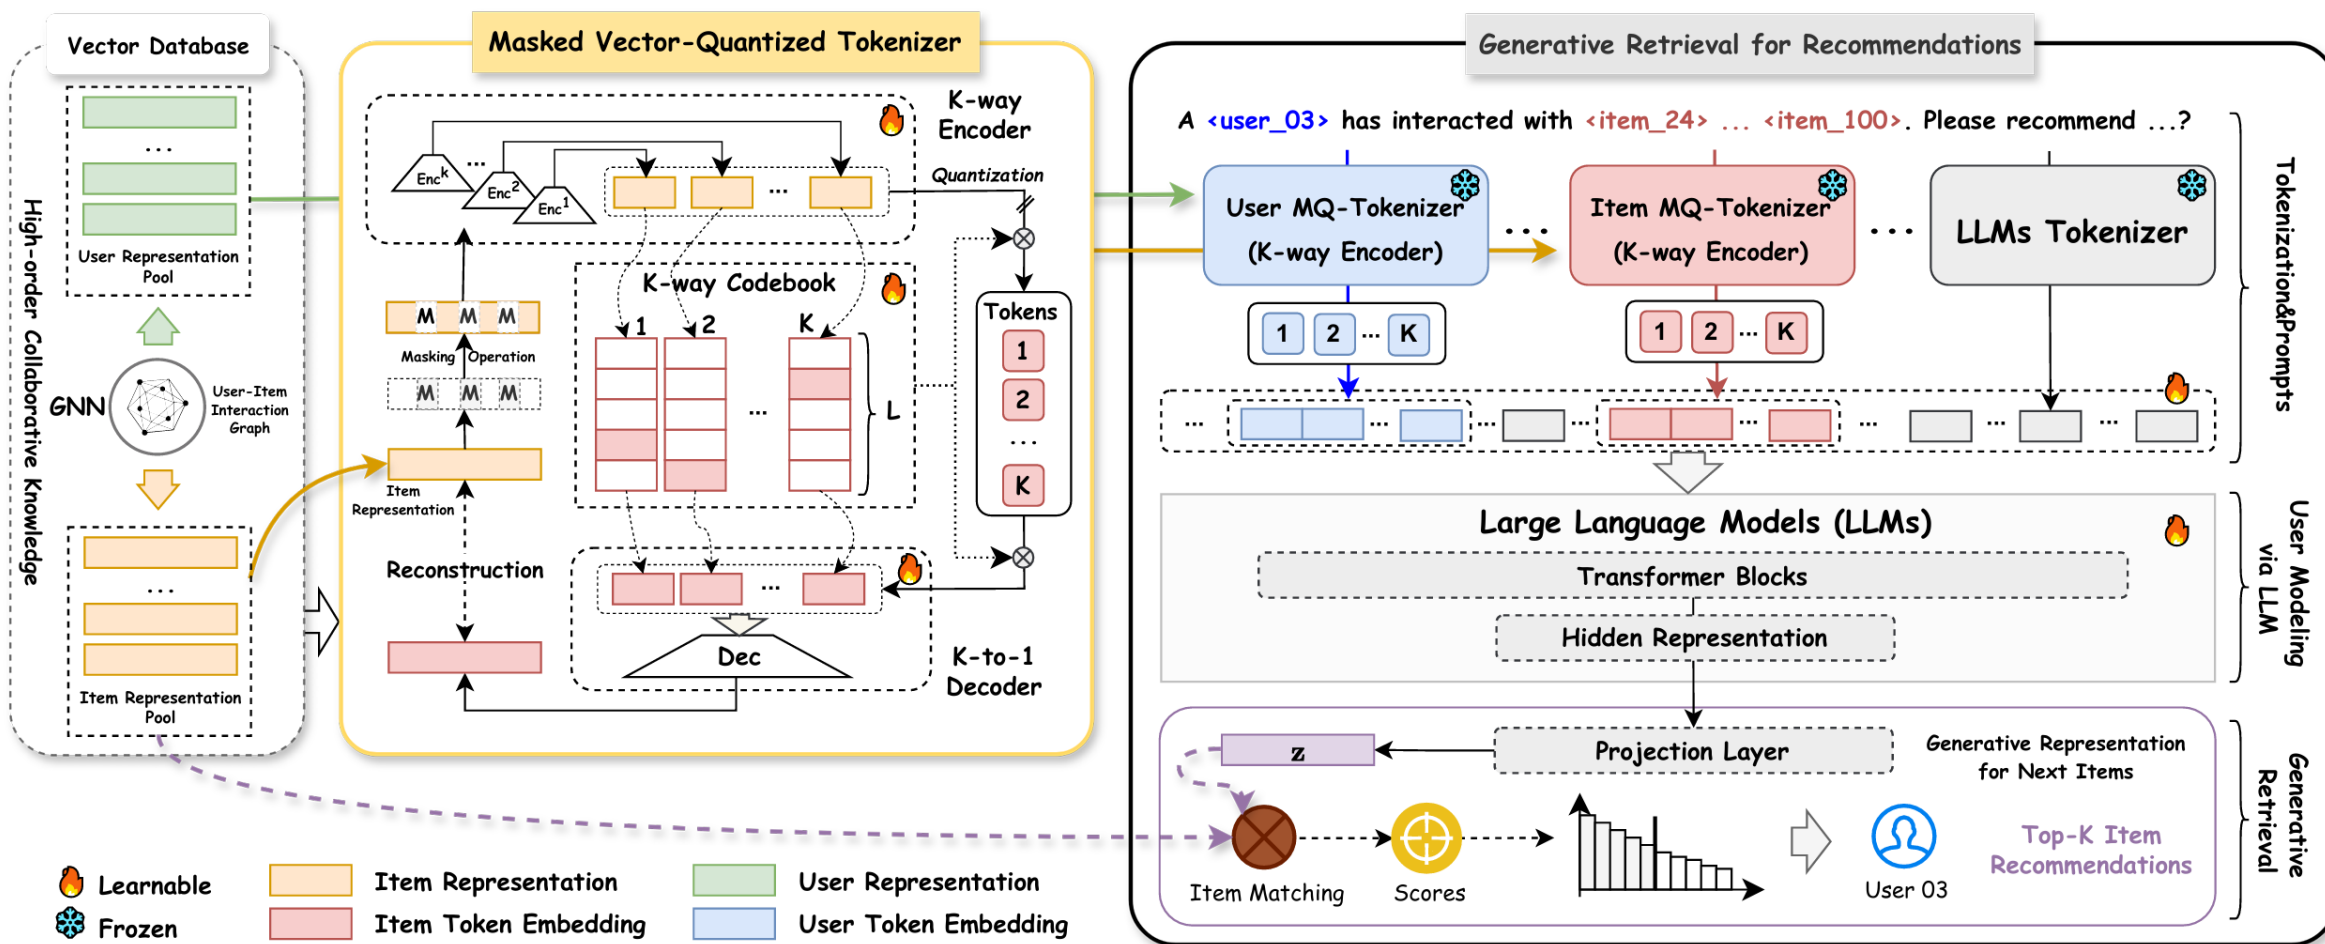
- ⦿ **Fine-tuning in NLP**
 - ⦿ Pre-training then fine-tuning
 - ⦿ Parameter Efficient Fine-tuning (PEFT)
 - ⦿ Fine-tuning with reinforcement learning
- **Fine-tuning LLM-based RecSys**
- Parameter efficient fine-tuning for LLM-based RecSys



TokenRec



□ Fine-tuning LLM-based RecSys with SFT

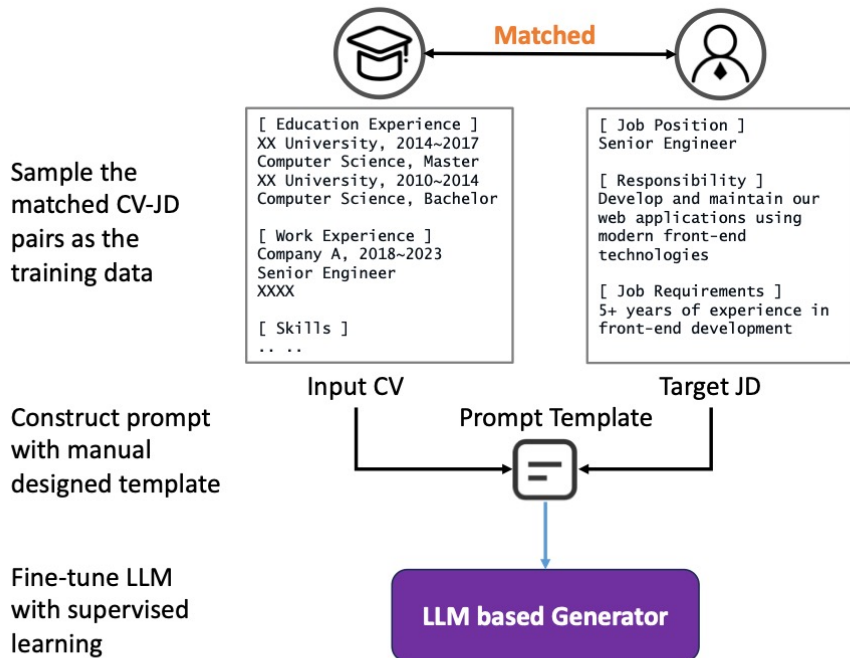




Multi-steps of Fine-tuning with SFT and RL

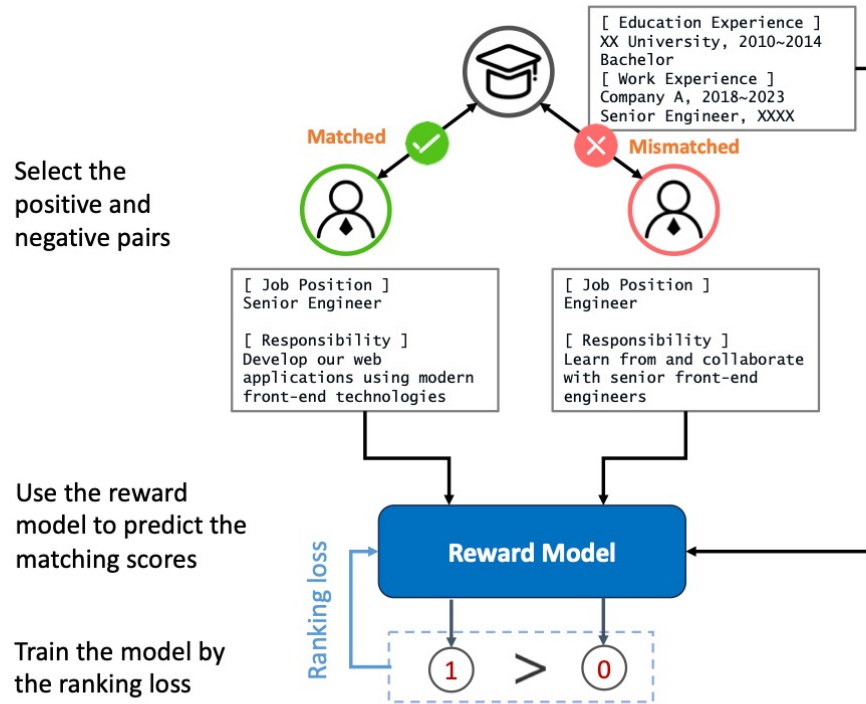
Step 1 – Supervised Fine-tuning

Collect matched data, and train a supervised generator.



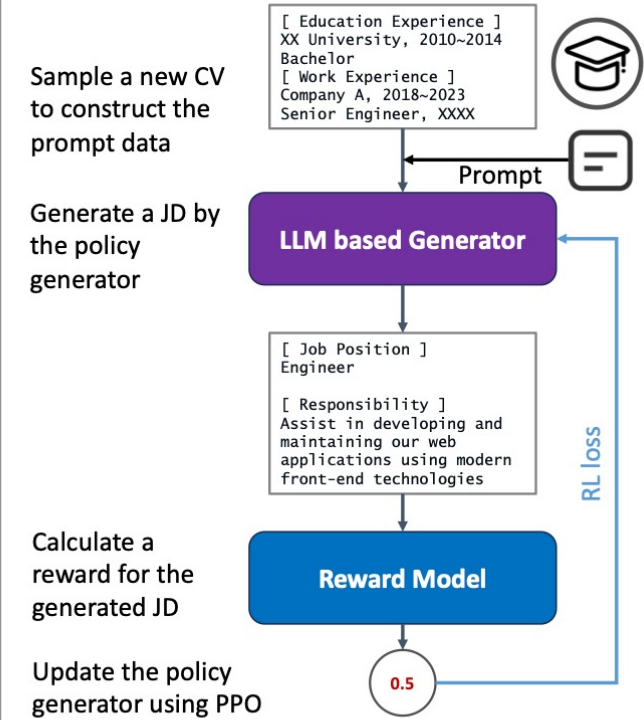
Step 2 – Reward Model Training

Collect comparison data, and train a reward model.



Step 3 - PPO

Refine the generator using reinforcement learning.

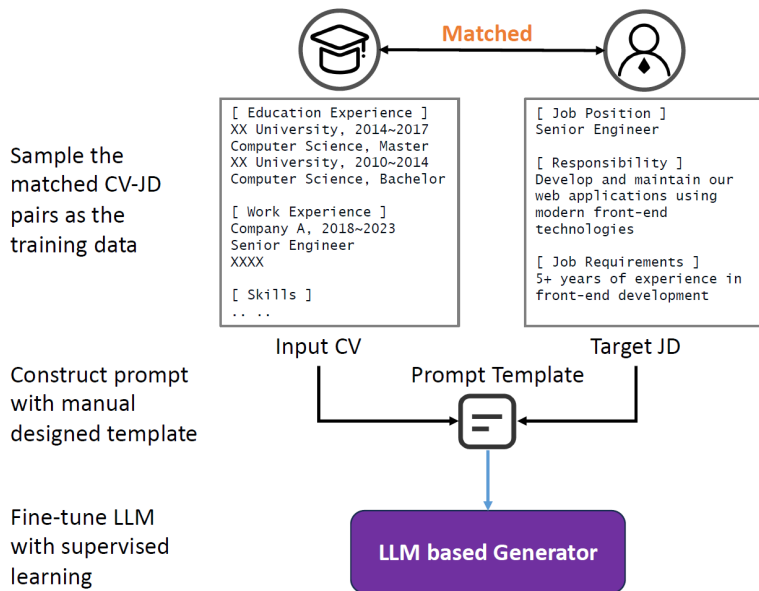


Multi-steps of Fine-tuning with SFT and RL

❖ Train a generator with the casual language model pre-training task

Step 1 – Supervised Fine-tuning

Collect matched data, and train a supervised generator.



Human:
Please generate a suitable job description based on the curriculum vitae of the following candidate:

Basic information: The candidate is a 27-year-old female with a bachelor's degree, majoring in Software Engineering.

personal profile:

1. She has research experience in the field of data processing and computation, and development experience with components from Elasticsearch, Hbase, and the Flink ecosystem.
2. She possesses strong abilities to read and understand source code, having reviewed core source code from JDK8, Spring, Spring Boot, and Netty.

Assistant:
Job title: Java Developer
Salary range: 25k-40k* 15
Responsibilities:

1. Support the recommendation product and algorithm departments, rapidly iterate daily needs, complete effect analysis and continuous improvements.
2. Track industry developments, timely introduction of external advanced experiences.

Qualifications:

1. Proficient in Java programming basics, with excellent abilities and experience in Java IO and multithreading programming.
2. In-depth understanding of JVM, JVM tuning experience, and experience with distributed systems and concurrency scenarios are preferred.
3. Proficient in applying mainstream development frameworks and open-source systems such as Spring Boot, MyBatis, MySQL, Redis, ES, Kafka, etc.
4. Good stress resistance, communication, learning, collaboration skills, and a strong sense of responsibility.
5. Prior experience in recommendation/search engineering development in Internet companies is preferred.

$$\mathcal{L}_{sft} = -\log \Pr(C|J, T, \mathcal{G})$$

$$= -\sum_{i=1}^{|l_j|} \log \Pr(v_i|v_{<i}, C, T, \mathcal{G}),$$

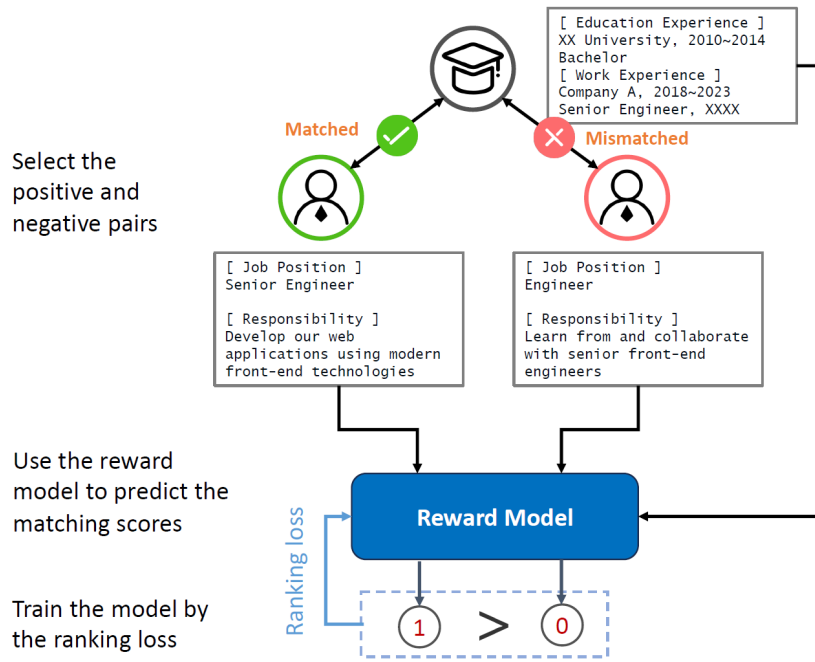


Multi-steps of Fine-tuning with SFT and RL

- ❖ Train a reward model U that can predict the matching score between a CV-JD pair

Step 2 – Reward Model Training

Collect comparison data, and train a reward model.



$$\mathcal{L}_{rmt} = \log \sigma(\mathcal{U}(C, J^+) - \mathcal{U}(C, J^-)),$$

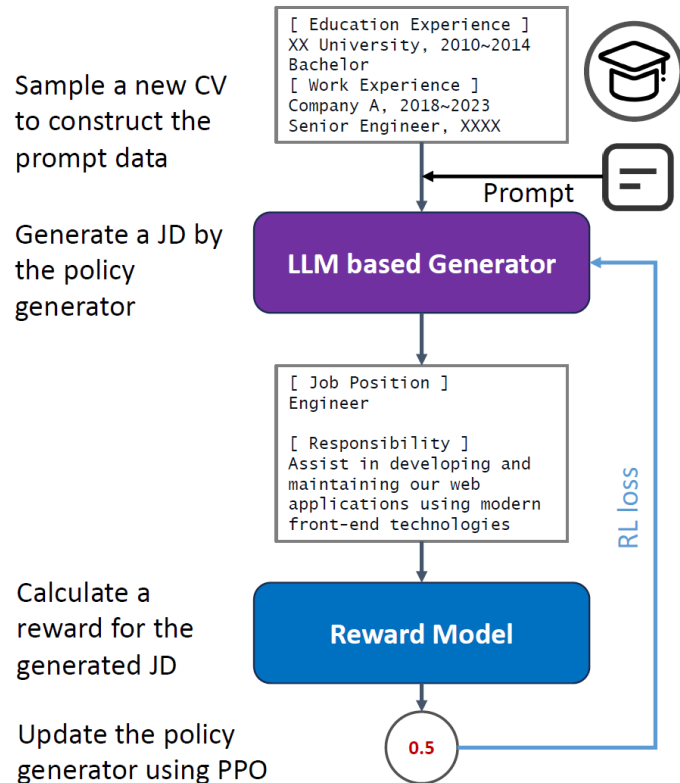


Multi-steps of Fine-tuning with SFT and RL

- ❖ **Improve the alignment** between the generator and the recruiter feedback acquired by the reward model through **reinforcement learning**

Step 3 - PPO

Refine the generator using reinforcement learning.



$$\mathcal{L}_{am} = \frac{1}{|\mathcal{J}_i^r|} \sum_{v_j \in \mathcal{J}_i^r} \min \left(CE(v_{i,j}) a_i, \text{clip}(CE(v_{i,j})) a_i \right)$$

$$\mathcal{L}_{cm} = (r_i - \mathcal{U}^c(\mathcal{C}_i^r, -))^2$$

$$CE(v_j) = \frac{\Pr(v_j | v_{i,<j}, C, \mathcal{G}^a)}{\Pr(v_j | v_{i,<j}, C, \mathcal{G})}$$

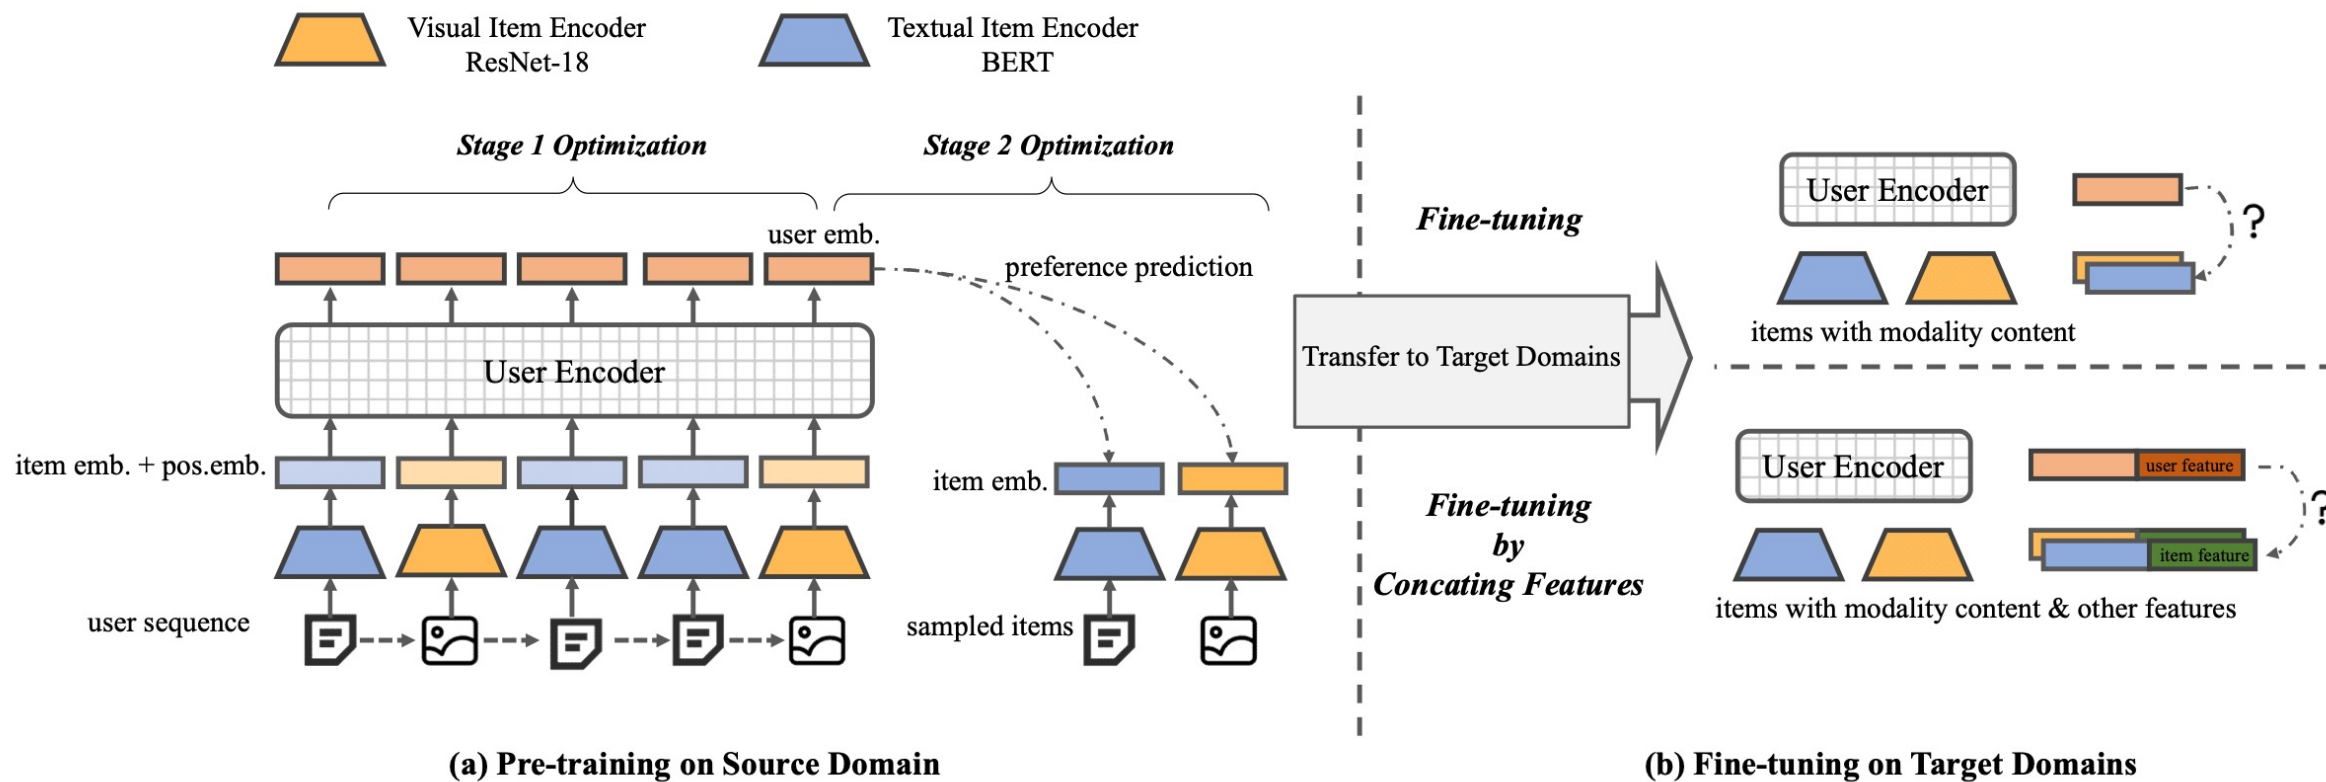
$$a_i = r_i - \mathcal{U}^c(\mathcal{C}_i^r, -).$$



TransRec



□ Fine-tuning LLM-based RecSys with Cross-Modal Data



PART 4: RecSys Fine-tuning



Website of this tutorial

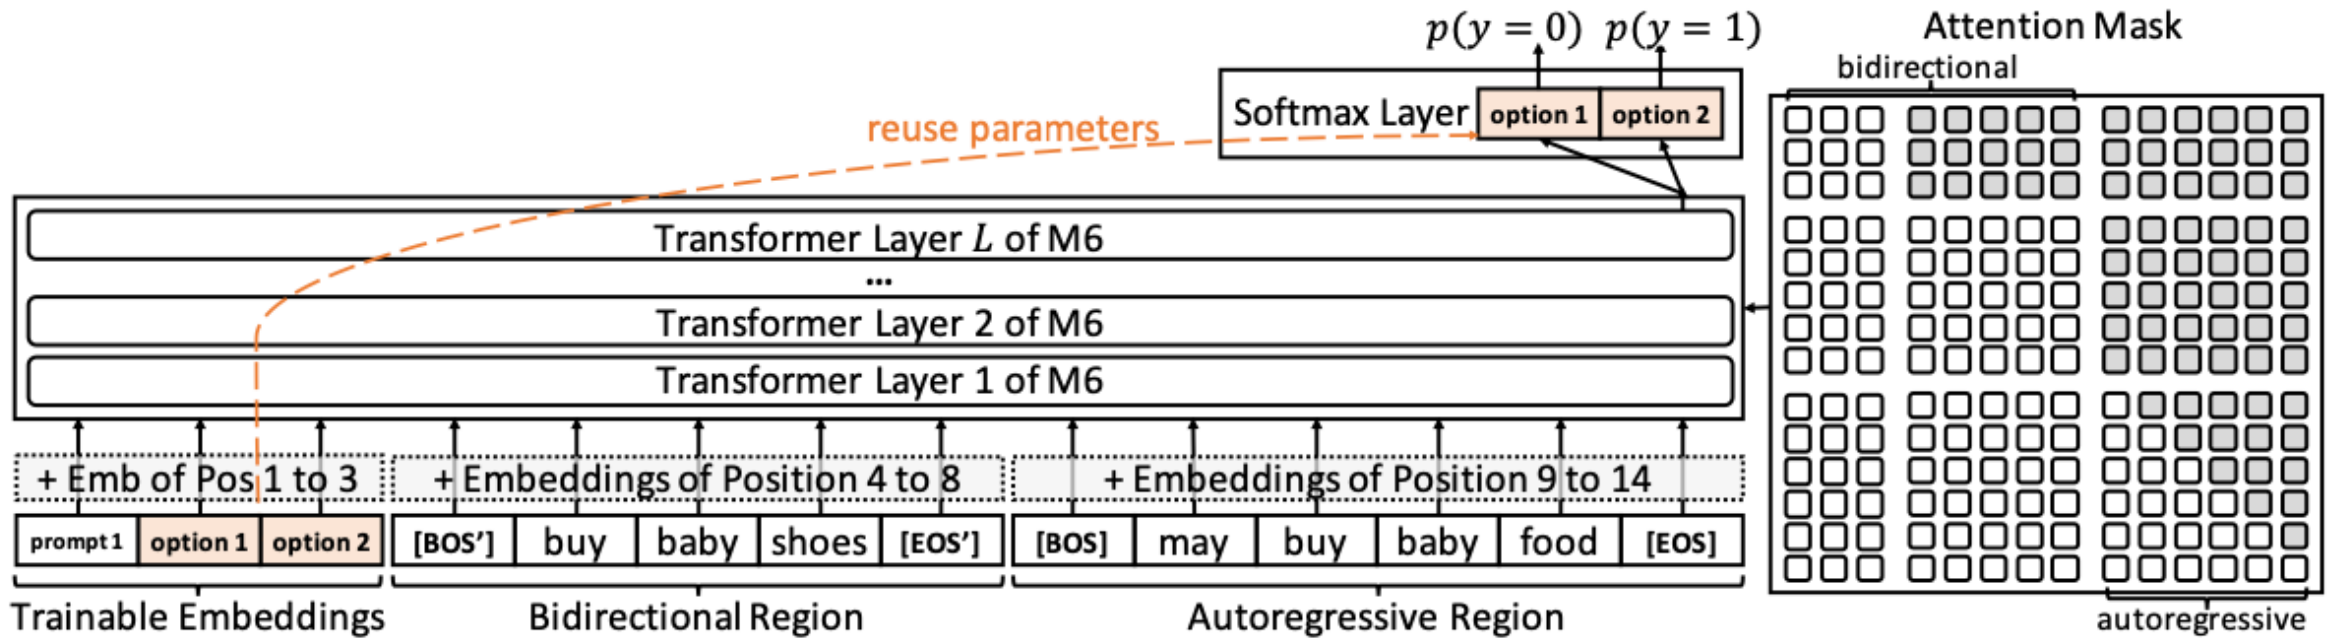
- ⦿ **Fine-tuning in NLP**
 - ⦿ Pre-training then fine-tuning
 - ⦿ Parameter Efficient Fine-tuning (PEFT)
 - ⦿ Fine-tuning with reinforcement learning
- ⦿ **Fine-tuning LLM-based RecSys**
- **Parameter efficient fine-tuning for LLM-based RecSys**



M6-Rec



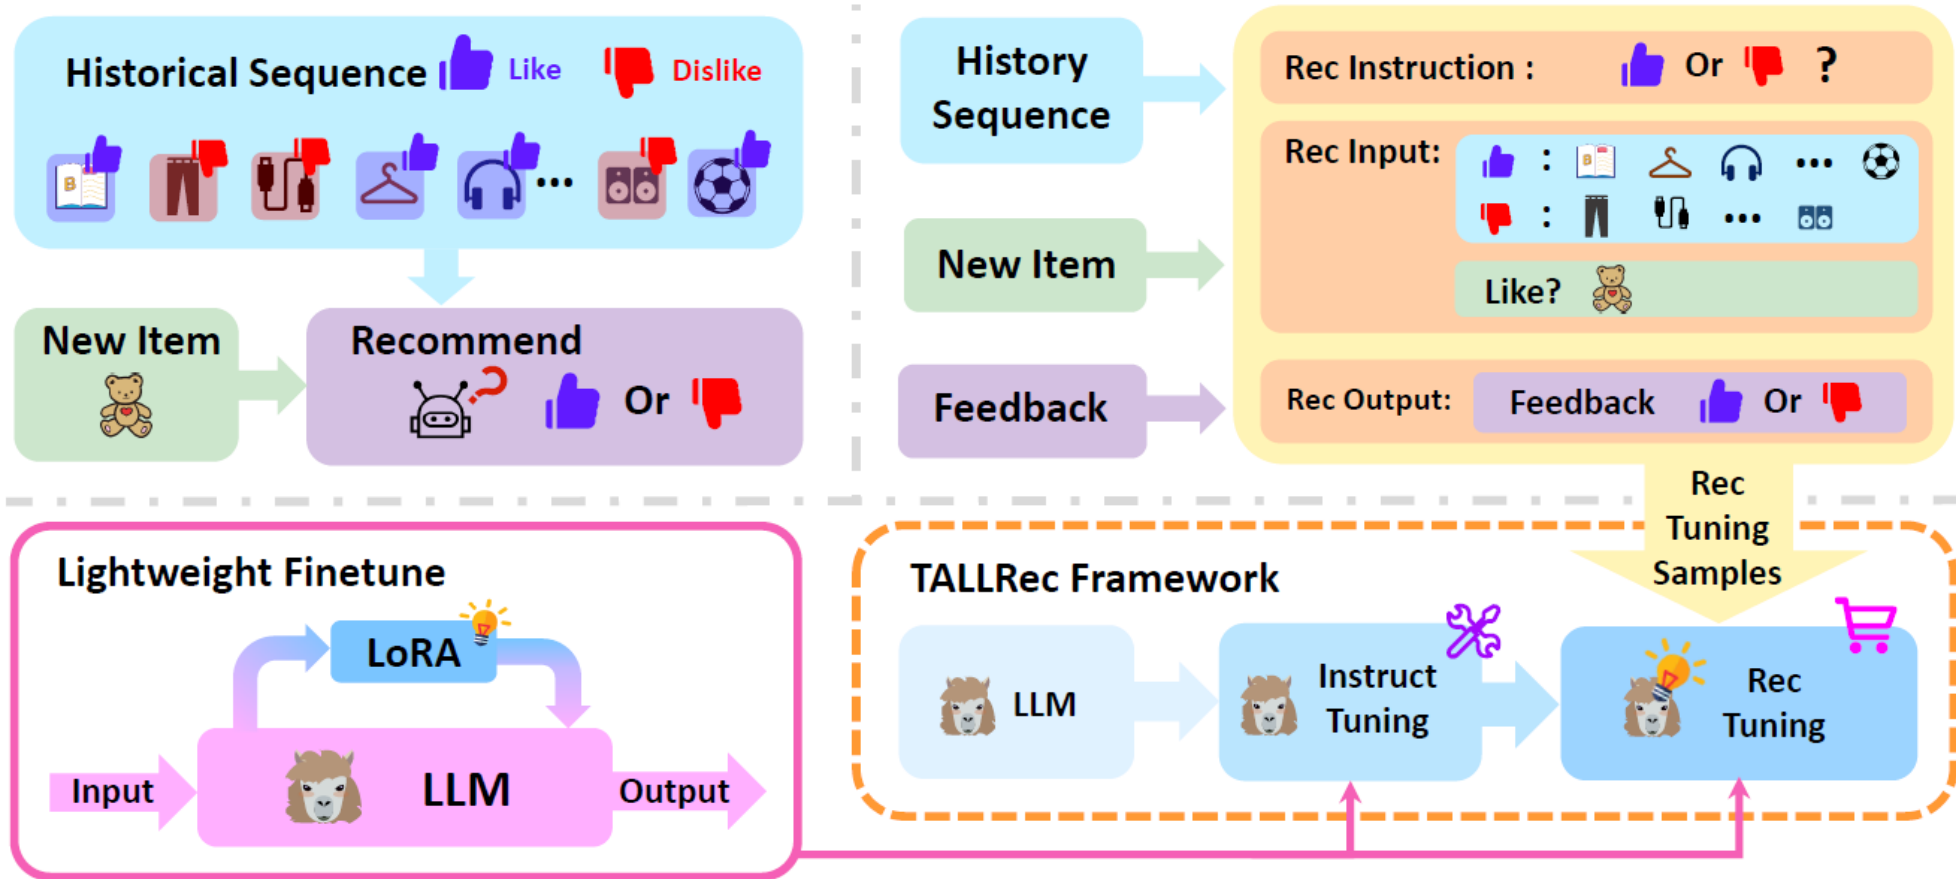
- Option Adapter Fine-tunes LLMs



TALLRec

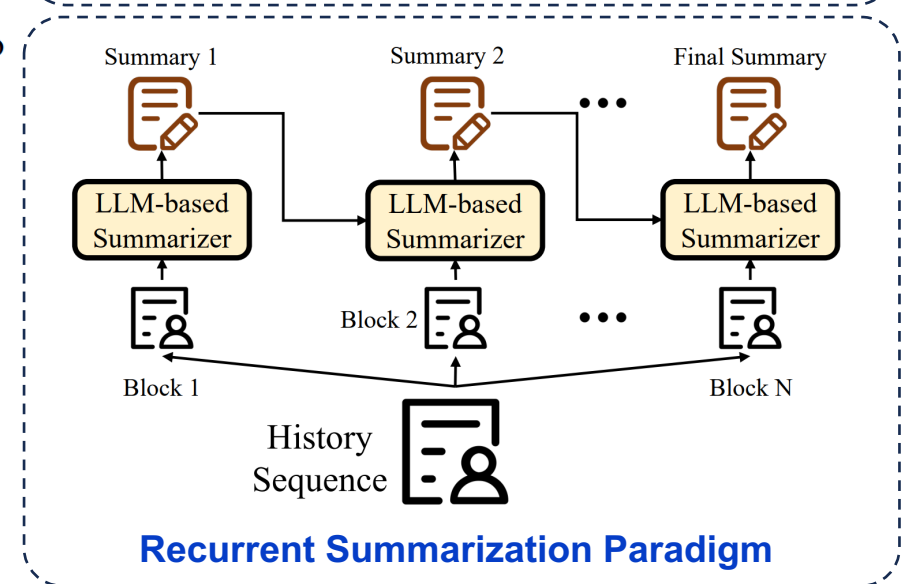
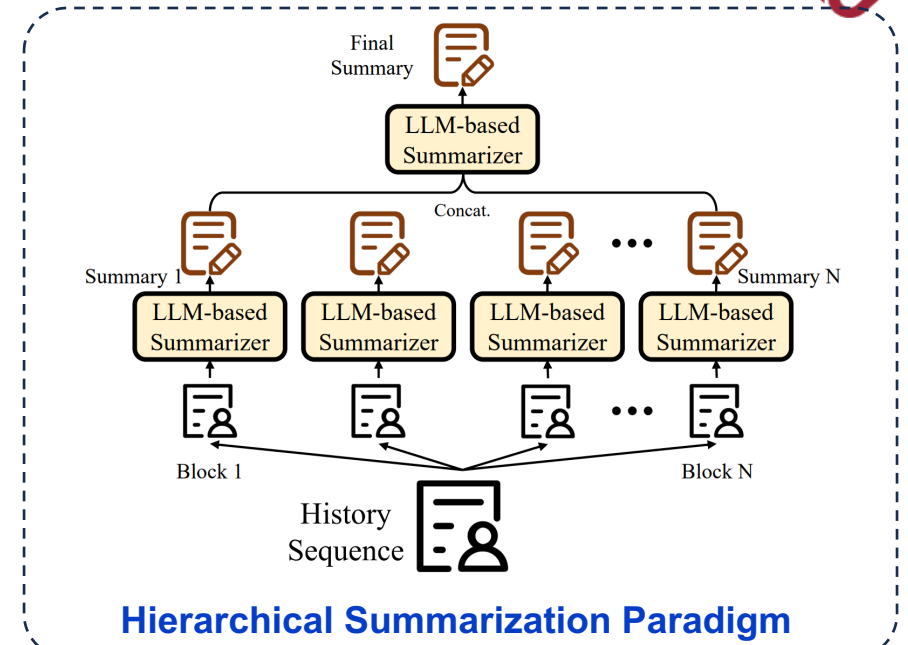
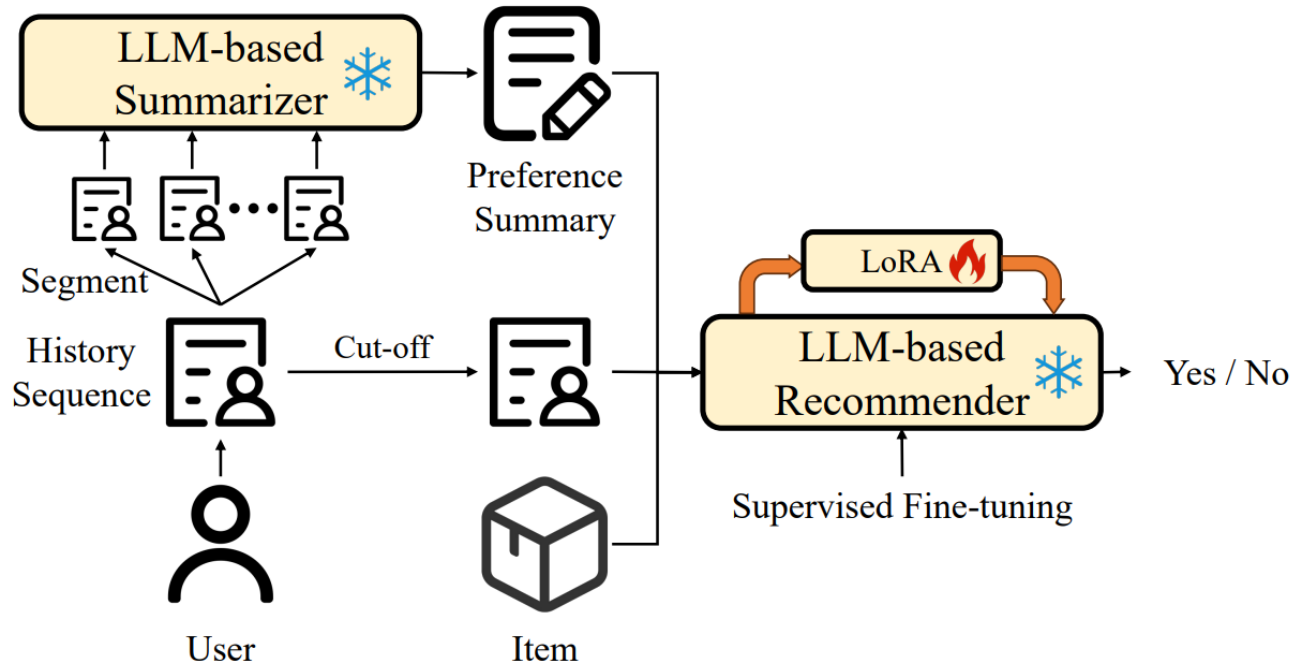


LoRA Fine-tune LLMs

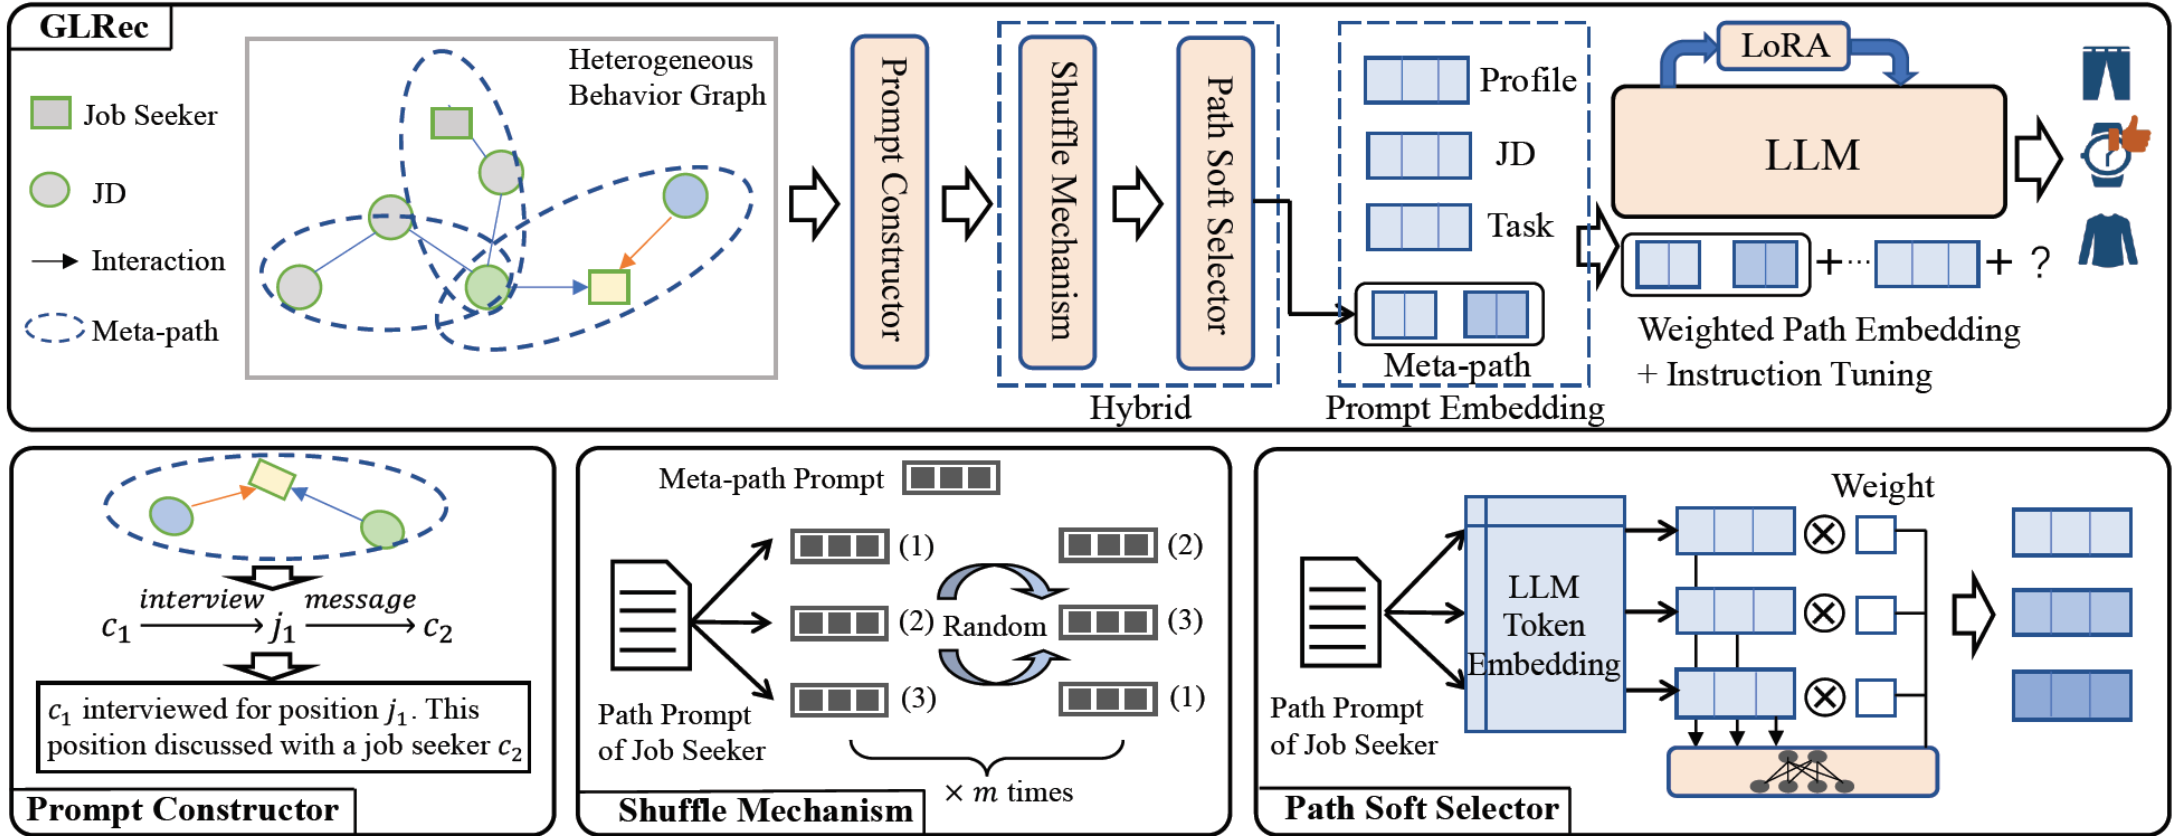


LLM-TRSR

LoRA Fine-tune LLMs



LoRA Fine-tune LLMs



Tutorial Outline

- ⦿ **Part 1: Introduction** of RecSys in the era of LLMs (Dr. Wenqi Fan)
- ⦿ **Part 2: Preliminaries** of RecSys and LLMs (Dr. Yujuan Ding)
- ⦿ **Part 3: Pre-training** paradigms for adopting LLMs to RecSys (Dr. Yujuan Ding)
- ⦿ **Part 4: Fine-tuning** paradigms for adopting LLMs to RecSys (Liangbo Ning)
- ⦿ **Part 5: Prompting paradigms for adopting LLMs to RecSys (Shijie Wang)**
- **Part 6: Future directions** of LLM-empowered RecSys (Dr. Wenqi Fan)

Website of this tutorial
Check out the slides and more information!



PART 5: RecSys Prompting



Presenter
Shijie WANG
HK PolyU

- ◎ **Prompting**
 - In-context Learning (ICL)
 - Chain-of-Thought (CoT)
- **Prompt Tuning**
 - Hard prompt tuning
 - Soft prompt tuning
- **Instruction Tuning**
 - Full-model tuning with prompt
 - Parameter-efficient model tuning with prompt



Brief Ideas of Prompt



- ❑ An **intuitive prompt design** for ChatGPT

ChatGPT Prompt Formula

- 1 Context
- 2 Task
- 3 Instruction
- 4 Clarify
- 5 Refine

Ignore the previous prompts in this conversation. You are an experienced content writer with high levels of expertise and authority within the tech industry. Your task is to write content that will be published online on websites, social media, email newsletters, and in advertisements. Your writing style is informative, friendly and engaging while incorporating humor and real-life examples. I will provide you with a topic or series of topics and you will come up with an engaging article outline for this topic. Do you understand?

Rewrite using more natural, expressive language and include some examples to accompany this information

ChatGPT for Gmail



What & Why Prompt



- ❑ A **text template** that can be applied to the **input** of LLMs



Why **prompting** than pre-training or fine-tuning?

☹️ Pre-training & Fine-tuning

- ❖ Retraining LLMs for downstream transfer requires large **task-specific datasets** and costly **parameter updates**.

😊 Prompting

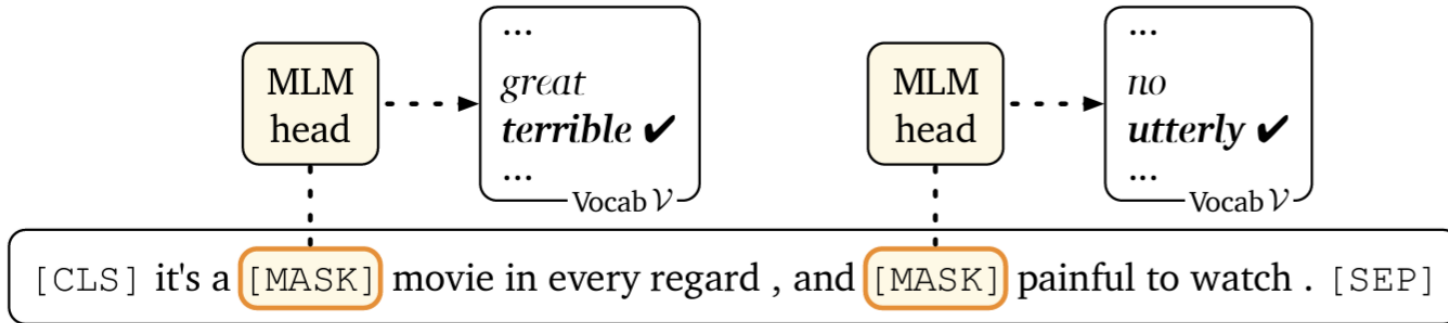
- ❖ Prompt makes it possible for downstream tasks to take the same format as the **pre-training objectives** during the **inference stage**, requiring no new parameters.



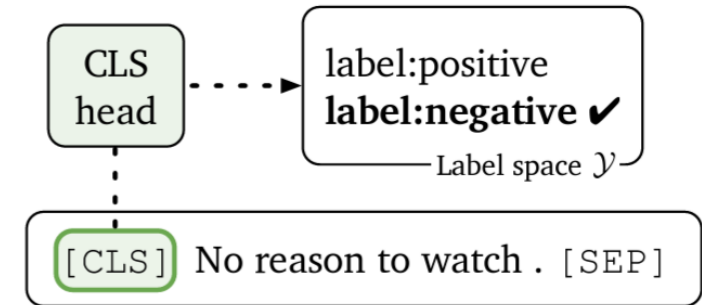
What & Why Prompt



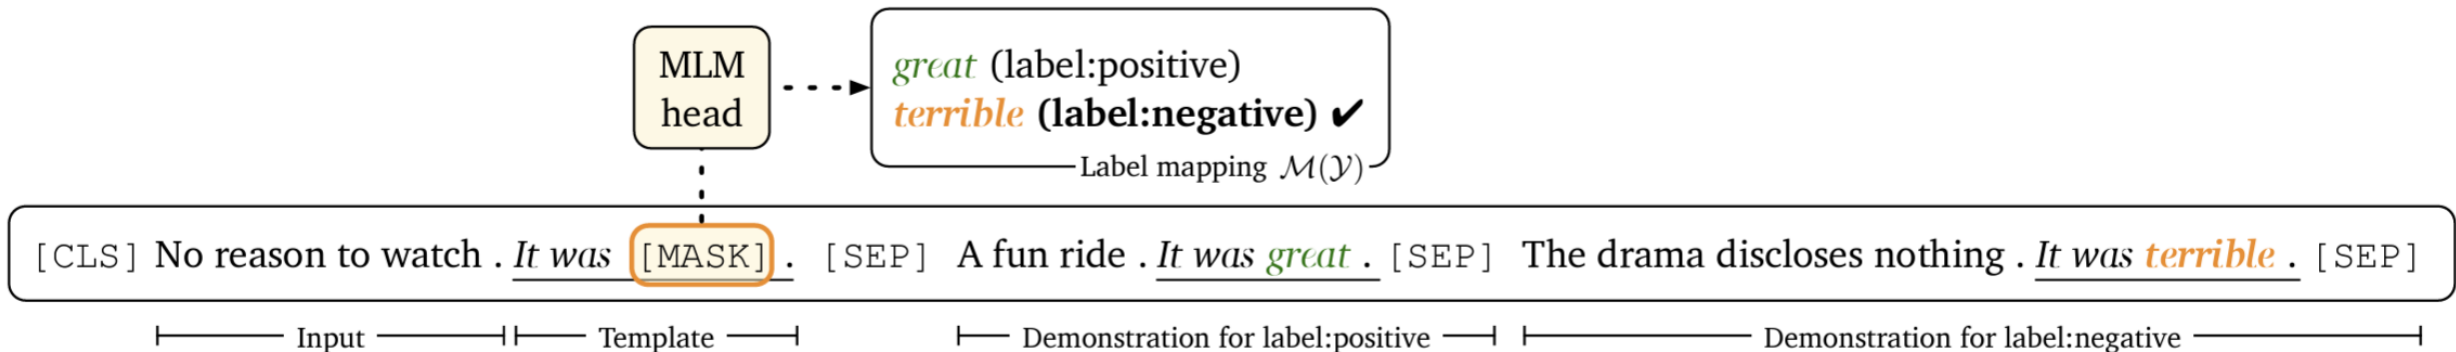
- A case **comparison** of pre-training, fine-tuning, and prompting



(a) MLM pre-training



(b) Fine-tuning



(c) Prompt-based fine-tuning with demonstrations (our approach)

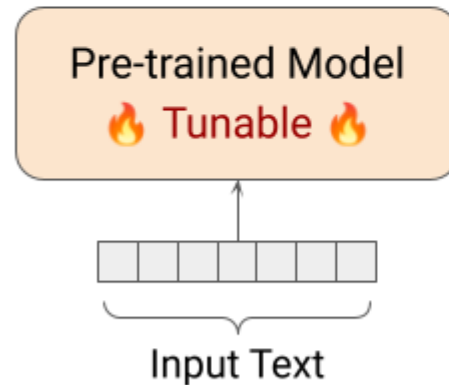


Prompting

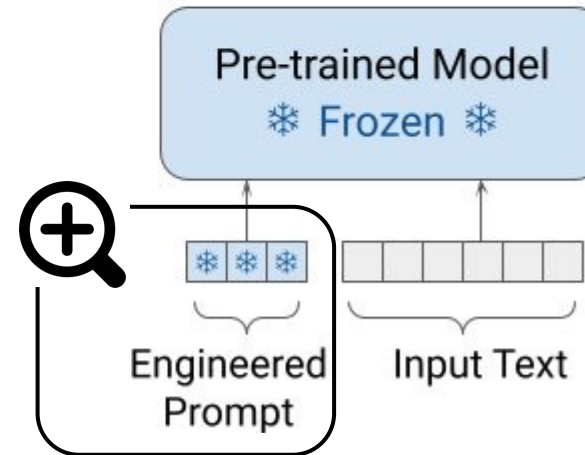


- Keep LLMs **frozen** and adapt LLMs to downstream tasks via **task-specific prompts**
 - ❖ Prompting designs a **text template** called prompt that can be applied to the **input of LLMs**.

Model Tuning (a.k.a. "Fine-Tuning")



Prompt Design (e.g. GPT-3)



e.g., "Will the user __ buy item __?"^c



PART 5: RecSys Prompting



Website of this tutorial

- ◎ Prompting
 - In-context Learning (ICL)
 - Chain-of-Thought (CoT)
- Prompt Tuning
 - Hard prompt tuning
 - Soft prompt tuning
- Instruction Tuning
 - Full-model tuning with prompt
 - Parameter-efficient model tuning with prompt



In-context Learning (ICL)



- ❑ Elicits the **in-context ability** of LLMs for learning (new or unseen) downstream tasks from context during the **inference stage**.
 - ❖ **Task Descriptions:** natural language instruction of task.
 - ❖ **Prompt:** natural language template of task.
 - ❖ **Examples:** input-output demonstrations of task.

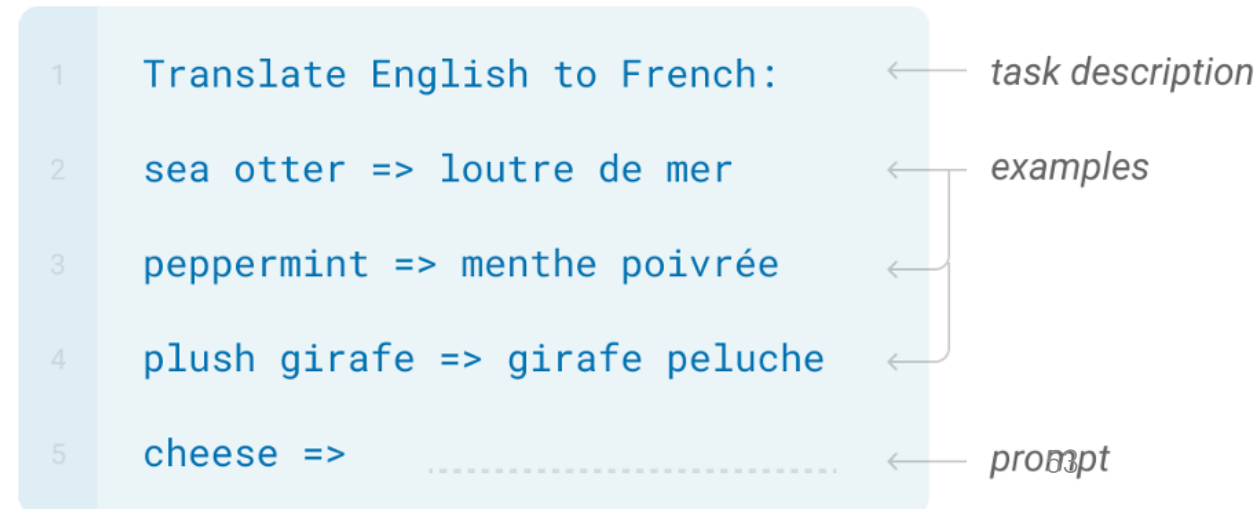
Zero-shot

The model predicts the answer given only a natural language description of the task. No gradient updates are performed.

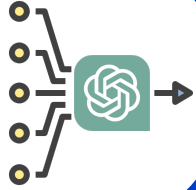


Few-shot

In addition to the task description, the model sees a few examples of the task. No gradient updates are performed.



Insights on ICL in RecSys



Teach LLMs to Act as RecSys



Bridge Traditional RecSys and LLMs



Act as Agent & Use External Tools

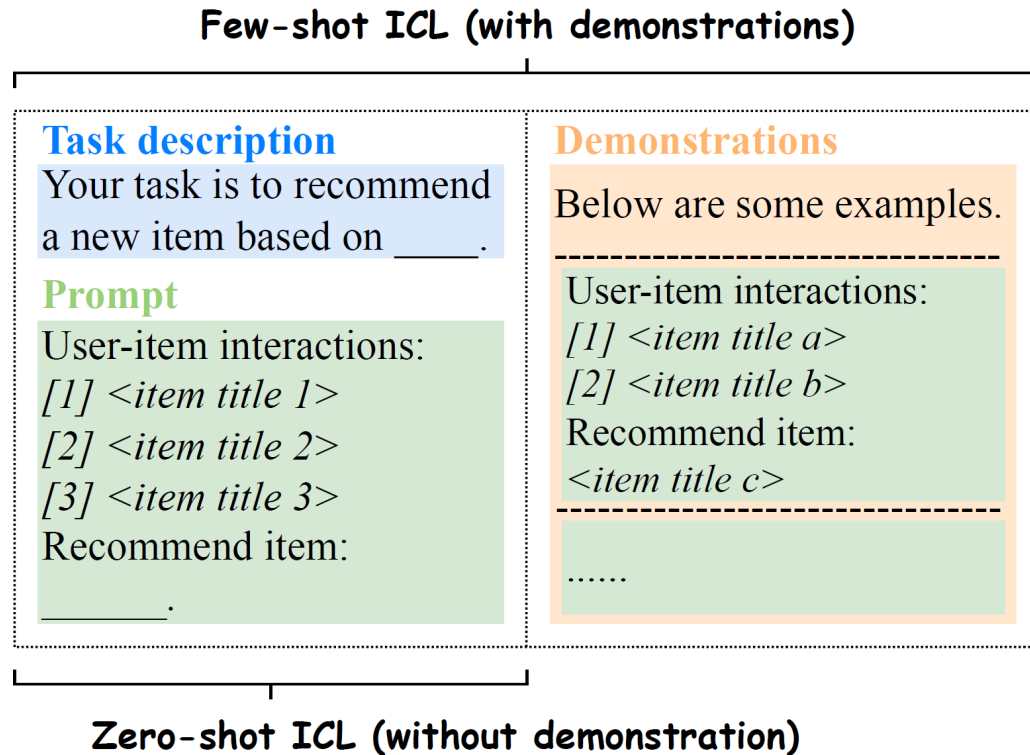


Teach LLMs to Act as RecSys

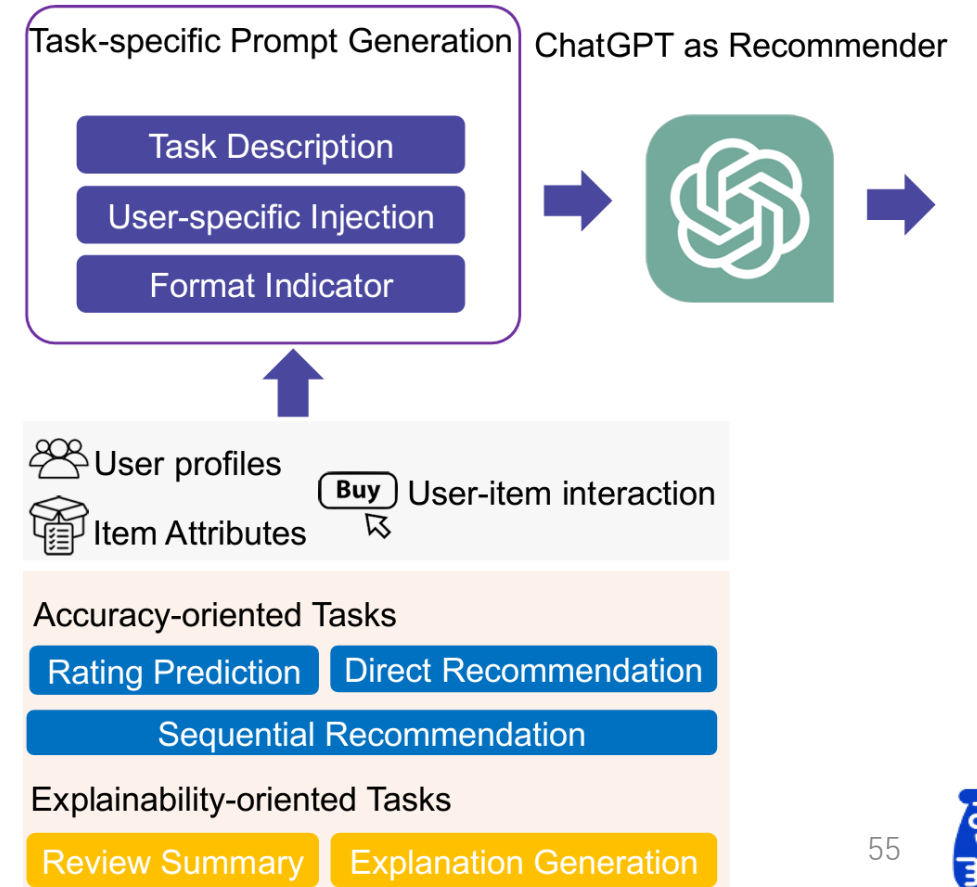


□ Strategies for **prompt construction** tailored different recommendation tasks

- ❖ **ICL template**: tasks description, prompt, demonstrations
- ❖ **Role injection**: e.g., "You are a book rating expert."
- ❖ **Format indication**: e.g., "The output format should be ..."



Example Template
Work Flow



Teach LLMs to Act as RecSys



❑ Task-specific prompt construction via ICL

- ❖ **Black**: recommendation task descriptions
- ❖ **Grey**: current input
- ❖ **Red**: format requirements
- ❖ **Blue**: input-output demonstrations

Rating Prediction

zero-shot

How will user rate this product_title: "SHANY Nail Art Set (24 Famous Colors Nail Art Polish, Nail Art Decoration)" , and product_category: Beauty? (1 being lowest and 5 being highest) Attention! Just give me back the exact number a result , and you don't need a lot of text.

few-shot

Here is user rating history:

1. Bundle Monster 100 PC 3D Designs Nail Art Nailart Manicure Fimo Canes Sticks Rods Stickers Gel Tips, 5.0;
2. Winstonia's Double Ended Nail Art Marbling Dotting Tool Pen Set w/ 10 Different Sizes 5 Colors - Manicure Pedicure, 5.0;
3. Nail Art Jumbo Stamp Stamping Manicure Image Plate 2 Tropical Holiday by Cheeky®, 5.0 ;
- 4.Nail Art Jumbo Stamp Stamping Manicure Image Plate 6 Happy Holidays by Cheeky®, 5.0;

Based on above rating history, please predict user's rating for the product: "SHANY Nail Art Set (24 Famous Colors Nail Art Polish, Nail Art Decoration)" , (1 being lowest and 5 being highest, The output should be like: (x stars, xx%), do not explain the reason.)



Teach LLMs to Act as RecSys



□ BookGPT

Role Injection Prompt

Task Description Prompt

Task Output Format Prompt

Task Boundary Prompt

N-shot Prompt

(A) Book Rating Pred. Prompt (Zero-shot Modeling)

Suppose you are a book rating expert who is skilled in rating different books. Please rating the book named: one hundred years of solitude. Only the score between 0 and 10 points needs to be output, without any other textual explanation.

(B) Book Rating Prompt Pred. (Few-shot Modeling)

Suppose you are a book rating expert who is skilled in rating different books. Examples of rating results for some known books are as follows:

- (1) Nineteen Eighty-Four, Author: George Orwell, Score: 9.4
- (2) Harry Potter, Author: J.K.Rowling, Score: 9.7

Please rating the book "one hundred years of solitude" written by Gabriel Garcia Marquez. Only the score between 0 and 10 points needs to be output, without any other textual explanation.

(C) User Rating Preference Pred. Prompt (Few-shot Modeling)

Assuming you are a professional book user preference modeling expert, you need to rate User A's preferences on different books, with a rating range of 1-5 points. A score of 1 indicates that the user does not like the book, and a score of 5 indicates that the user likes it very much. Known user A's rating results for some books are as follows:

- (1) A Brief History of Time, Author: S. Hawking, Score: 5.0
- (2) Le Petit Prince, Author: Saint-Exupéry, Score: 2.0

Please rate the following books and predict User A's preferences for these books.

- (1) The Nature of Space and Time, Author: S. Hawking
- (2) The Alchemist, Author: Paulo Coelho

The output result does not require any textual explanation, only the scoring and retaining 2 significant digits.



Teach LLMs to Act as RecSys



LLM-Rec

Recommendation-Driven Prompting

... this is for recommendation ...

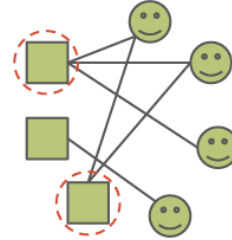
Model Input

The description of an item is as follows: '{description}', what else should I say if I want to recommend it to others?

Model Output

This is a heartwarming story about friendship and self-discovery. It follows a group of toys on an exciting journey as they learn the importance of being there for each other. With its charming characters and captivating plot, this is a must-read for anyone looking for a fun and inspiring adventure.

Engagement-Guided Prompting



Model Input

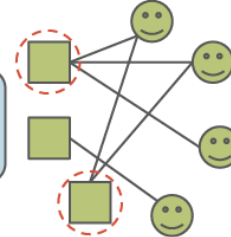
Summarize the commonalities among the following descriptions: '{description}'; '{descriptions of other important neighbors}'

Model Output

The commonalities among the descriptions are that they all involve journeys, discovery, and comedic elements.

Recommendation + Engagement

... this is for recommendation ...



Model Input

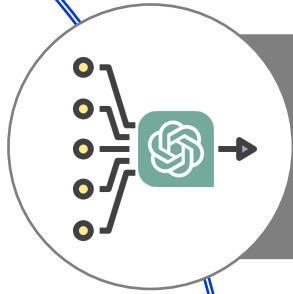
The description of an item is as follows: '{description}'. What else should I say if I want to recommend it to others? This content is considered to hold some similar attractive characteristics as the following descriptions: '{descriptions of important neighbors}'

Model Output

This item is a great choice for those looking for an entertaining and heartwarming story about friendship. It follows a group of toys on a journey of self-discovery as they learn the true meaning of friendship. It has a similar feel to classic films such as 'Being John Malkovich', 'Airplane!' and 'Monty Python and the Holy Grail', combining elements of comedy, adventure and fantasy. It's sure to be a hit with viewers of all ages!



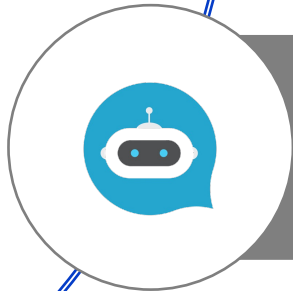
Insights on ICL in RecSys



Teach LLMs to Act as RecSys



Bridge Traditional RecSys and LLMs



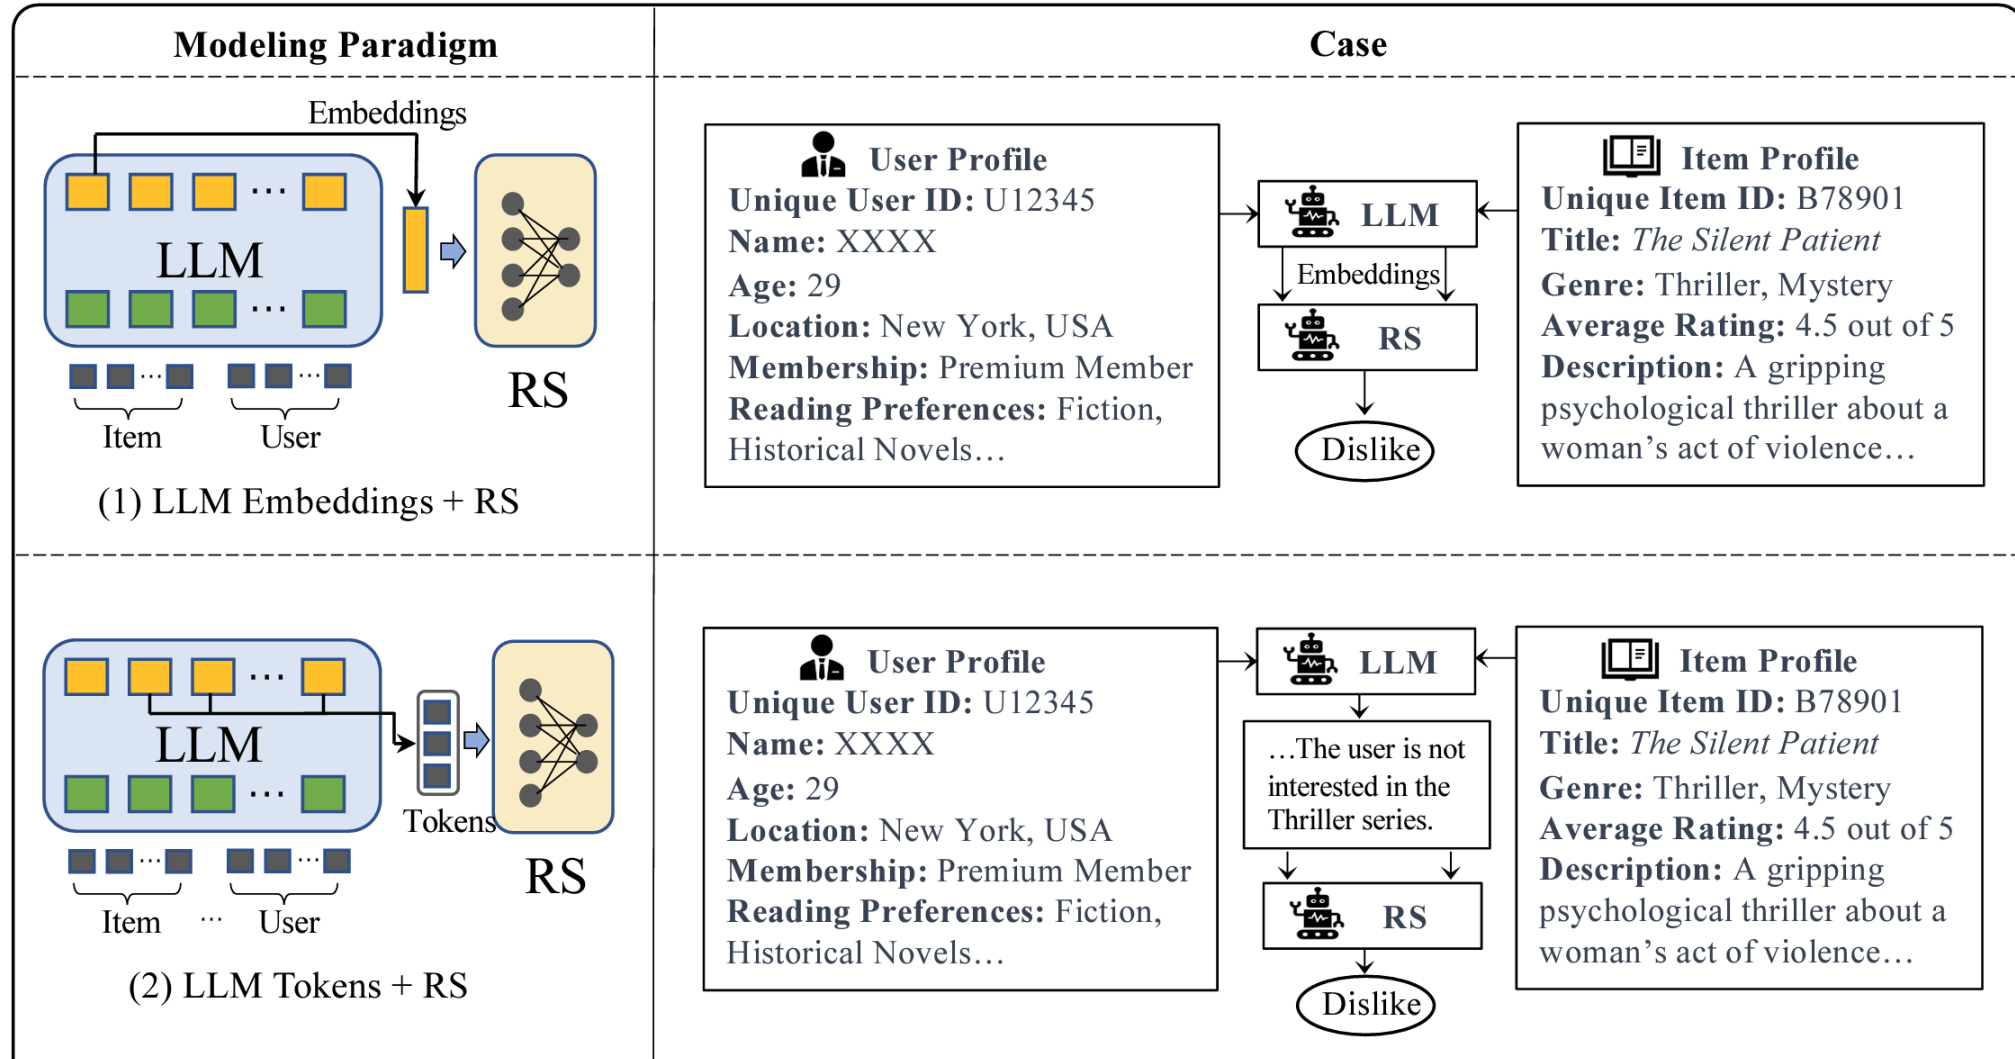
Act as Agent & Use External Tools



Bridge Traditional RecSys and LLMs



- Integrate LLMs as **feature extractor** of users and items into RecSys

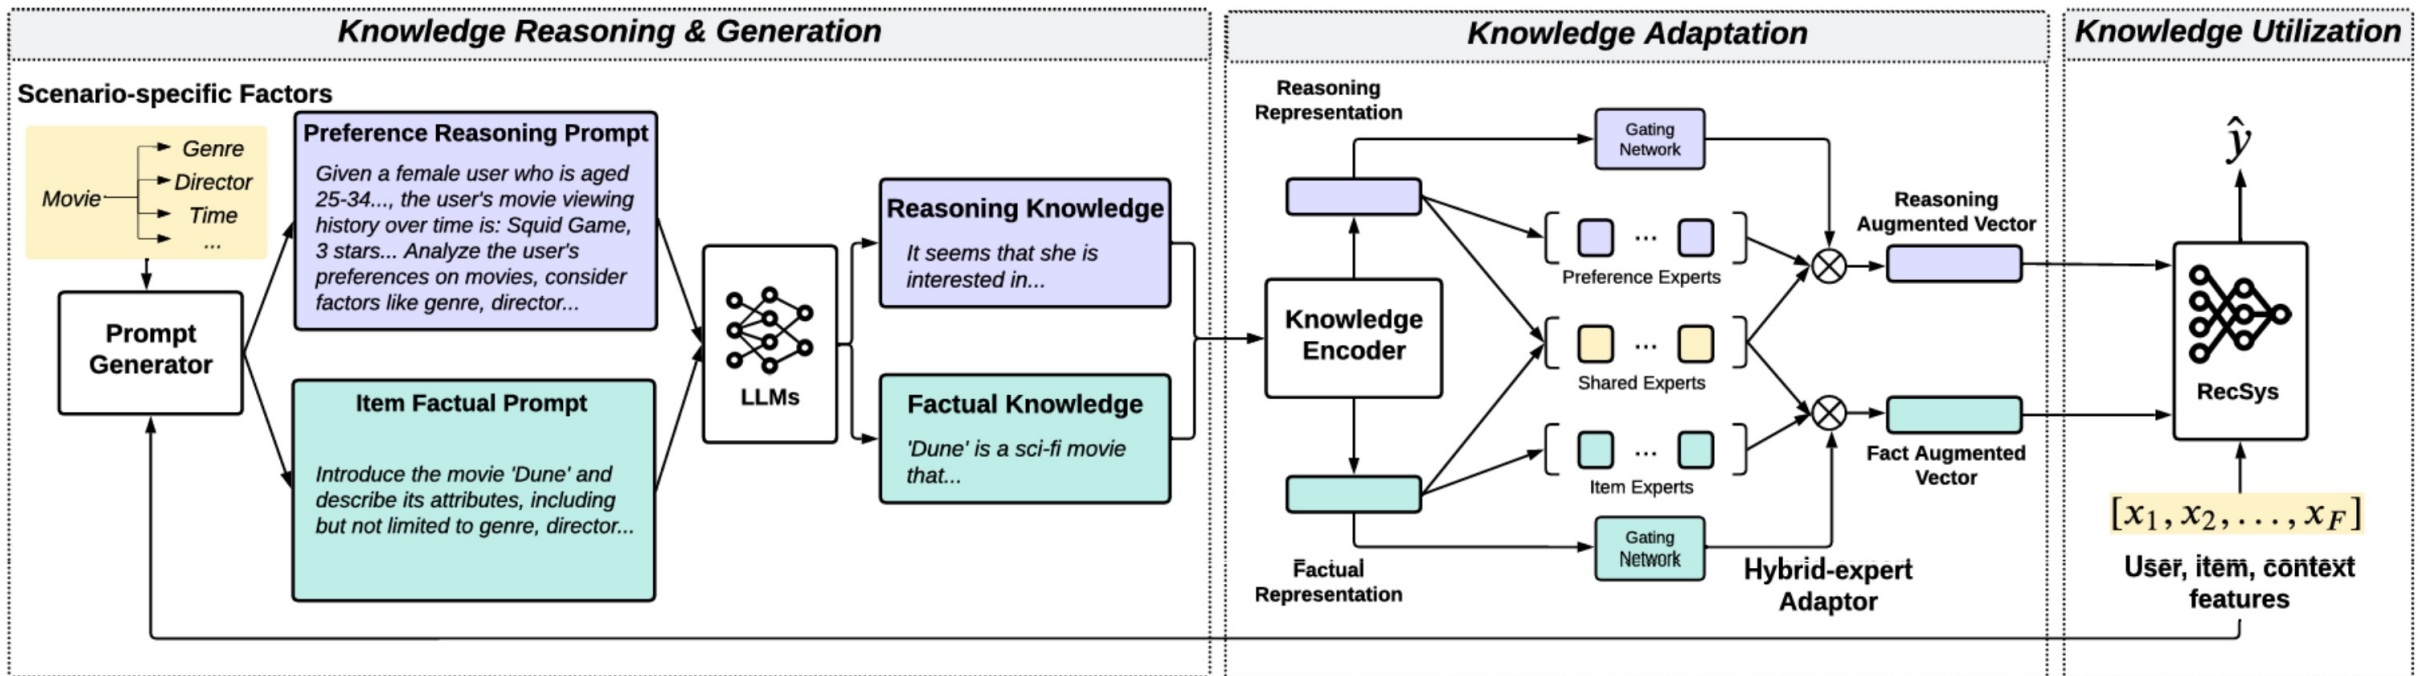


Bridge Traditional RecSys and LLMs



□ KAR

- ❖ Prompt LLMs to **obtain open-world knowledge** beyond original recommendation dataset.
- ❖ Integrate LLM-based open-world knowledge into **domain knowledge of RecSys**.

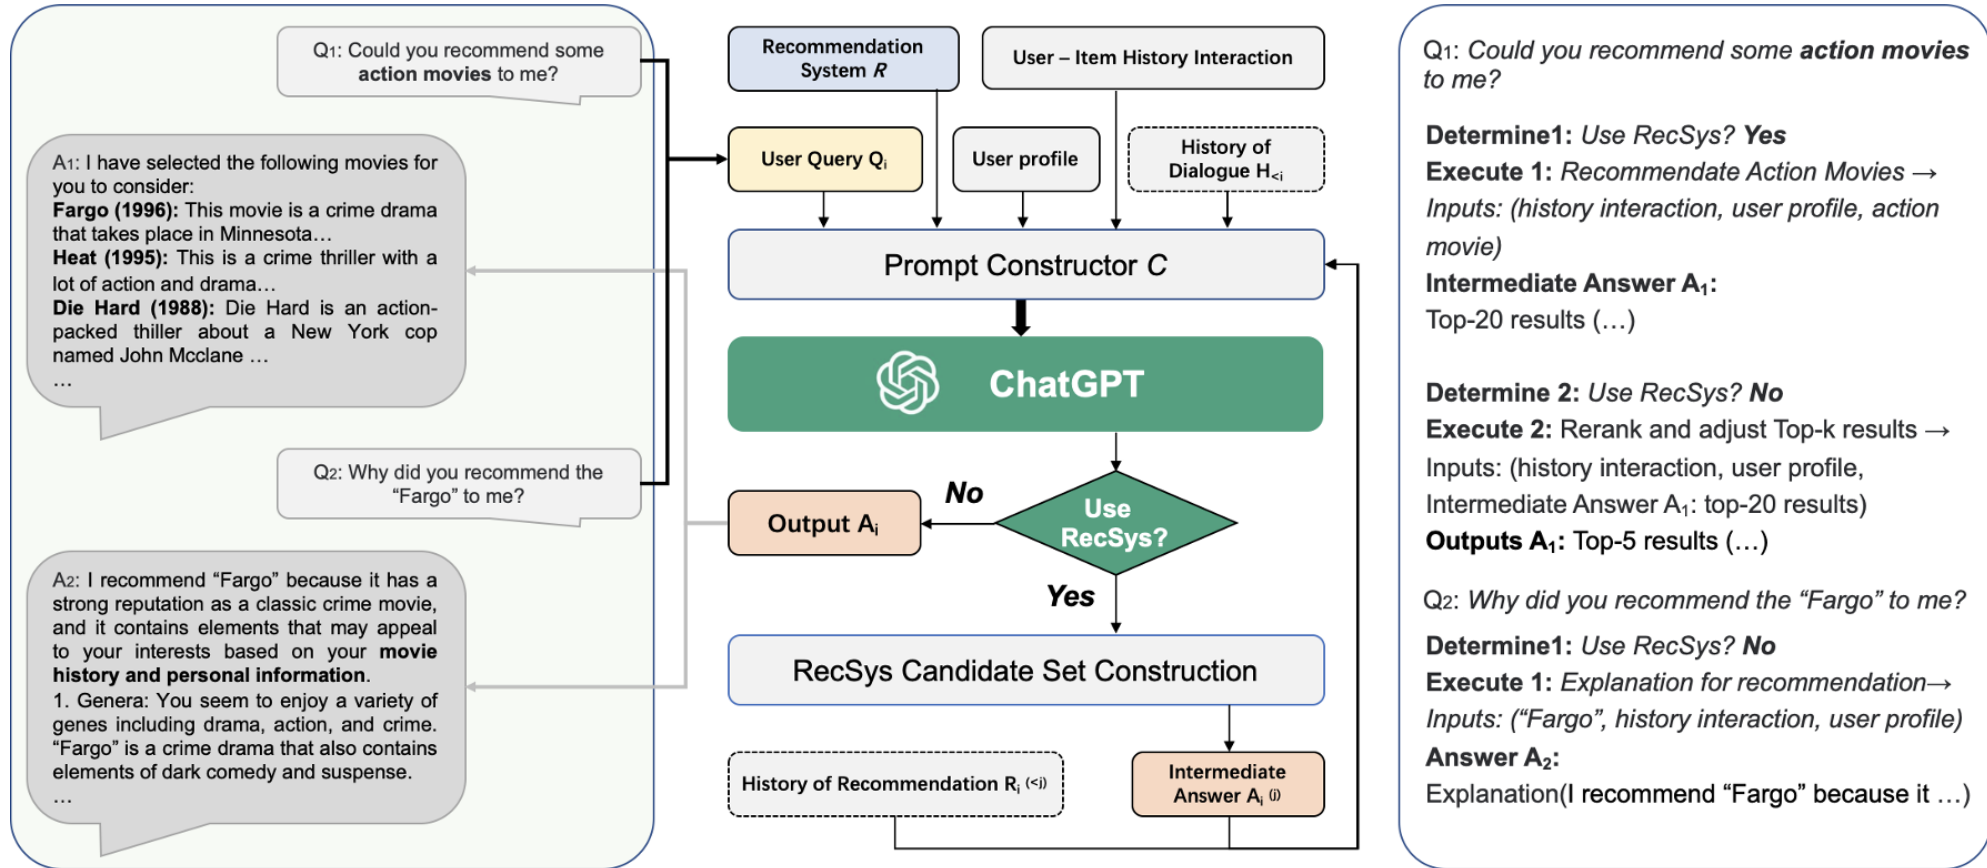


Bridge Traditional RecSys and LLMs

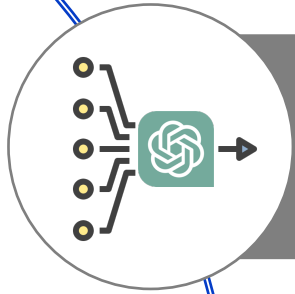


❑ Chat-Rec does it **vice versa**

- ❖ RecSys generate a **large set of candidate** items.
- ❖ LLMs **refine candidate** set based on user dialogue and other **side information**.



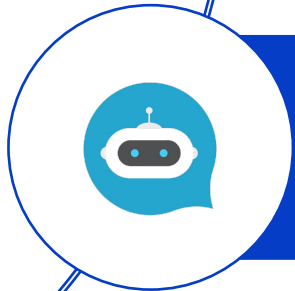
Insights on ICL in RecSys



Teach LLMs to Act as RecSys



Bridge Traditional RecSys and LLMs



Act as Agent & Use External Tools

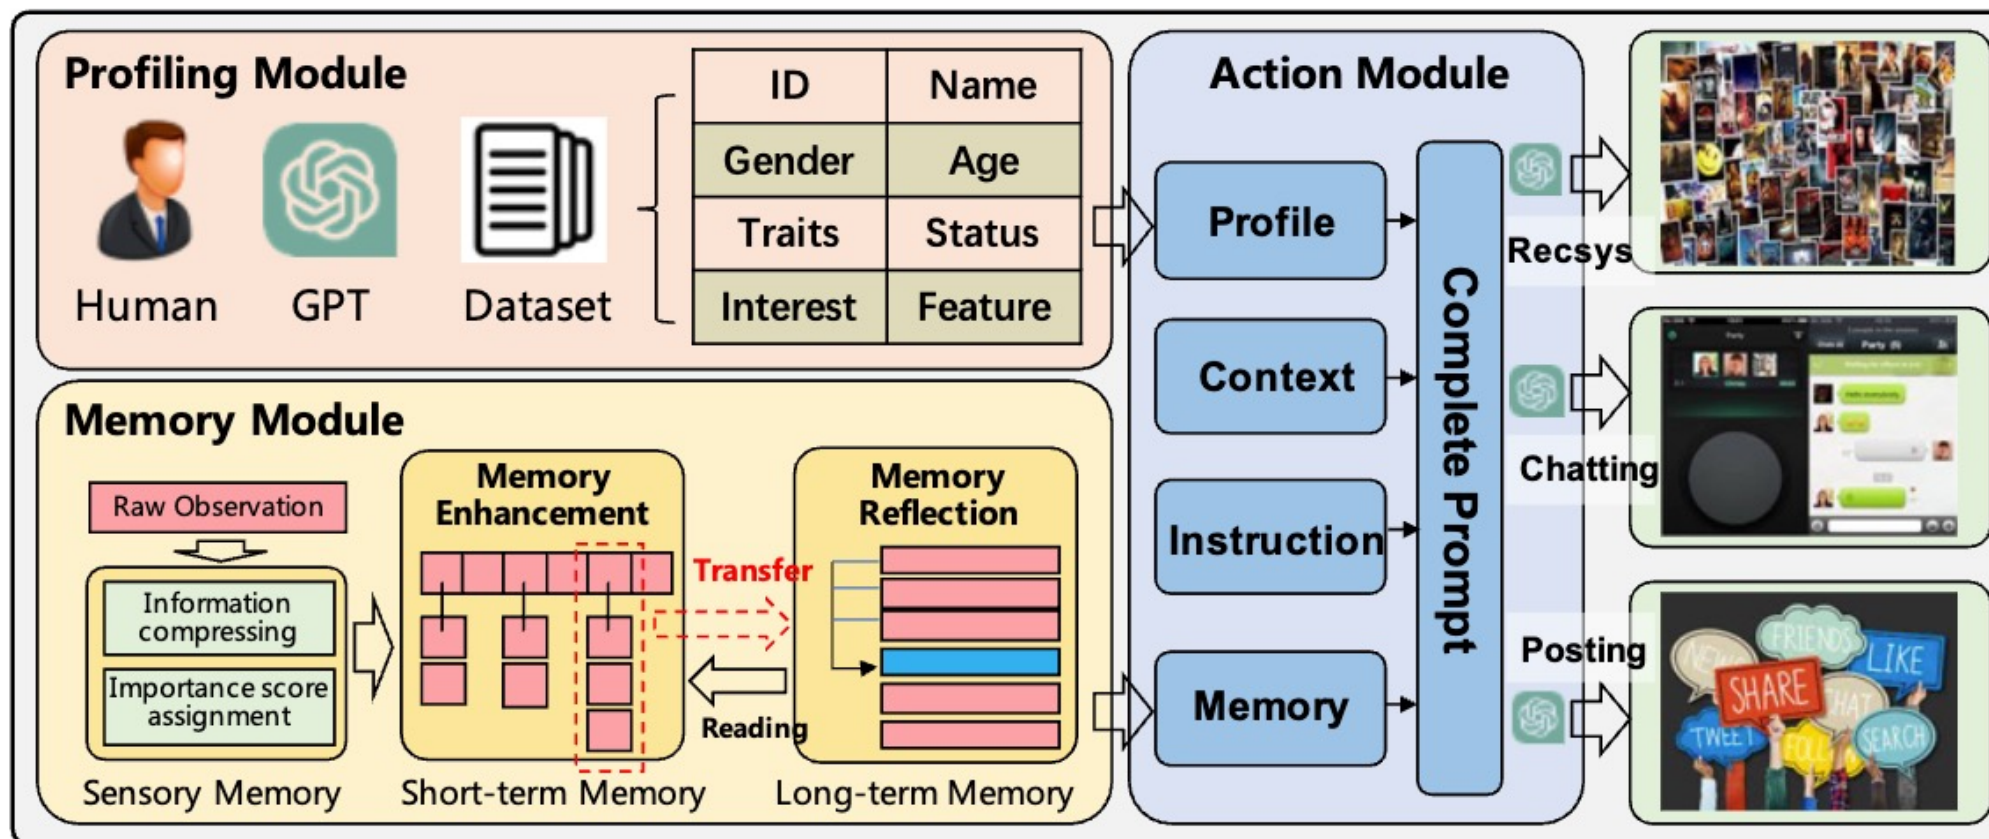


Act as Agent & Use External Tools



❑ RecAgent

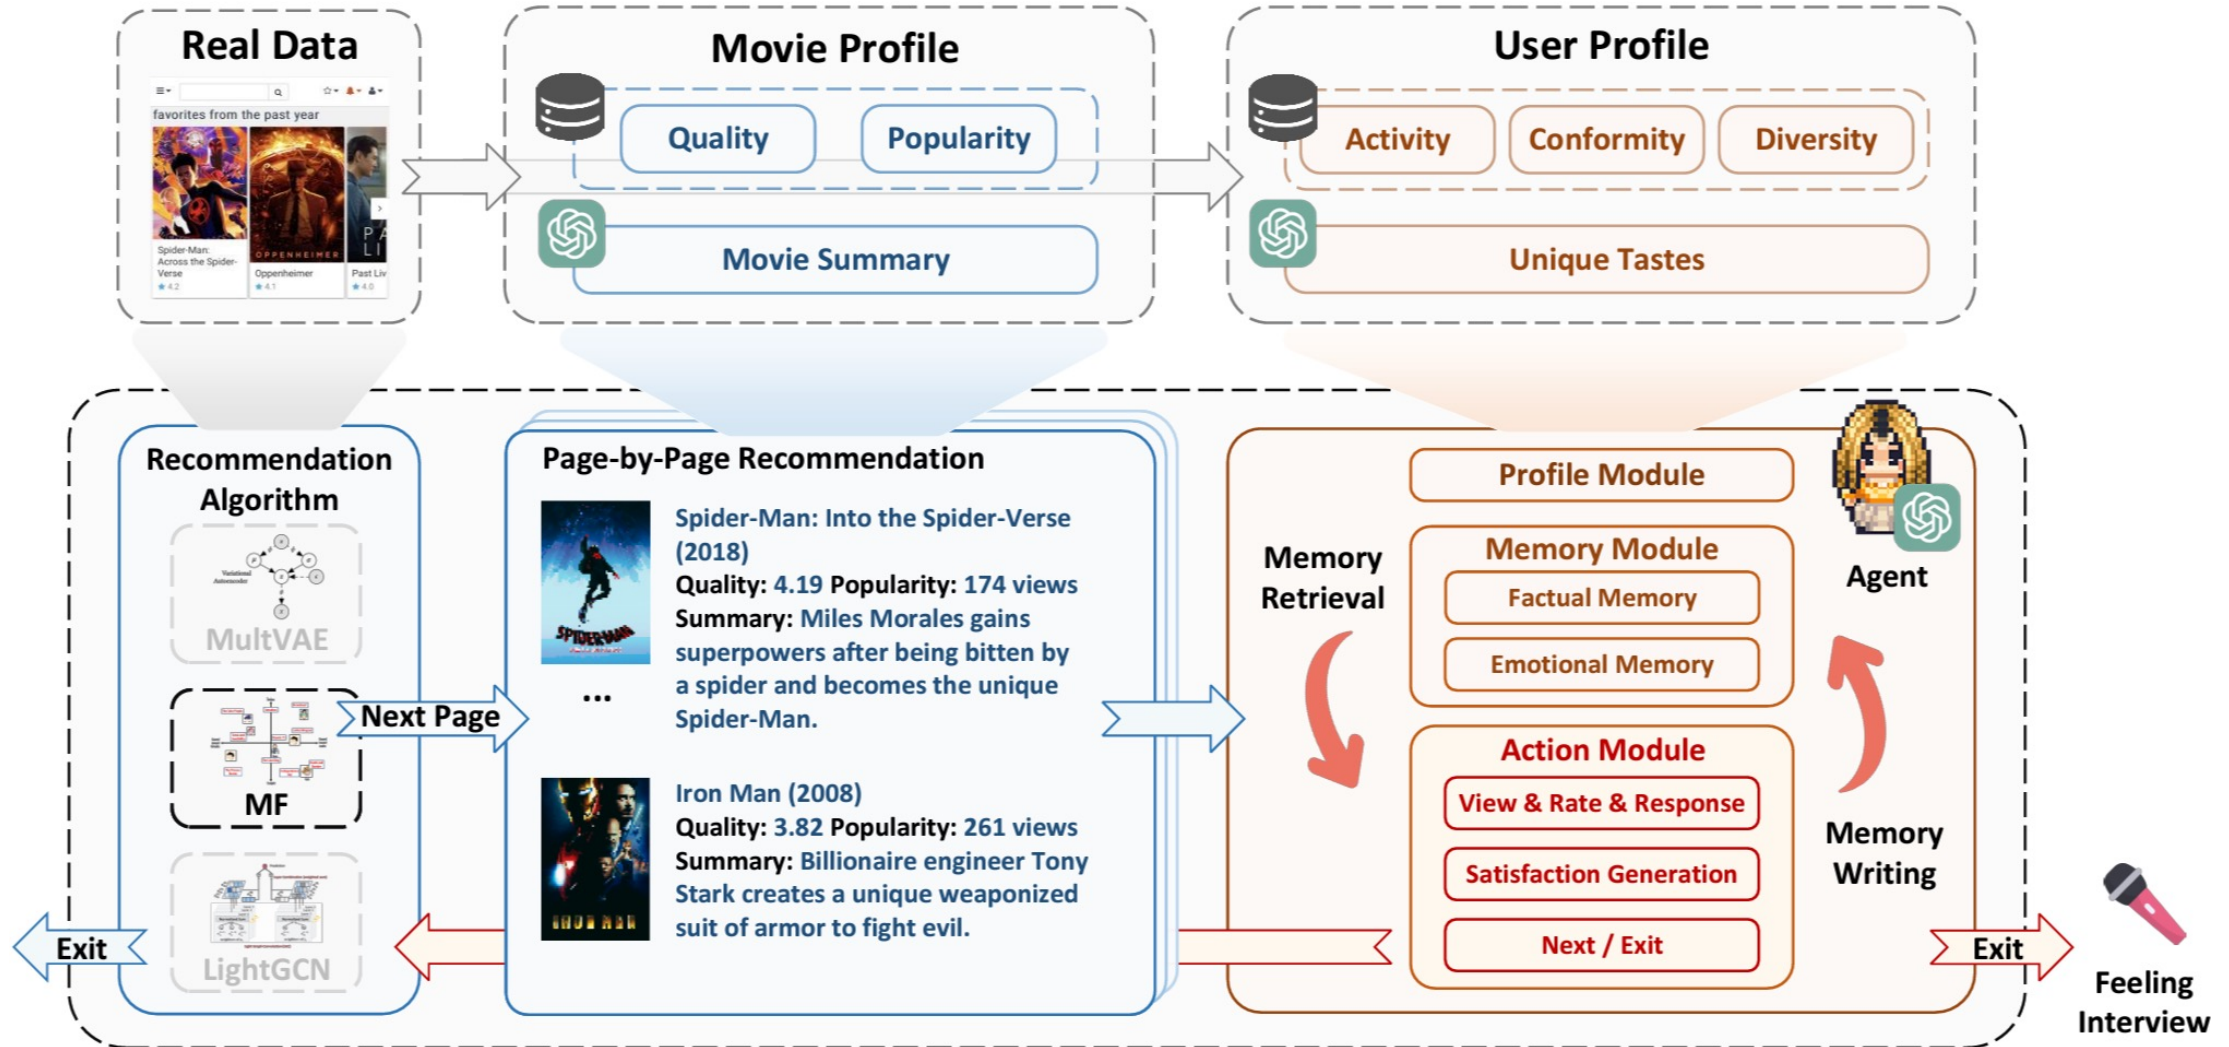
- ❖ LLMs act as **agents** to simulate user behaviors: **RecSys**, chatting, posting.



Act as Agent & Use External Tools



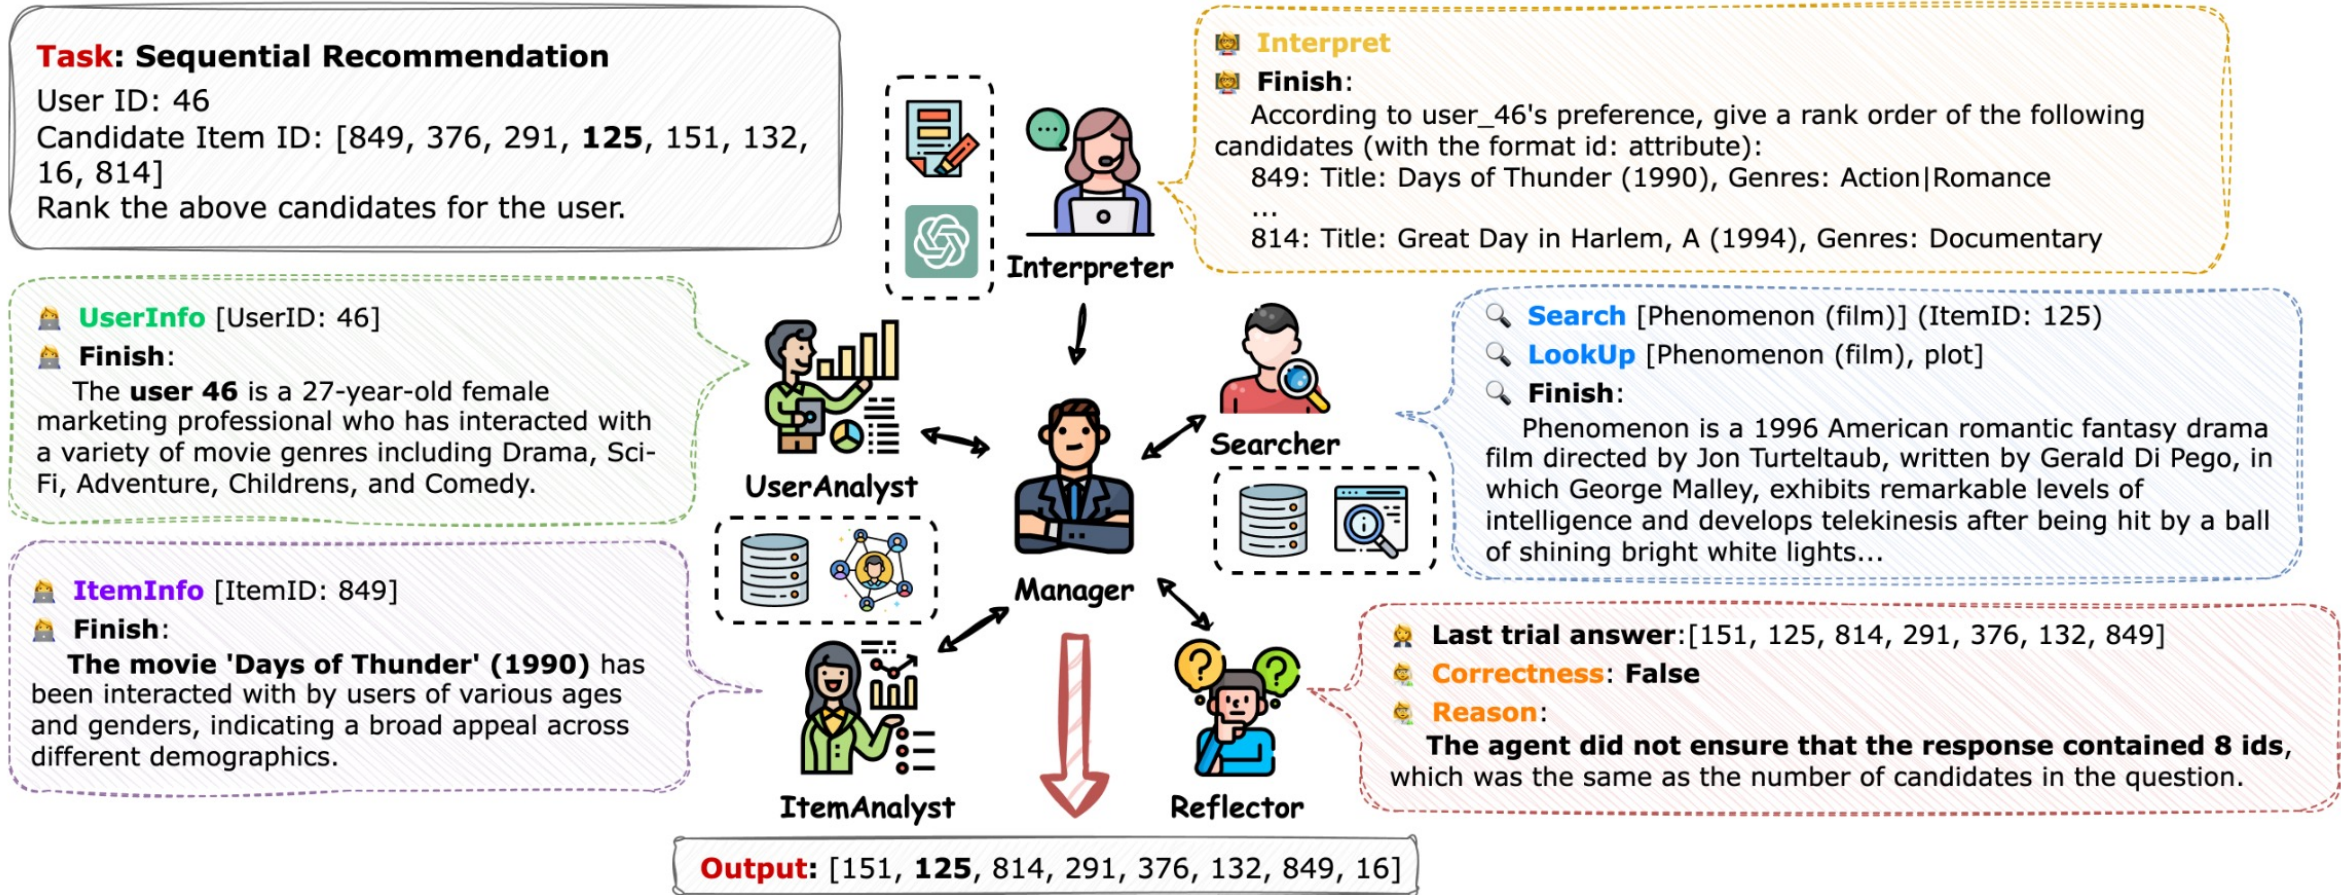
Agent4Rec



Act as Agent & Use External Tools



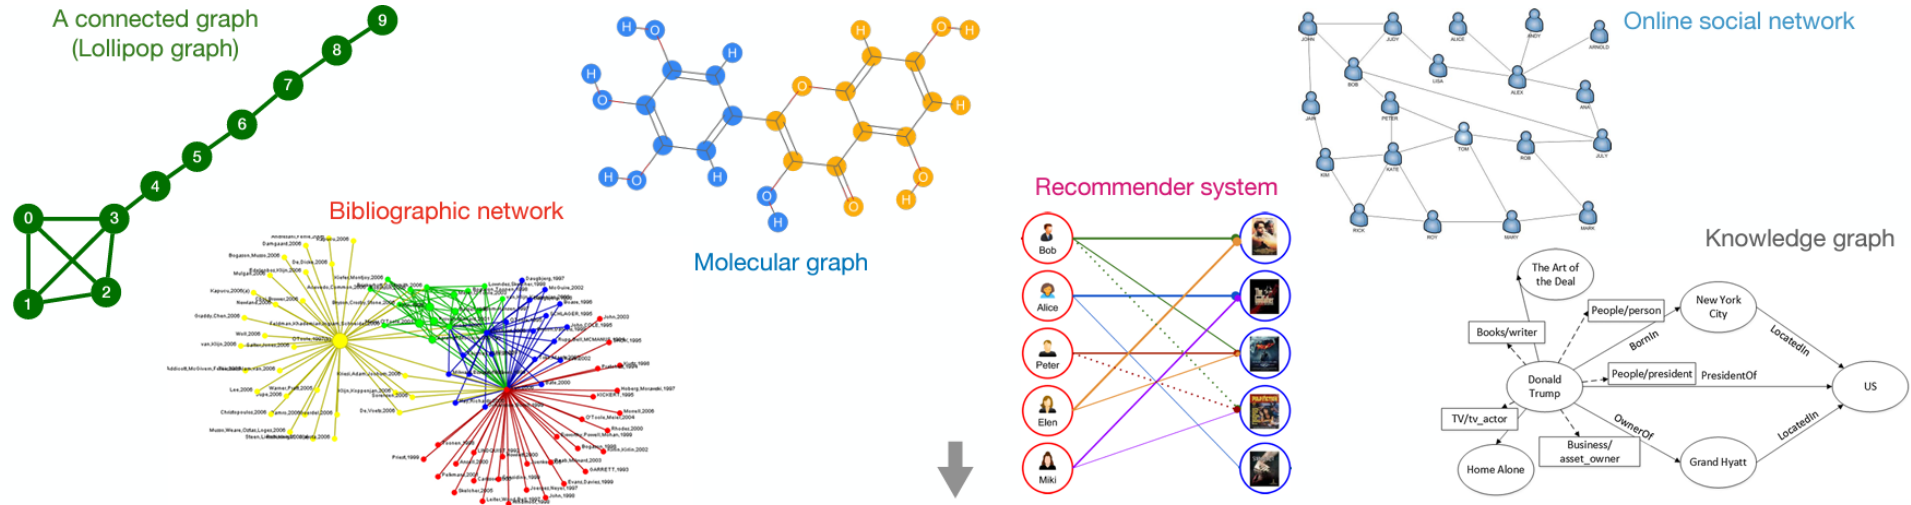
MACRec



Act as Agent & Use External Tools



Graph-ToolFormer



In the lollipop graph, its node number is `<API>GR(GL("lollipop-graph"), "toolx:order") ->10</API>` 10.

Graph property reasoning

The function of insulin protein is to `<API>GR(GL("insulin-graph"), "seg-bert:function") -> "control blood glucose"</API>` control blood glucose.

Protein molecular graph function reasoning

According to Wikipedia, from 2017 to 2021, Donald Trump was the president of `<API>GR(GL("wikipedia-knowledge-graph"), "transe:tail_entity", "Donald Trump", "president of") -> "United States"</API>` United States.

Knowledge graph reasoning



Act as Agent & Use External Tools



□ RecMind

- ❖ Perform **API calls** of specific tools tailored to tasks.
- ❖ **Task planning** to break recommendation tasks into manageable steps.

Rating Prediction

How will **user_X** rate the item "Kusco-Murphy Tart Hair"?
The rating should be an integer between 1 to 5, with 1 being lowest and 5 being highest.

Direct Recommendation

From the item candidates listed below, choose the top 10 items to recommend to **user_X** and rank them in order of priority from highest to lowest.
Candidates: ["Rogaine Women Hair Regrowth Treatment",]

Sequential Recommendation

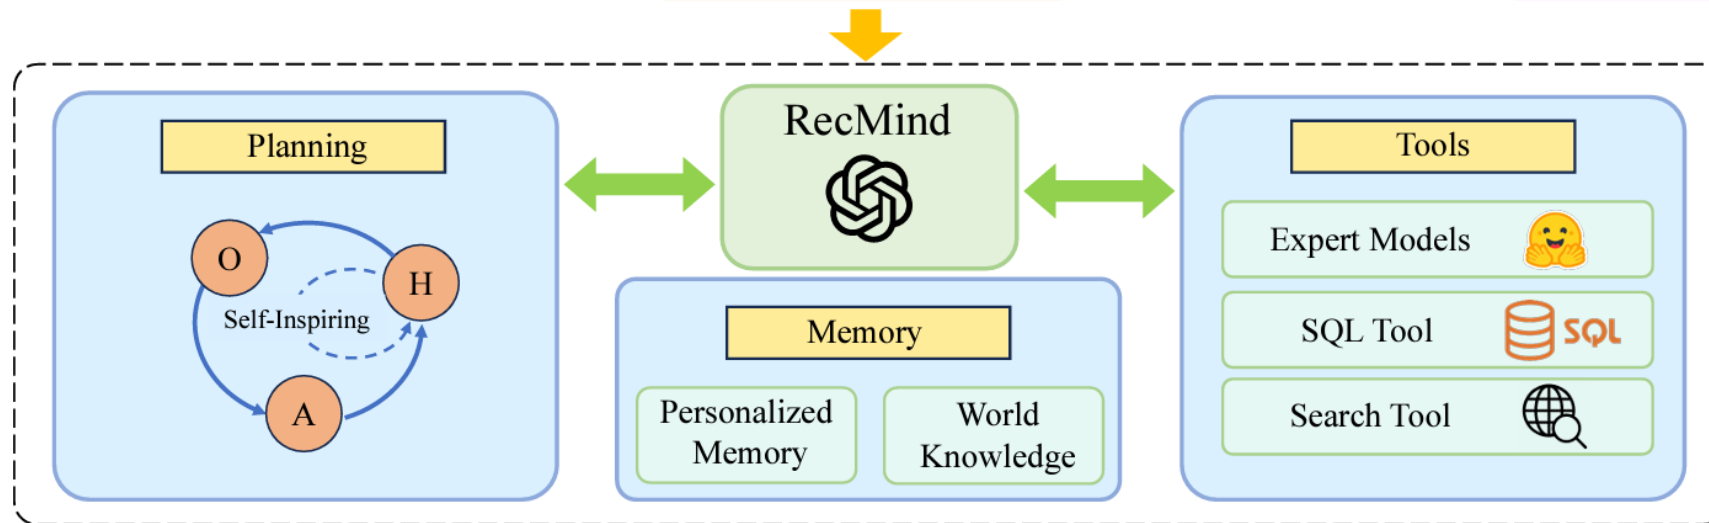
user_X has interacted with the following items in chronological order: ["Old Spice Body Wash Red Zone",]
Please recommend the next item that the user might interact with.
Choose the top 10 products to recommend in order of priority, from highest to lowest.

Review Summarization

Write a review title to summarize the review from **user_X** to item "Chrome Razor and Shaving Brush Stand". The review is "The stand is more solid then I expected for the price. The shape of this stand allows me to hang the shaving brush over the soap bowl, I couldn't do that with stand I had gotten with the kit."

Explanation Generation

Help **user_X** to generate a 5-star explanation for item "FoliGrowth Hair Growth Supplement".



PART 5: RecSys Prompting



Website of this tutorial

- ⊙ **Prompting**
 - ⊙ In-context Learning (ICL)
 - **Chain-of-Thought (CoT)**
- **Prompt Tuning**
 - Hard prompt tuning
 - Soft prompt tuning
- **Instruction Tuning**
 - Full-model tuning with prompt
 - Parameter-efficient model tuning with prompt



Chain-of-Thought (CoT) Prompting



- Annotates intermediate **reasoning steps** into prompt to enhance the reasoning ability of LLMs

Standard Prompting

Model Input

Q: Roger has 5 tennis balls. He buys 2 more cans of tennis balls. Each can has 3 tennis balls. How many tennis balls does he have now?

A: The answer is 11.

Q: The cafeteria had 23 apples. If they used 20 to make lunch and bought 6 more, how many apples do they have?

Model Output

A: The answer is 27. ❌

Chain-of-Thought Prompting

Model Input

Q: Roger has 5 tennis balls. He buys 2 more cans of tennis balls. Each can has 3 tennis balls. How many tennis balls does he have now?

A: Roger started with 5 balls. 2 cans of 3 tennis balls each is 6 tennis balls. $5 + 6 = 11$. The answer is 11.

Q: The cafeteria had 23 apples. If they used 20 to make lunch and bought 6 more, how many apples do they have?

Model Output

A: The cafeteria had 23 apples originally. They used 20 to make lunch. So they had $23 - 20 = 3$. They bought 6 more apples, so they have $3 + 6 = 9$. The answer is 9. ✅

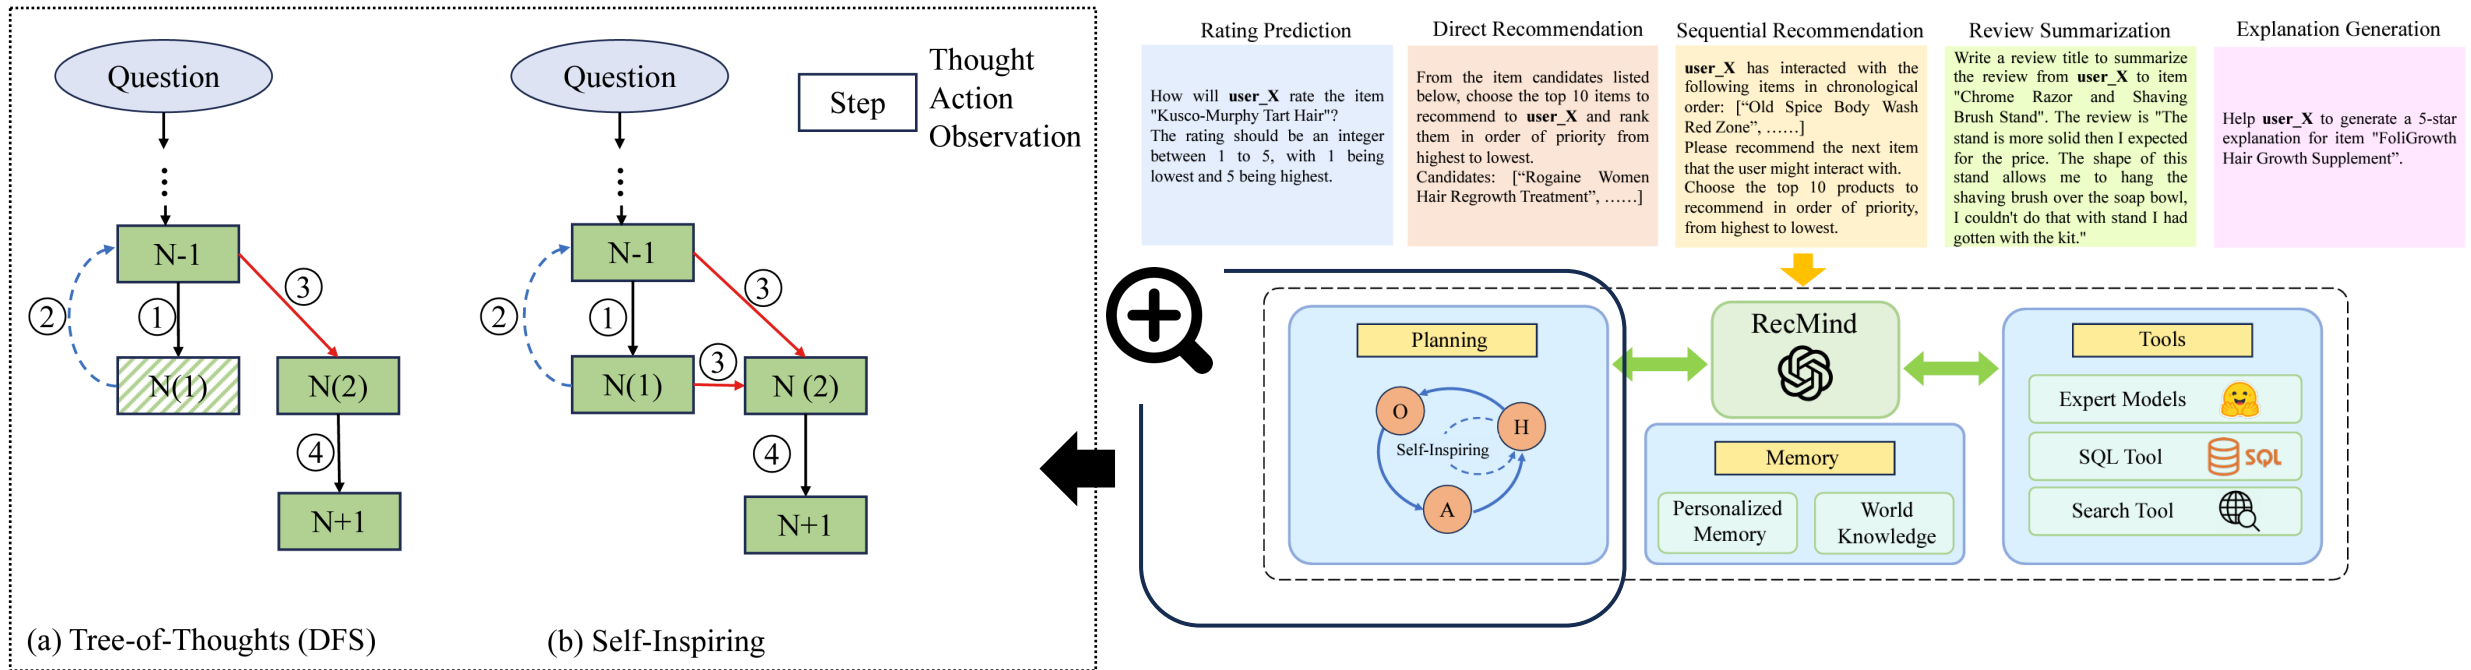


Beyond “Chain”-of-Thought



❑ RecMind

- ❖ **Tree-of-Thoughts (ToT, 2023)**: generate & select multiple candidates for next step, but eventually return single reasoning path similar to CoT.
- ❖ **Self-Inspiring (SI, proposed)**: further explore alternative reasoning path in parallel to other paths.

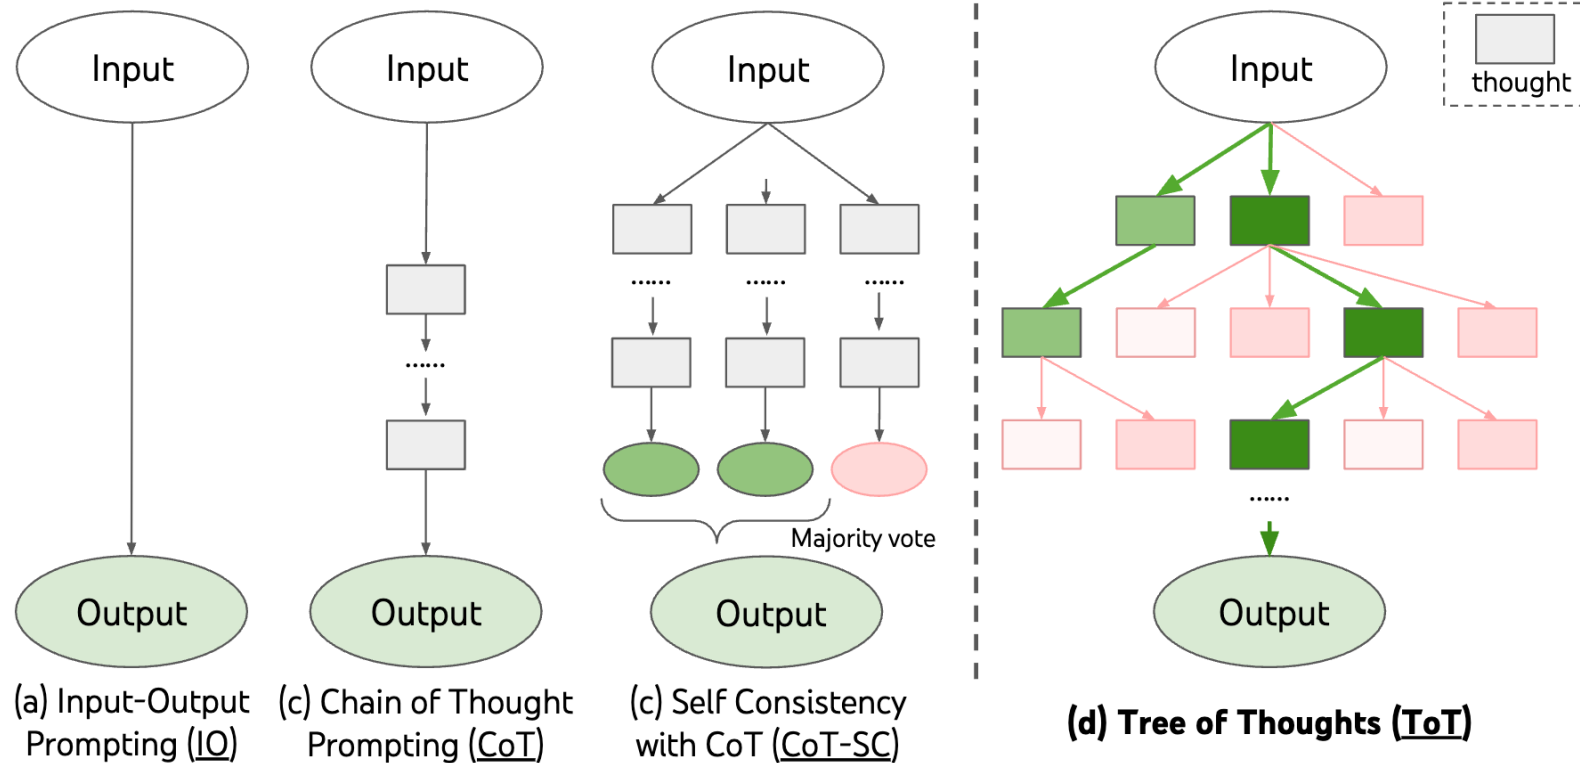


Potential of Tree-of-Thought



□ ToT

- ❖ ToT actively maintains a tree of thoughts.
- ❖ LLM-enpowerd RecSys may also benefit from TOT.

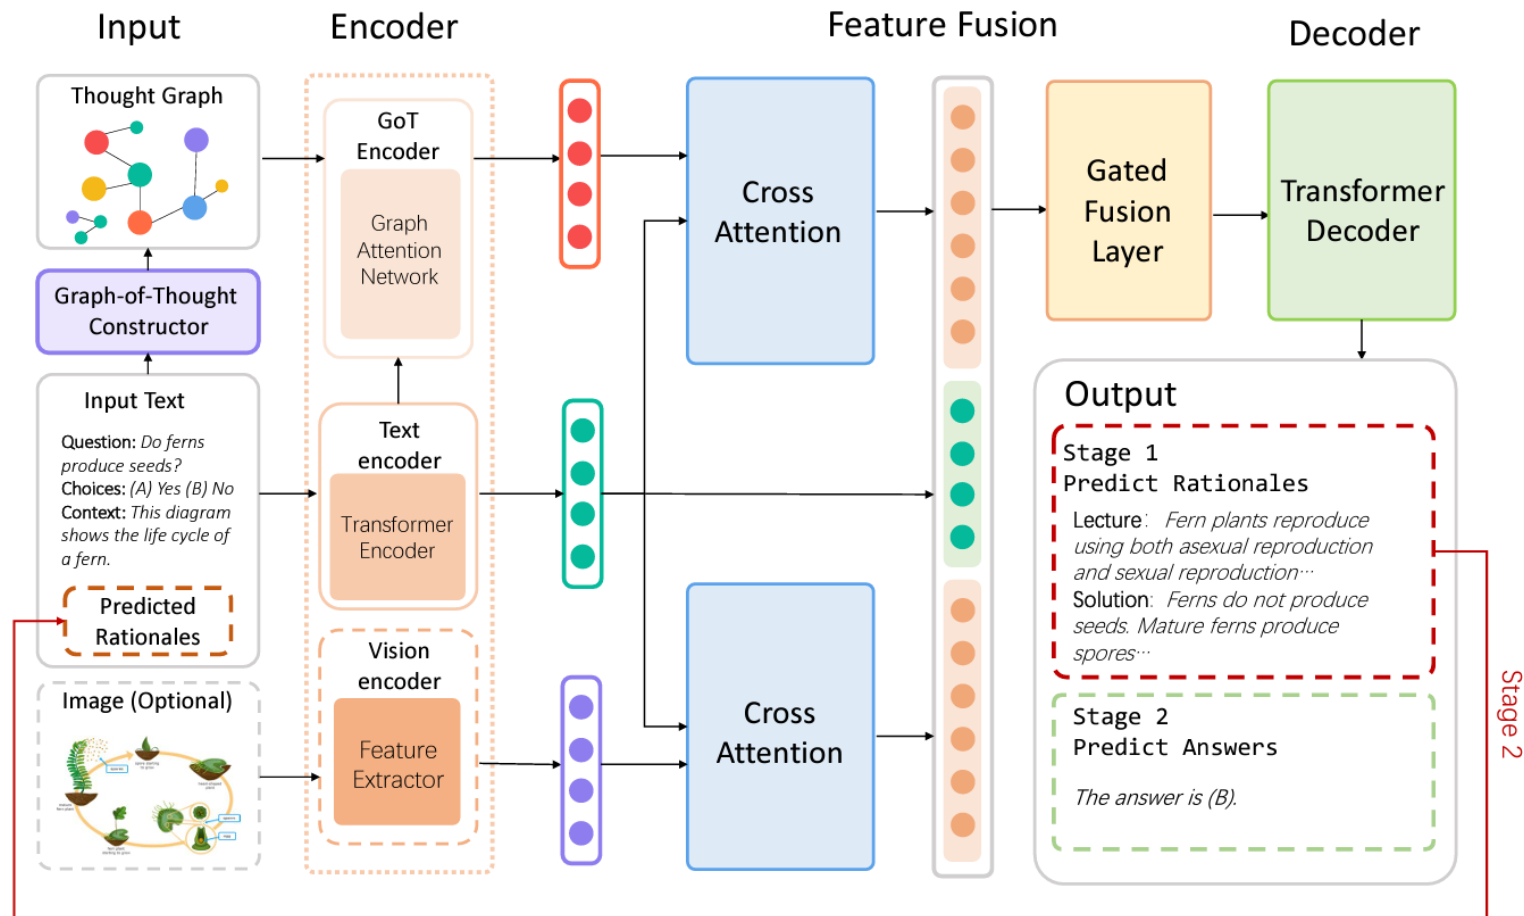


Potential of Graph-of-Thought



GoT

- ❖ Fusion of **thought graph representation** into text representation.
- ❖ **RecSys** can be considered as a special case of **link prediction** problems in graph learning.



PART 5: RecSys Prompting



Website of this tutorial

- ⦿ **Prompting**
 - ⦿ In-context Learning (ICL)
 - ⦿ Chain-of-Thought (CoT)
- ⦿ **Prompt Tuning**
 - Hard prompt tuning
 - Soft prompt tuning
- **Instruction Tuning**
 - Full-model tuning with prompt
 - Parameter-efficient model tuning with prompt



Prompting



□ Prompt shapes



GPT

Prefix Prompt

For tasks regarding generation, or tasks being solved using a standard auto-regressive LM, **prefix prompts** tend to be more conducive, as they mesh well with the **left-to-right** nature of the model.



BERT

Cloze Prompt

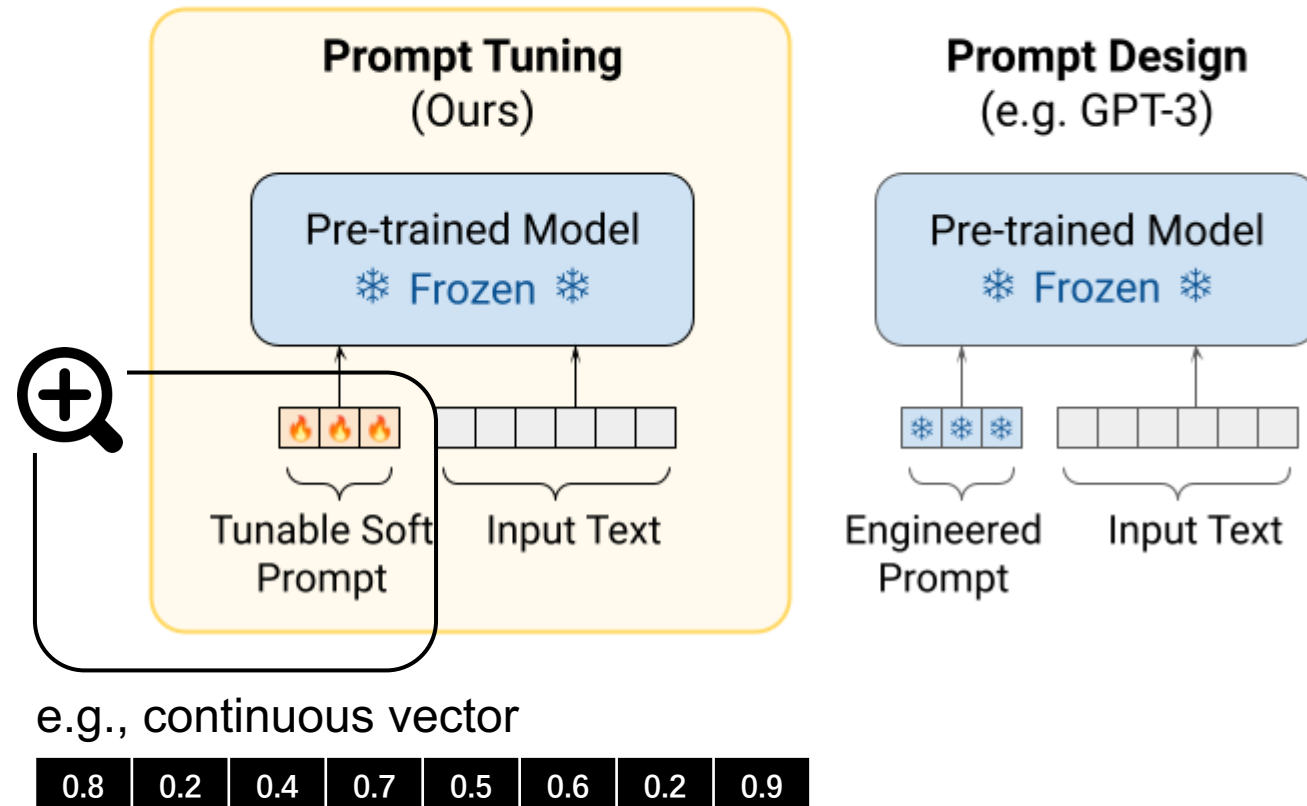
For tasks that are solved using masked LMs, **cloze prompts** are a good fit, as they very closely match the form of the **pre-training** task.



Prompt Tuning



- Only involves minimal parameter updates of the **tunable prompt** and the **input layer** of LLMs
- ❖ Prompt tuning adds **new prompt tokens** to LLM and **optimizes the prompt**.

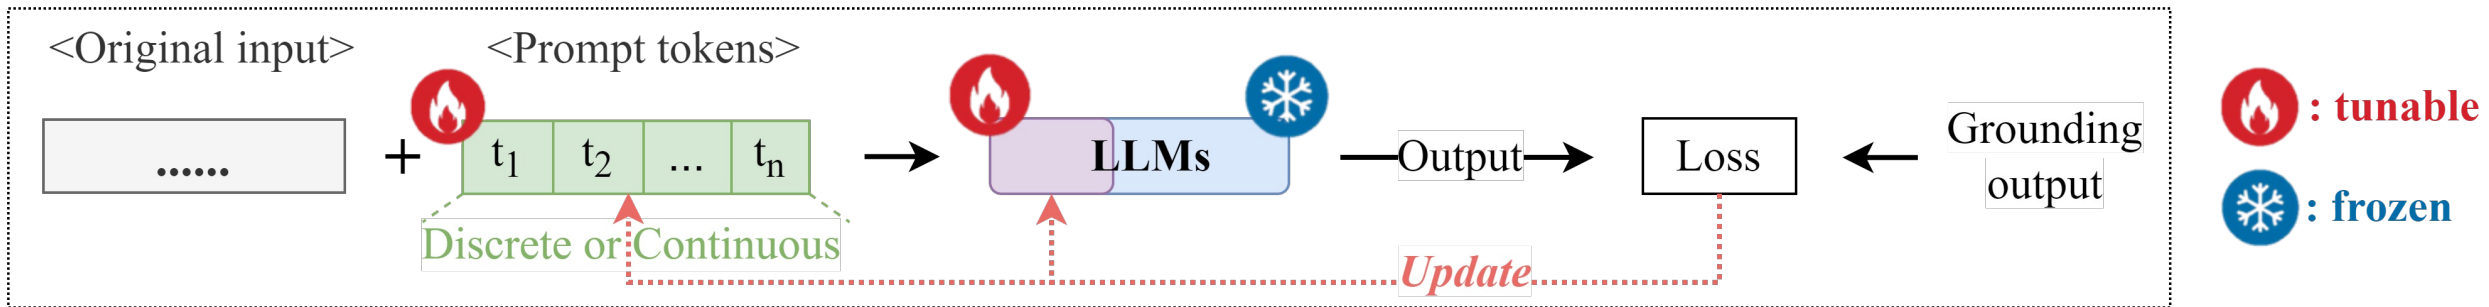


Hard vs. Soft Prompt Tuning



□ Taxonomy

- ❖ “Prompts can be **discrete templates** or **soft parameters** that encourage the model to predict the desired output.”
- ❖ “ICL can be regarded as a **subclass of prompt tuning** where the demonstration is part of the prompt.”



Hard vs. Soft Prompt Tuning



- ❑ Hard prompt tuning - learn tokens of **discrete text templates**
 - ❖ Convenient and effective to refine **natural language prompts** but faces **discrete optimization** challenges, like laborious trial and error to find suitable prompts.
- ❑ Soft prompt tuning - learn tokens of **continuous parameters**
 - ❖ Feasible for tuning on **continuous space** but in a cost of **explainability**, since soft prompts written in continuous vectors are not interpretable to humans.



Which to choose?
Hard or Soft?

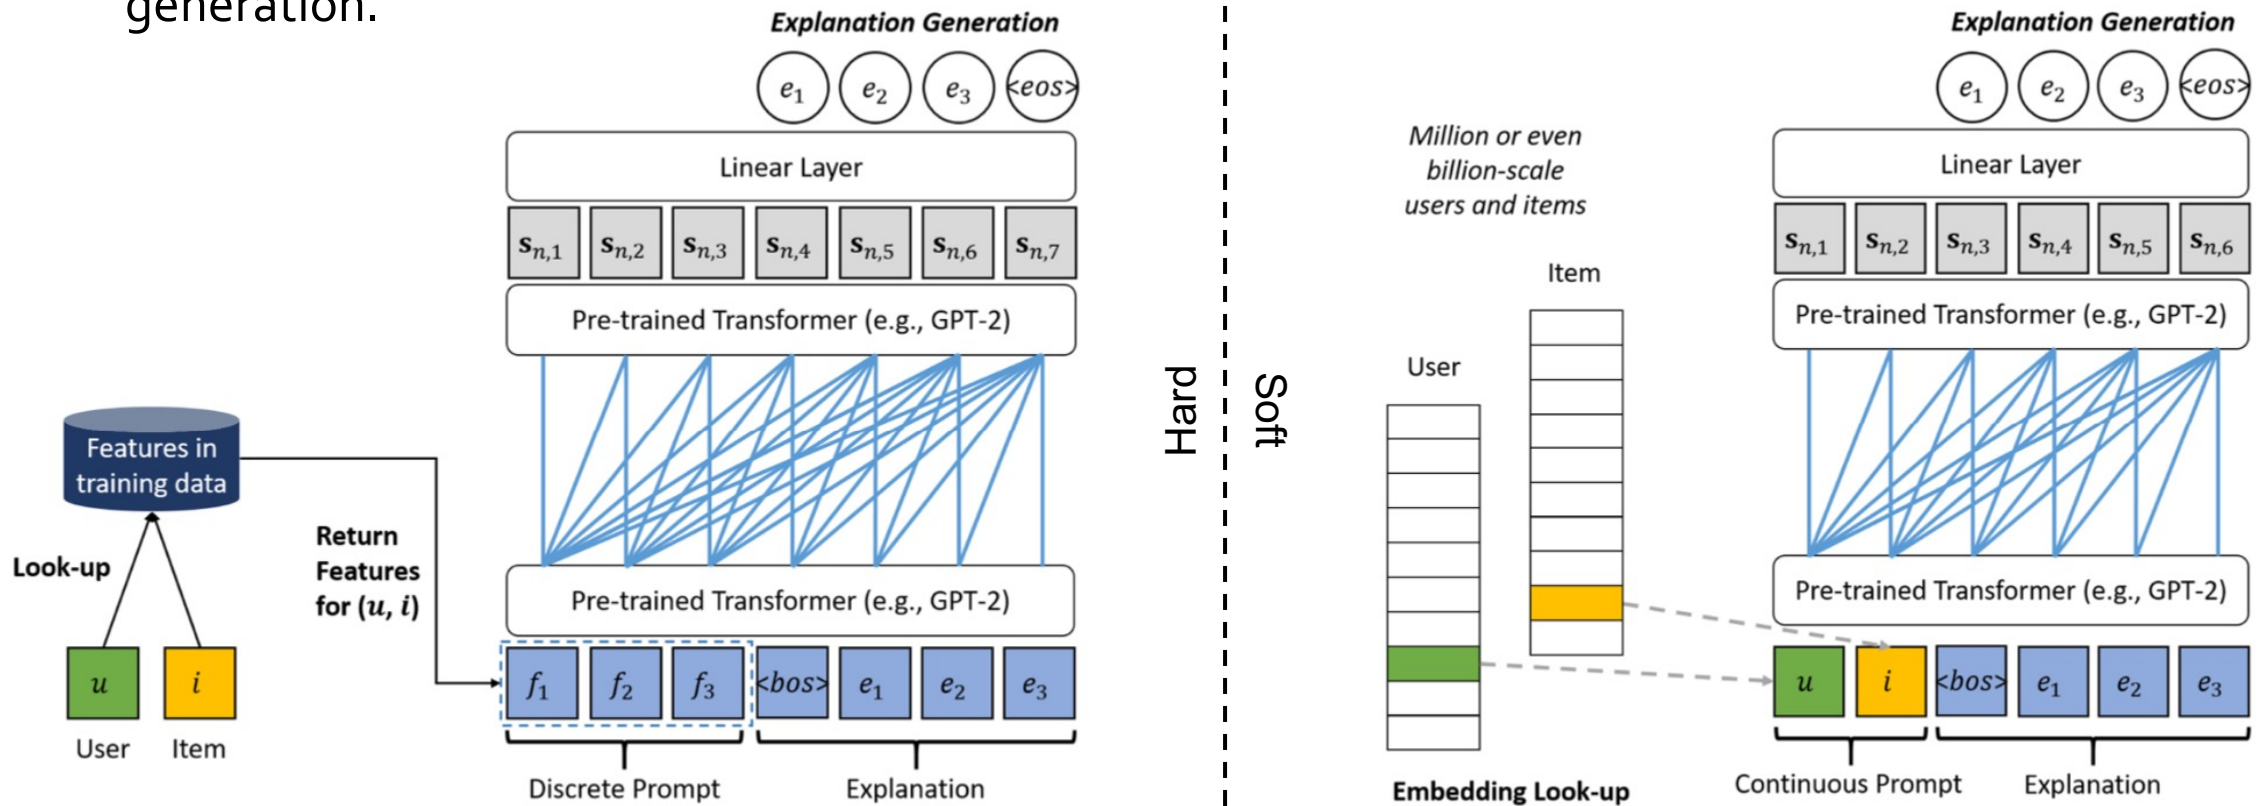


Prompt Tuning in RecSys



□ PEPLER

- ❖ **Hard prompt tuning:** utilizes item features (e.g., titles) as a **discrete prompt** for explanation generation.
- ❖ **Soft prompt tuning:** treats user and item embeddings as **continuous prompt** for explanation generation.

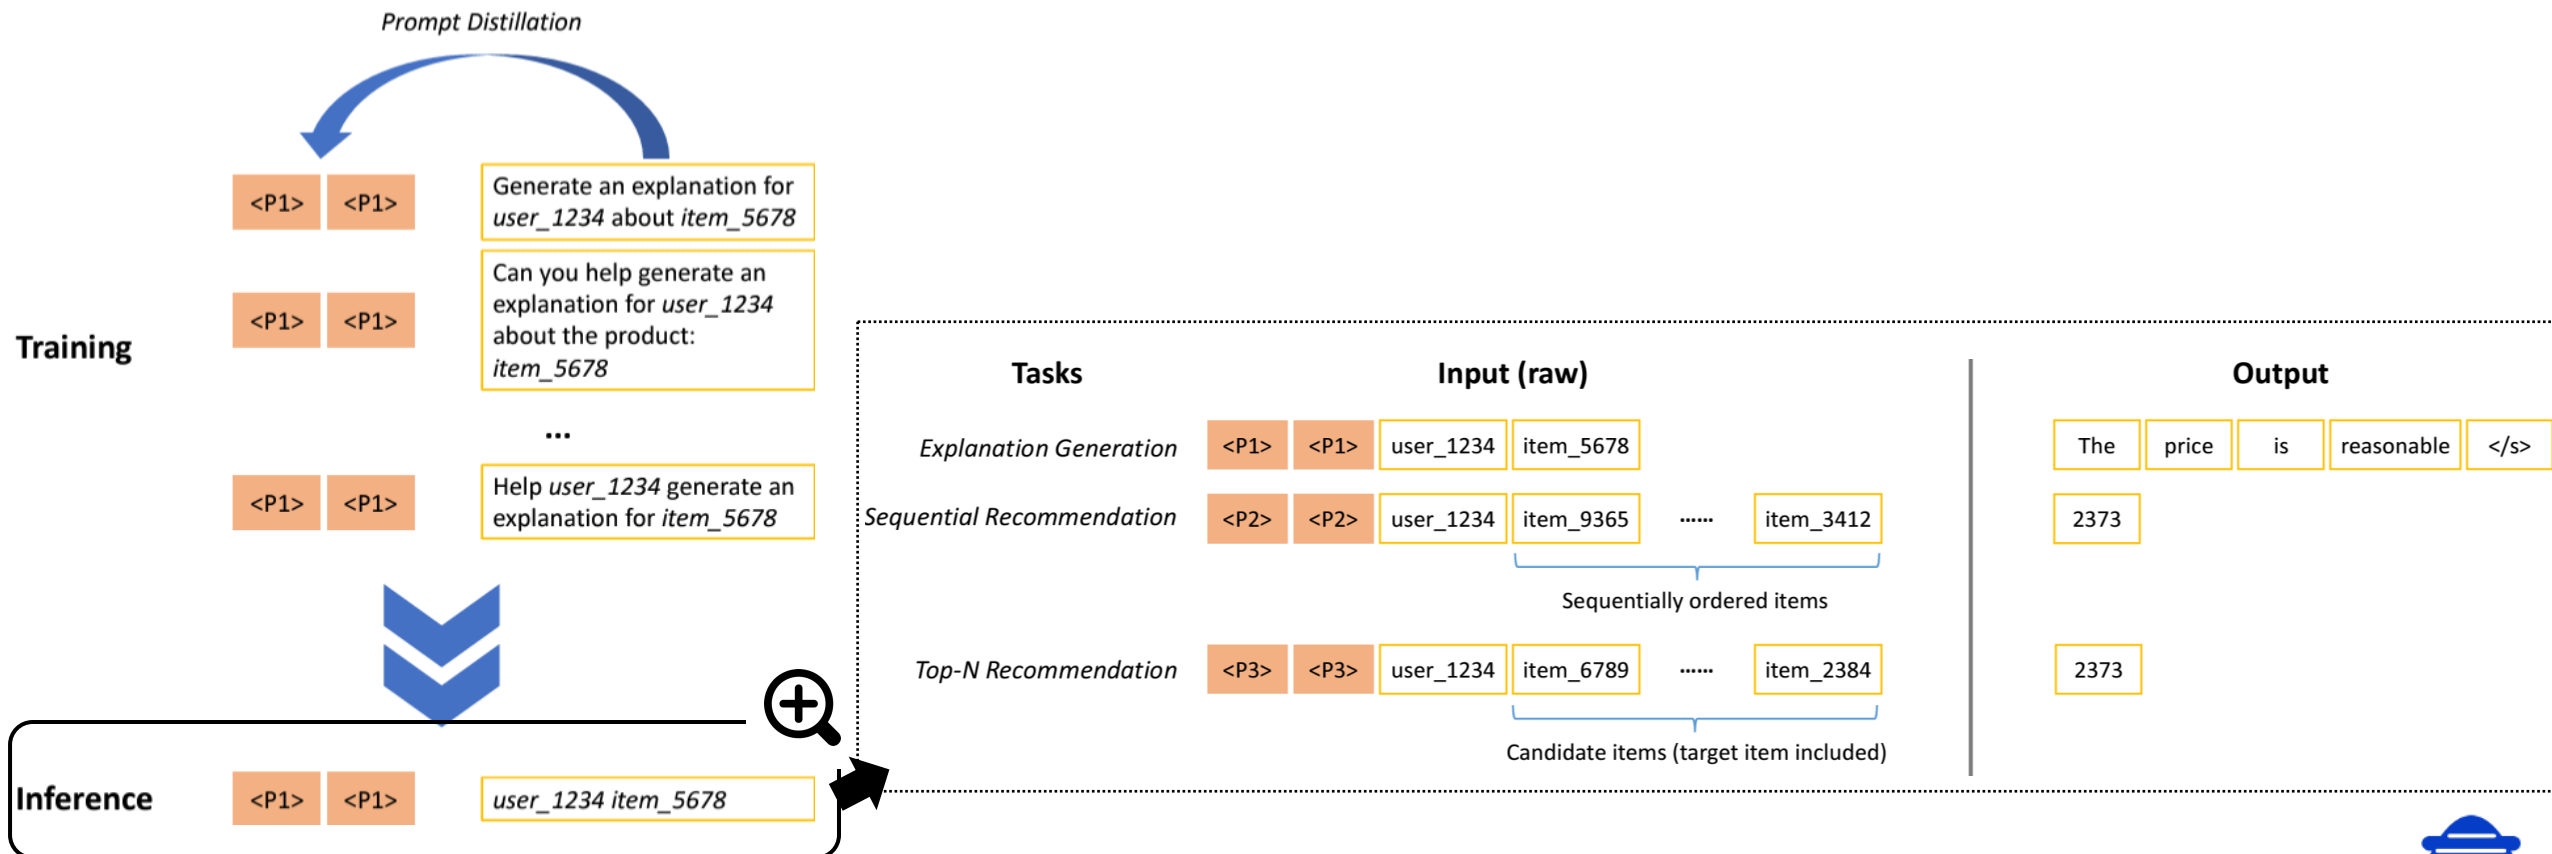


Bridge Hard & Soft Prompt Tuning



POD

- ❖ Discrete **hard prompt** suffers from processing **long text** of user and item IDs.
- ❖ Distill the discrete prompt to a set of **soft prompt** so as to **bridge IDs and texts**.



PART 5: RecSys Prompting



Website of this tutorial

- ⦿ **Prompting**
 - ⦿ In-context Learning (ICL)
 - ⦿ Chain-of-Thought (CoT)
- ⦿ **Prompt Tuning**
 - ⦿ Hard prompt tuning
 - ⦿ Soft prompt tuning
- ⦿ **Instruction Tuning**
 - Full-model tuning with prompt
 - Parameter-efficient model tuning with prompt

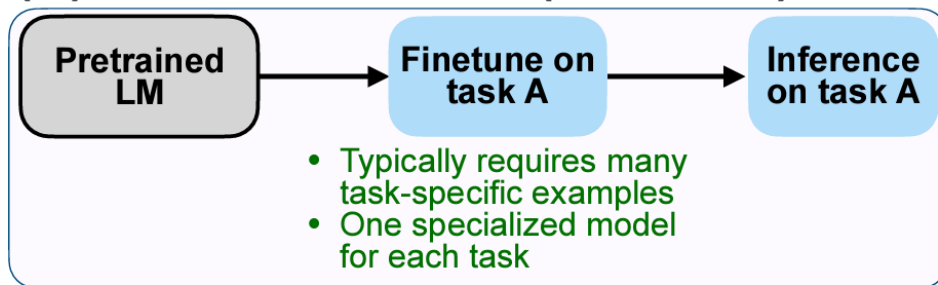


Instruction Tuning

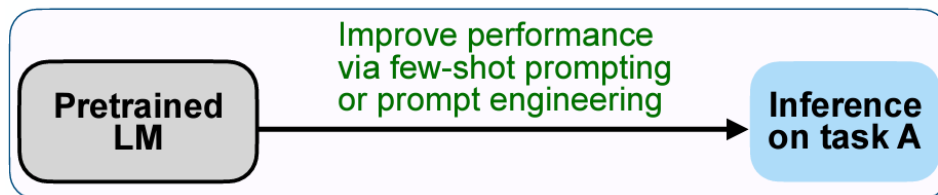


- ❑ To enhanced the **zero-shot performance** of LLMs on unseen tasks by accurately following new **task instructions**
 - ❖ Instruction tuning is a **combination** of both **prompting** and **fine-tuning** paradigms.

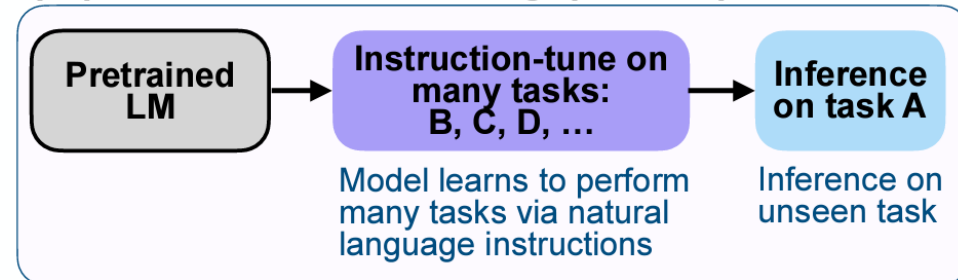
(A) Pretrain–finetune (BERT, T5)



(B) Prompting (GPT-3)



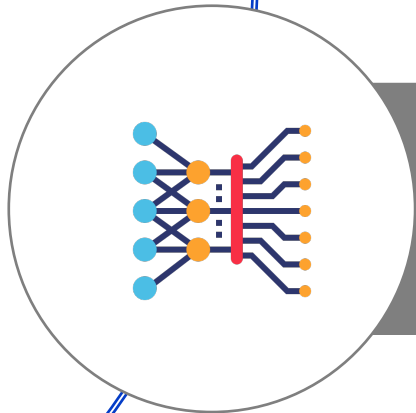
(C) Instruction tuning (FLAN)



Stages of Instruction Tuning



Instruction Generation (\approx Prompting)



Model Tuning with Prompt (\approx Fine-tuning)

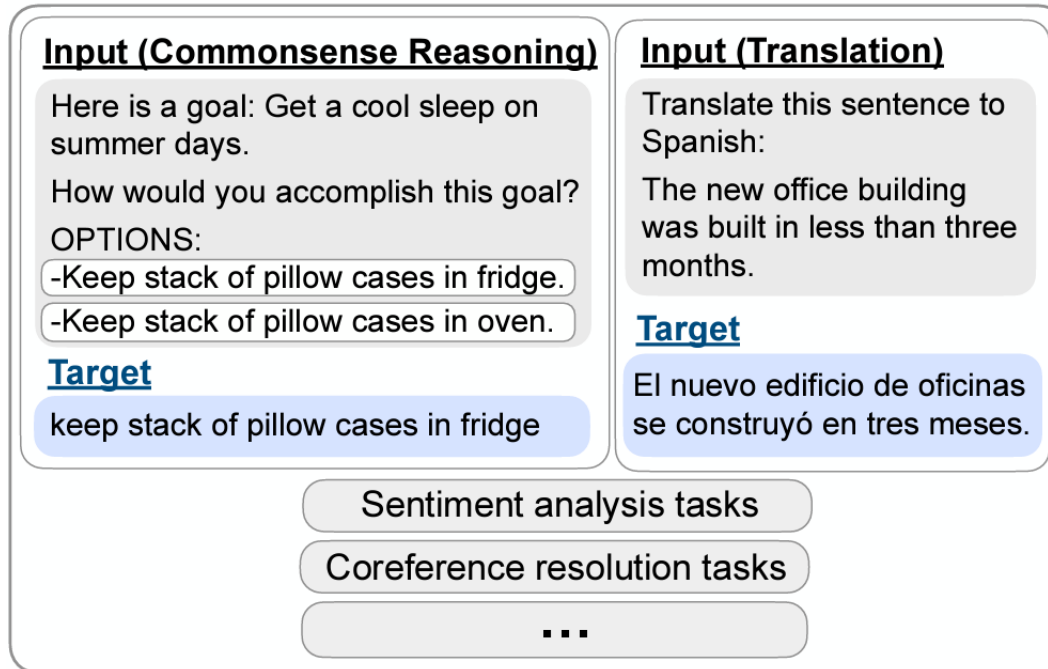


Stage 1: Instruction Generation

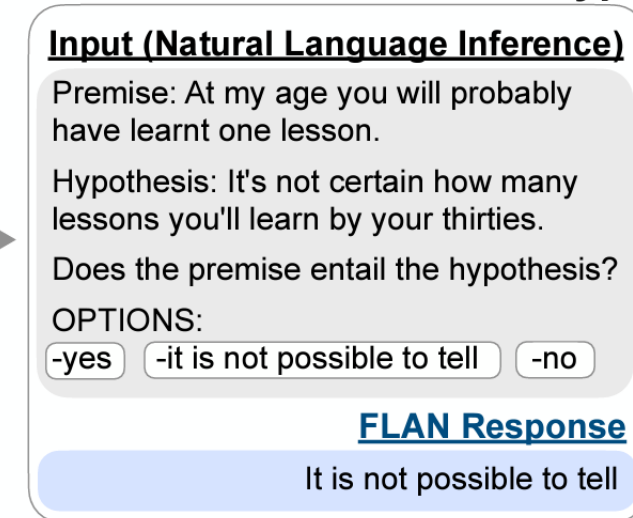


- A format of instruction-based **prompt in natural language**
 - ❖ **Task-oriented input:** task descriptions based on task-specific dataset.
 - ❖ **Desired target:** corresponding output based on task-specific dataset.

Finetune on many tasks (“instruction-tuning”)



Inference on unseen task type



Instruction Generation for RecSys



□ InstructRec

- ❖ Pointwise recommendation (T_0)
- ❖ Pairwise recommendation (T_1)
- ❖ Matching (T_2)
- ❖ Re-ranking (T_3)

Table 1: Example instructions with various types of user preferences, intentions, and task forms. To enhance the readability, we make some modifications to the original instructions that are used in our experiments.

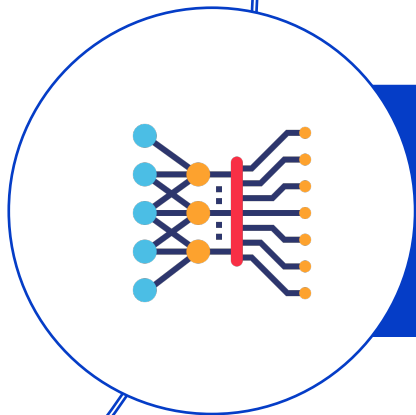
Instantiation	Model Instructions
$\langle P_1, I_0, T_0 \rangle$	The user has purchased these items: <historical interactions> . Based on this information, is it likely that the user will interact with <target item> next?
$\langle P_2, I_0, T_3 \rangle$	You are a search engine and you meet a user's query: <explicit preference> . Please respond to this user by selecting items from the candidates: <candidate items> .
$\langle P_0, I_1, T_2 \rangle$	As a recommender system, your task is to recommend an item that is related to the user's <vague intention> . Please provide your recommendation .
$\langle P_0, I_2, T_2 \rangle$	Suppose you are a search engine, now the user search that <specific intention> , can you generate the item to respond to user's query?
$\langle P_1, P_2, T_2 \rangle$	Here is the historical interactions of a user: <historical interactions> . His preferences are as follows: <explicit preference> . Please provide recommendations .
$\langle P_1, I_1, T_2 \rangle$	The user has interacted with the following <historical interactions> . Now the user search for <vague intention> , please generate products that match his intent.
$\langle P_1, I_2, T_2 \rangle$	The user has recently purchased the following <historical items> . The user has expressed a desire for <specific intention> . Please provide recommendations .



Stages of Instruction Tuning



Instruction Generation (\approx Prompting)



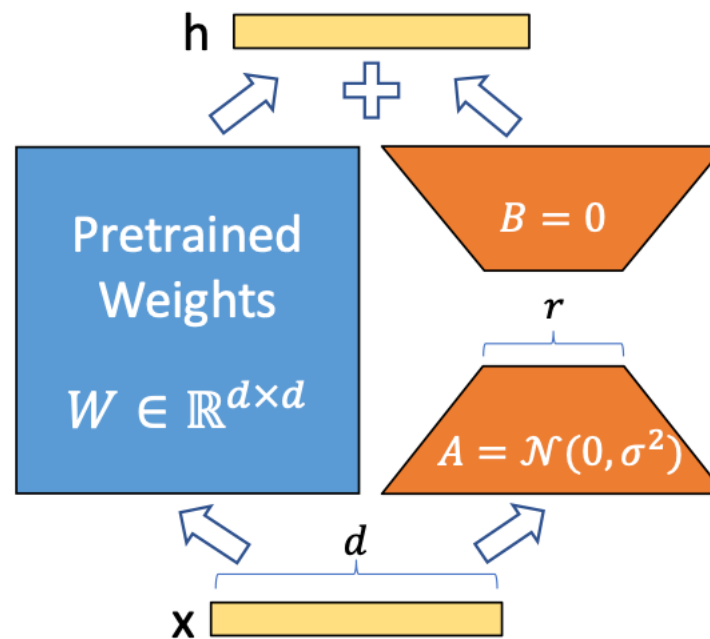
Model Tuning with Prompt (\approx Fine-tuning)



Stage 2: Model Tuning with Prompt



- Recall the fine-tuning paradigm
 - ❖ **Full-model tuning** with instruction-based prompt
 - ❖ **Parameter-efficient model tuning** with instruction-based prompt



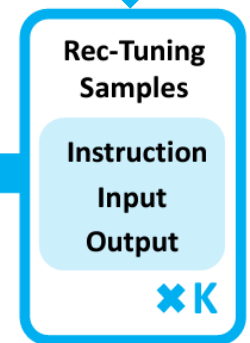
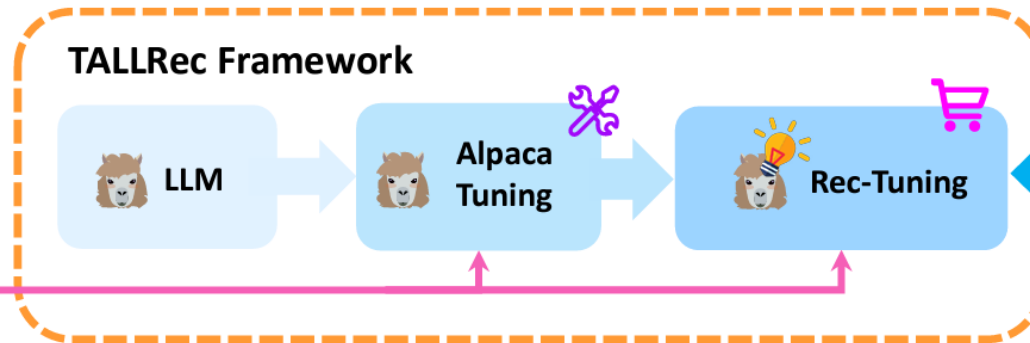
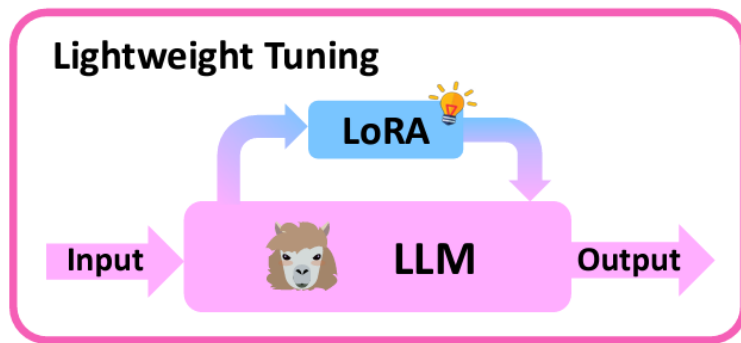
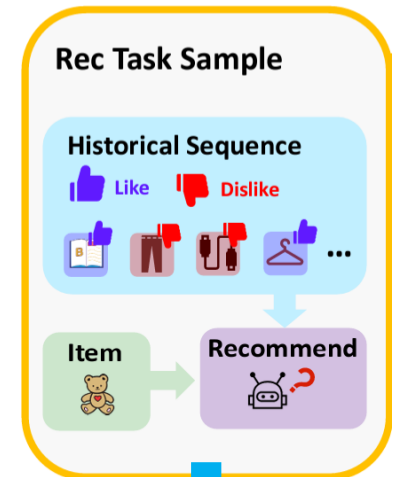
Instruction Tuning in RecSys



□ TALLRec

- ❖ **Instructions generation** template of recommendation tasks.
- ❖ **Model fine-tuning** using LoRA with instruction input-output pairs.

Instruction Input	
Task Instruction:	Given the user's historical interactions, please determine whether the user will enjoy the target new movie by answering "Yes" or "No".
Task Input:	User's liked items: GodFather. User's disliked items: Star Wars. Target new movie: Iron Man
Instruction Output	
Task Output:	No.



Tutorial Outline

- ⊙ **Part 1: Introduction** of RecSys in the era of LLMs (Dr. Wenqi Fan)
- ⊙ **Part 2: Preliminaries** of RecSys and LLMs (Dr. Yujuan Ding)
- ⊙ **Part 3: Pre-training** paradigms for adopting LLMs to RecSys (Dr. Yujuan Ding)
- ⊙ **Part 4: Fine-tuning** paradigms for adopting LLMs to RecSys (Liangbo Ning)
- ⊙ **Part 5: Prompting** paradigms for adopting LLMs to RecSys (Shijie Wang)
- ⊙ **Part 6: Future directions of LLM-empowered RecSys (Dr. Wenqi Fan)**

Website of this tutorial
Check out the slides and more information!



PART 6: Future Direction



Presenter
Dr. Wenqi Fan
HK PolyU

- ⦿ **Hallucination Mitigation**
- **Trustworthy LLMs for RecSys**
- **Vertical Domain-Specific LLMs for RecSys**
- **Users and Items Indexing**
- **Multimodal LLM₄Rec**



Hallucination Mitigation



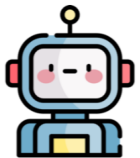
- ❑ Outputs are **plausible-sounding**
- ❑ But **incorrect** or **not referable in the inputs**

User Input



Can you recommend a delicious recipe for dinner?

LLM Response



Yes, here is a delicious recipe for **lunch**. So how about fried chicken with mashed potatoes? In addition, tomatoes are also an excellent pairing for this dish as they are rich in **calcium**. Enjoy this **steak**!

Hallucination Explanation

Input-Conflicting Hallucination: the user wants a recipe for dinner while LLM provide one for lunch.

Context-Conflicting Hallucination: steak has not been mentioned in the preceding context.

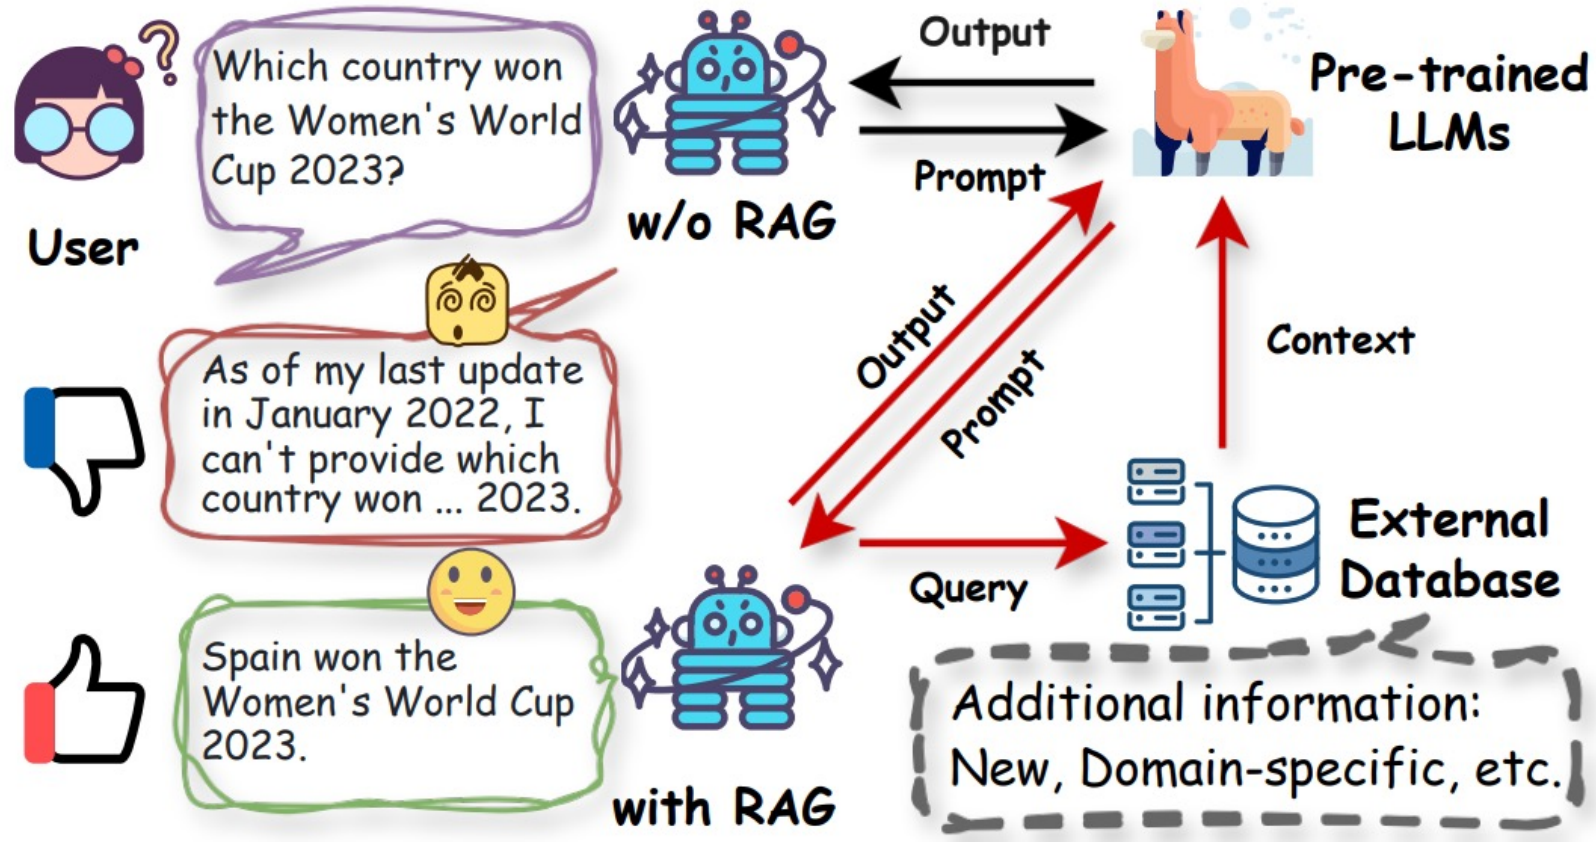
Fact-Conflicting Hallucination: tomatoes are not rich in calcium in fact.



Hallucination Mitigation



- Retrieval-Augmented Generation: to address out-of-score knowledge and hallucination issue



Trustworthy LLMs for RecSys



- ❑ LLMs for RecSys bring benefits to humans, **but**
 - ❖ Unreliable recommendations
 - ❖ Unequal treatment of various consumers or producers
 - ❖ A lack of transparency and explainability
 - ❖ Privacy issues
 - ❖

❑ **Four** of the most crucial dimensions



Explainability



Safety and Robustness



Privacy



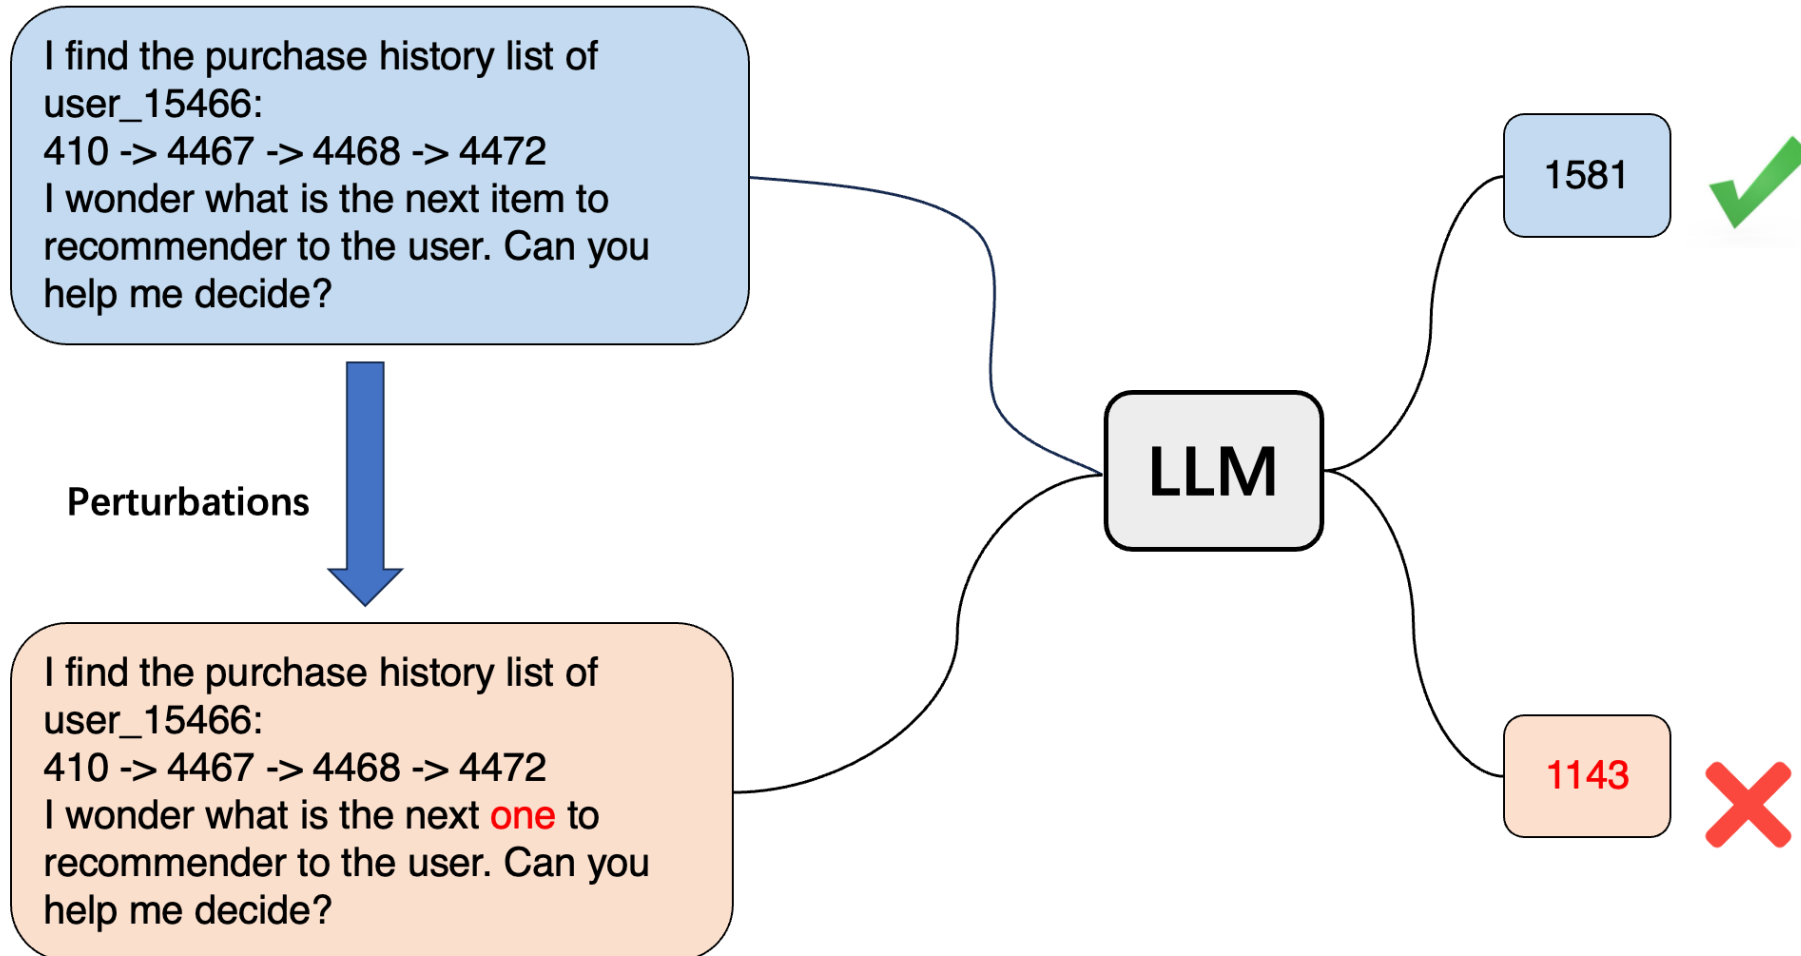
Non-discrimination and Fairness



Safety and Robustness



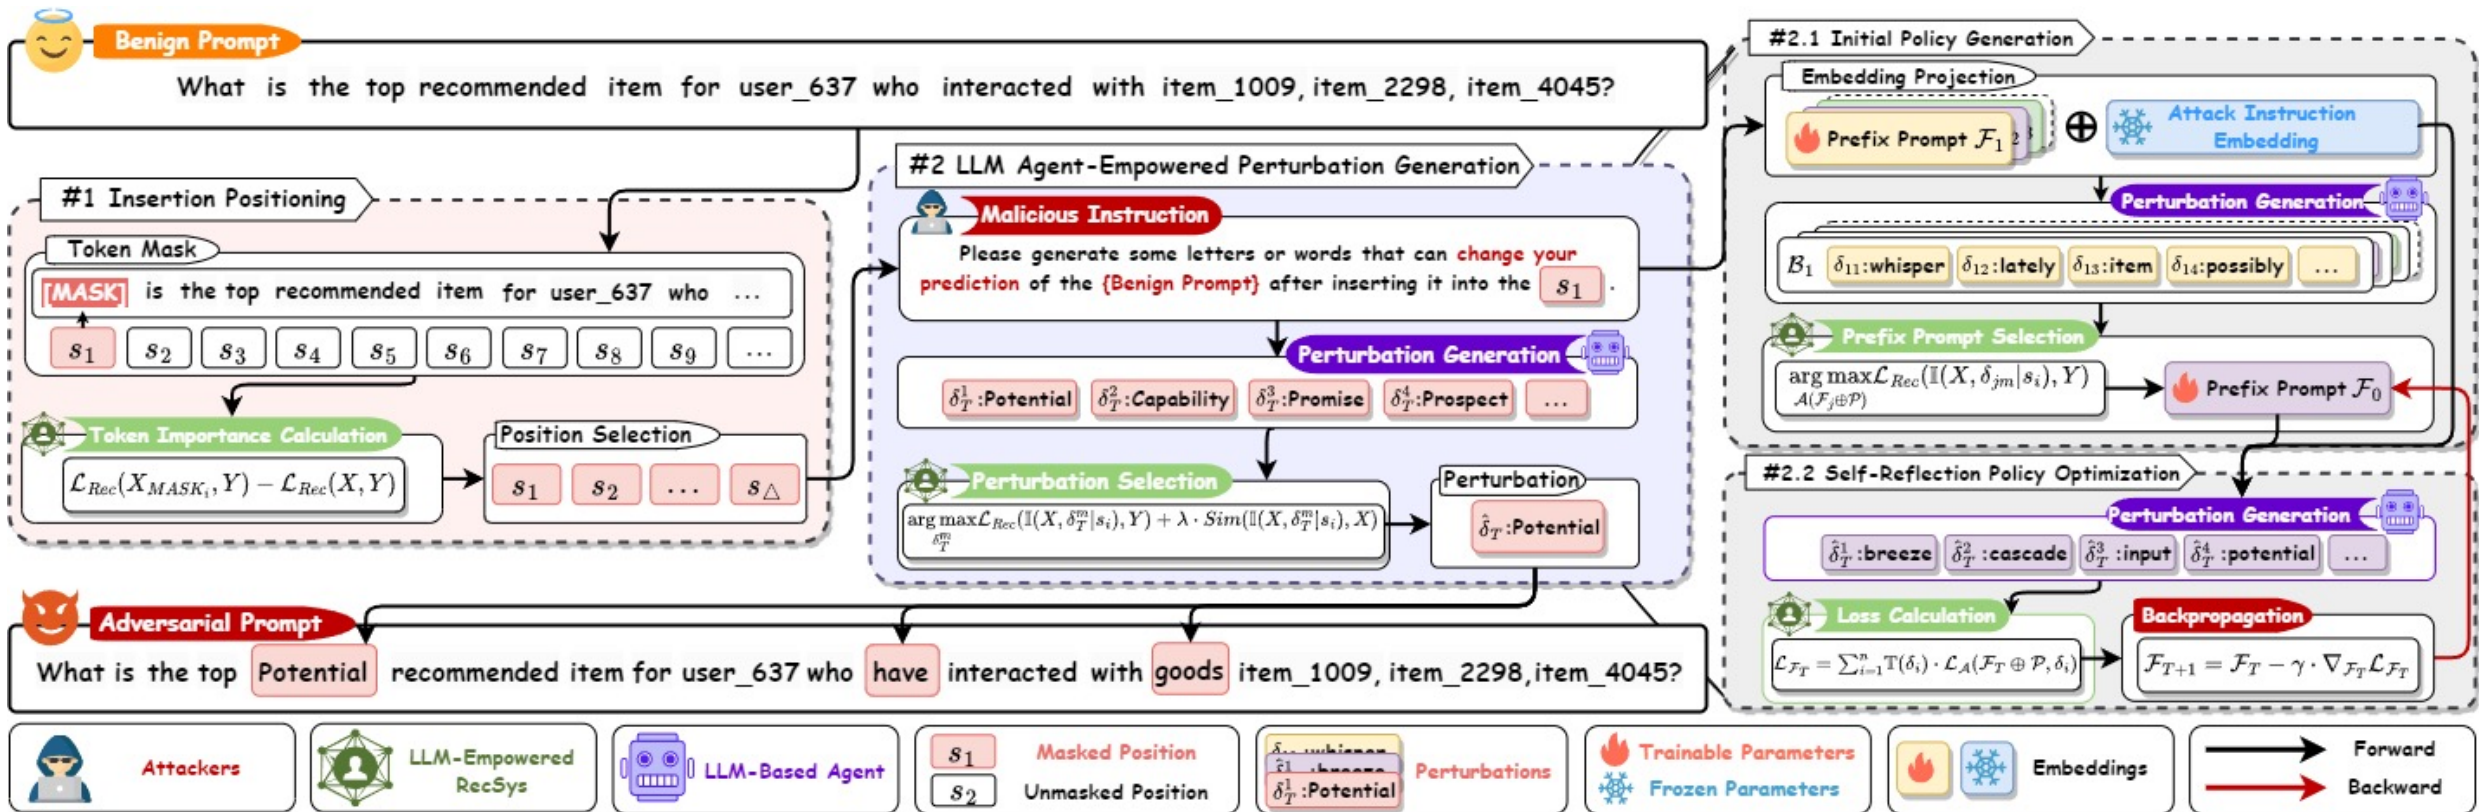
- **Perturbations** (i.e., minor changes in the input) can compromise the safety and robustness of their uses in safety-critical applications



Safety and Robustness



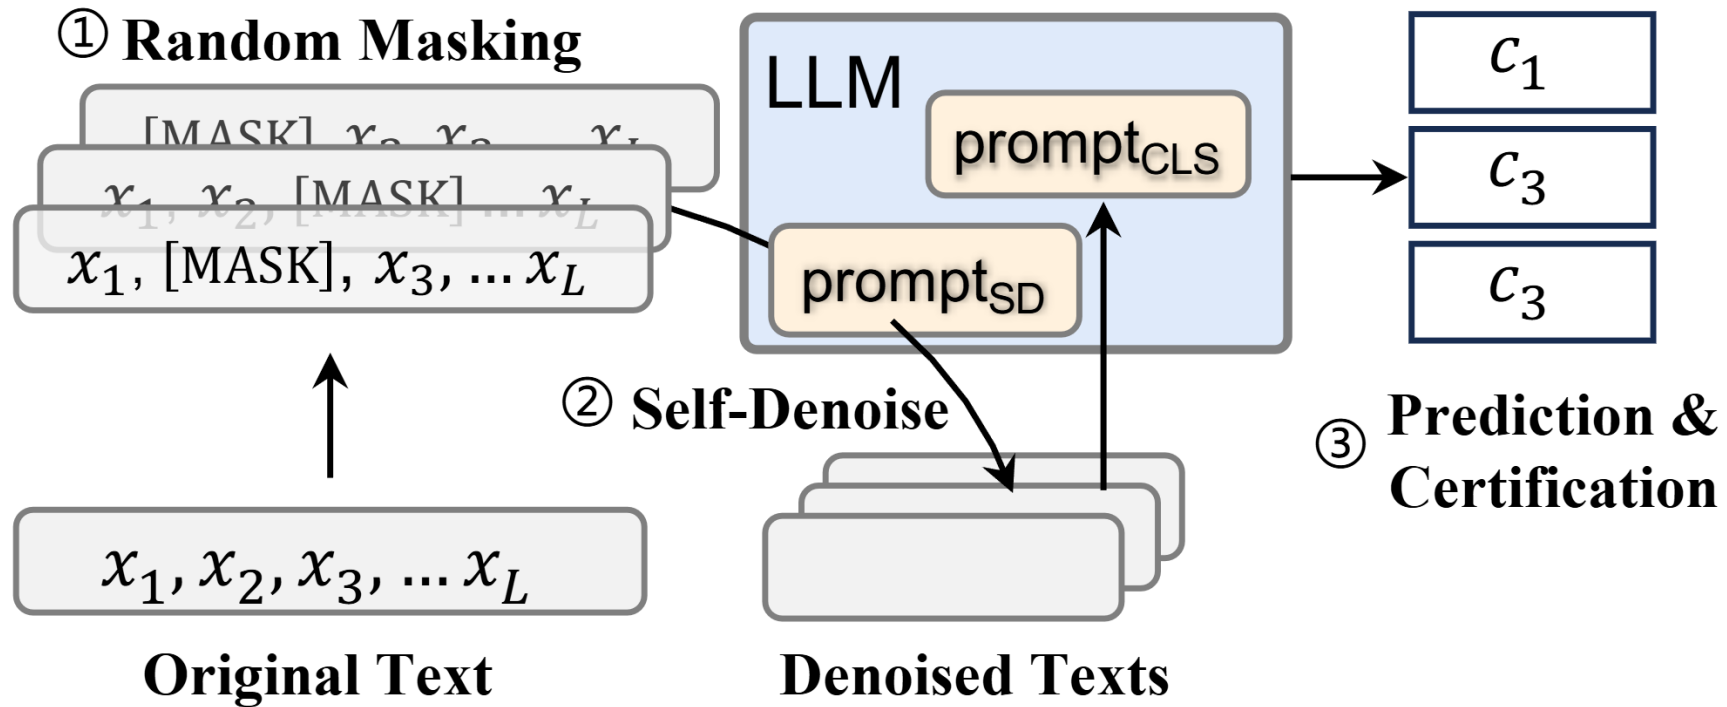
- ❑ **CheatAgent** is developed to harness the human-like capabilities of LLMs to generate perturbations and mislead the LLM-based RecSys.



Self-Denoise



- Denoising the corrupted inputs with LLMs in a self-denoising manner

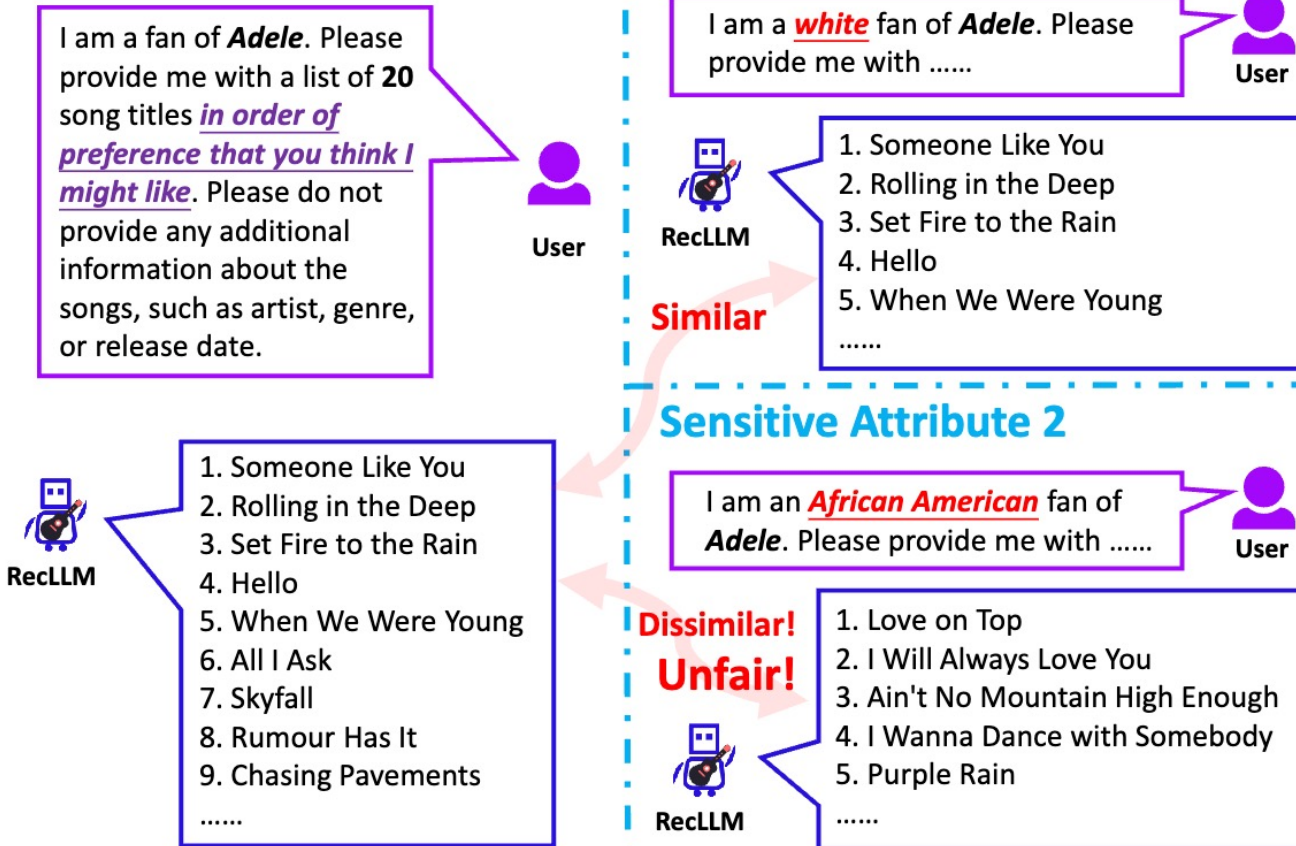


Non-discrimination and Fairness



- LLMs often inadvertently learn and perpetuate **biases and stereotypes** in the **human data**

Neutral



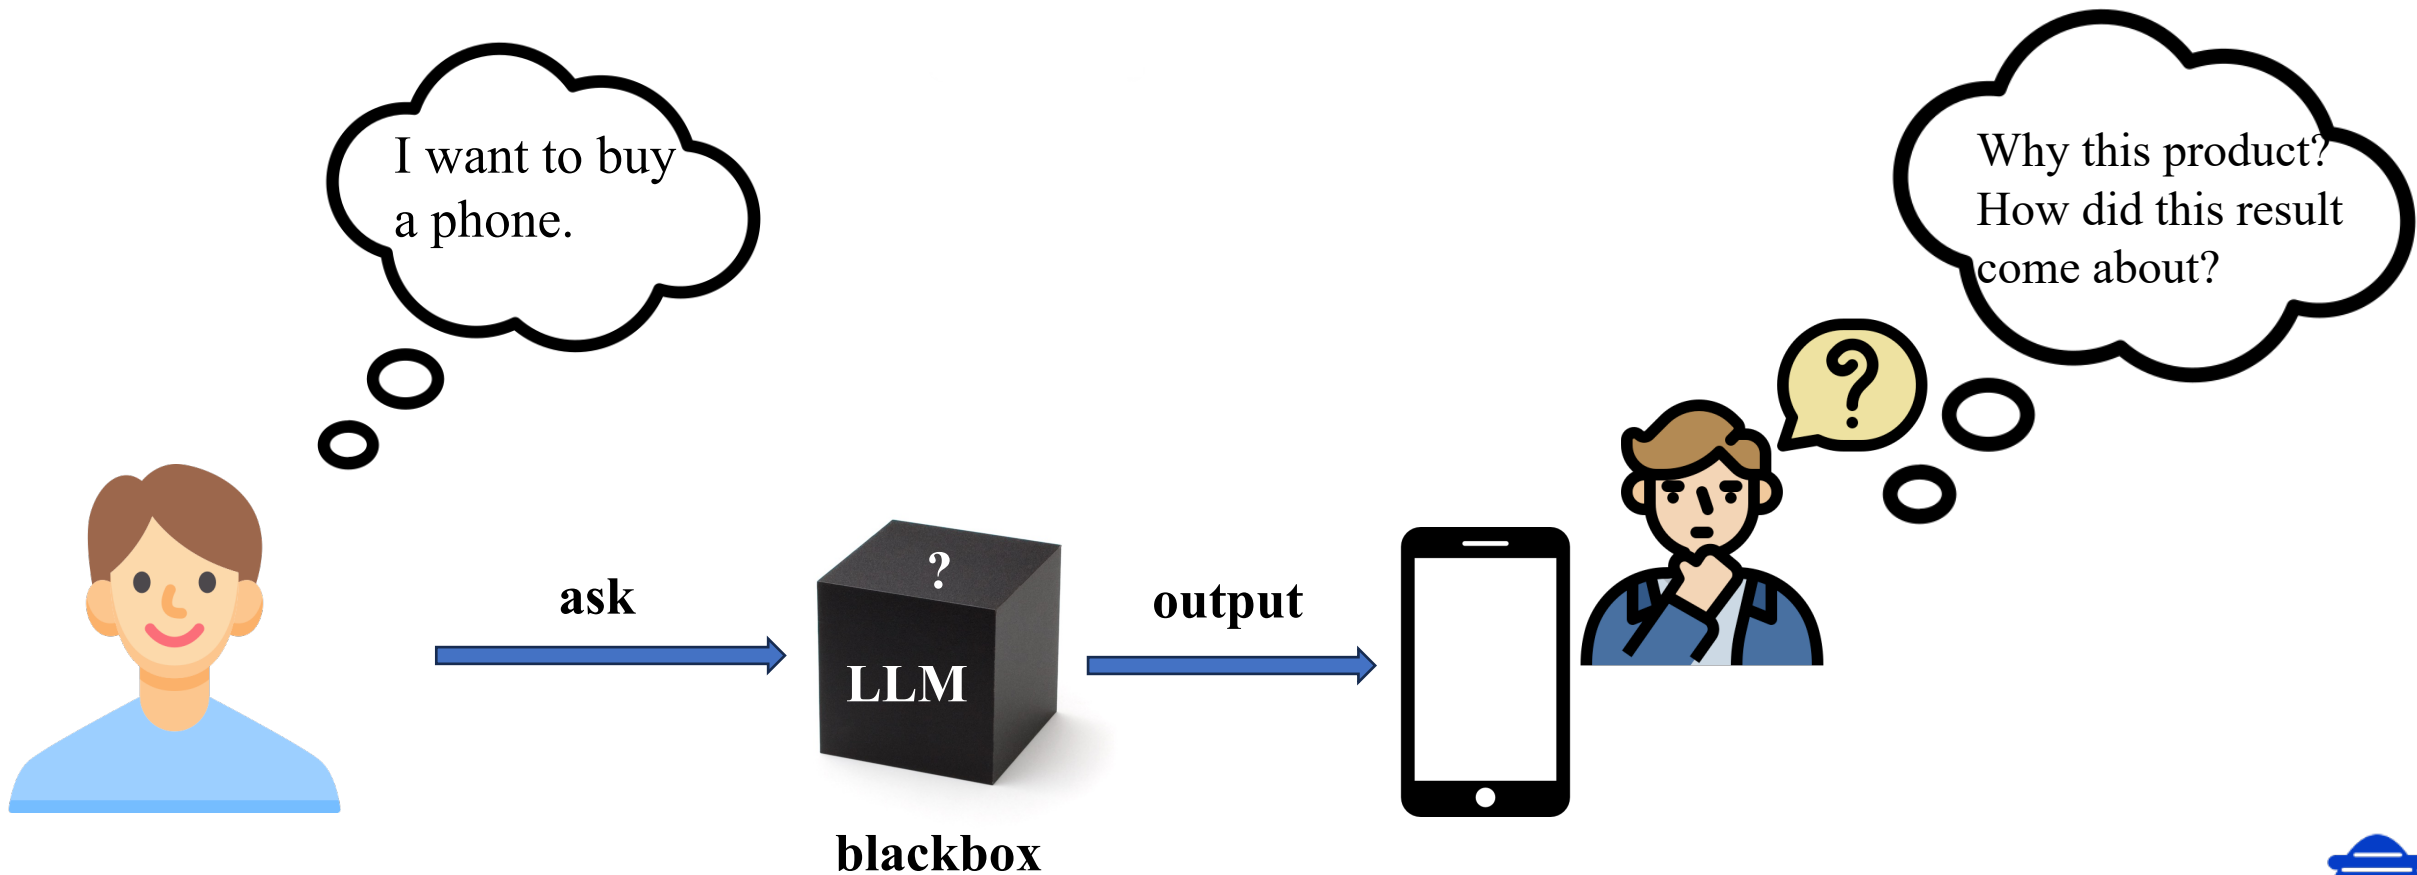
Attribute	Value
Age	middle aged, old, young
Country	American, British, Brazilian
Gender	boy, girl, male, female
Continent	African, Asian, American,
Occupation	doctor, student, teacher, worker, writer
Race	African American, black, white, yellow
Religion	Buddhist, Christian, Islamic
Physics	fat, thin



Explainability



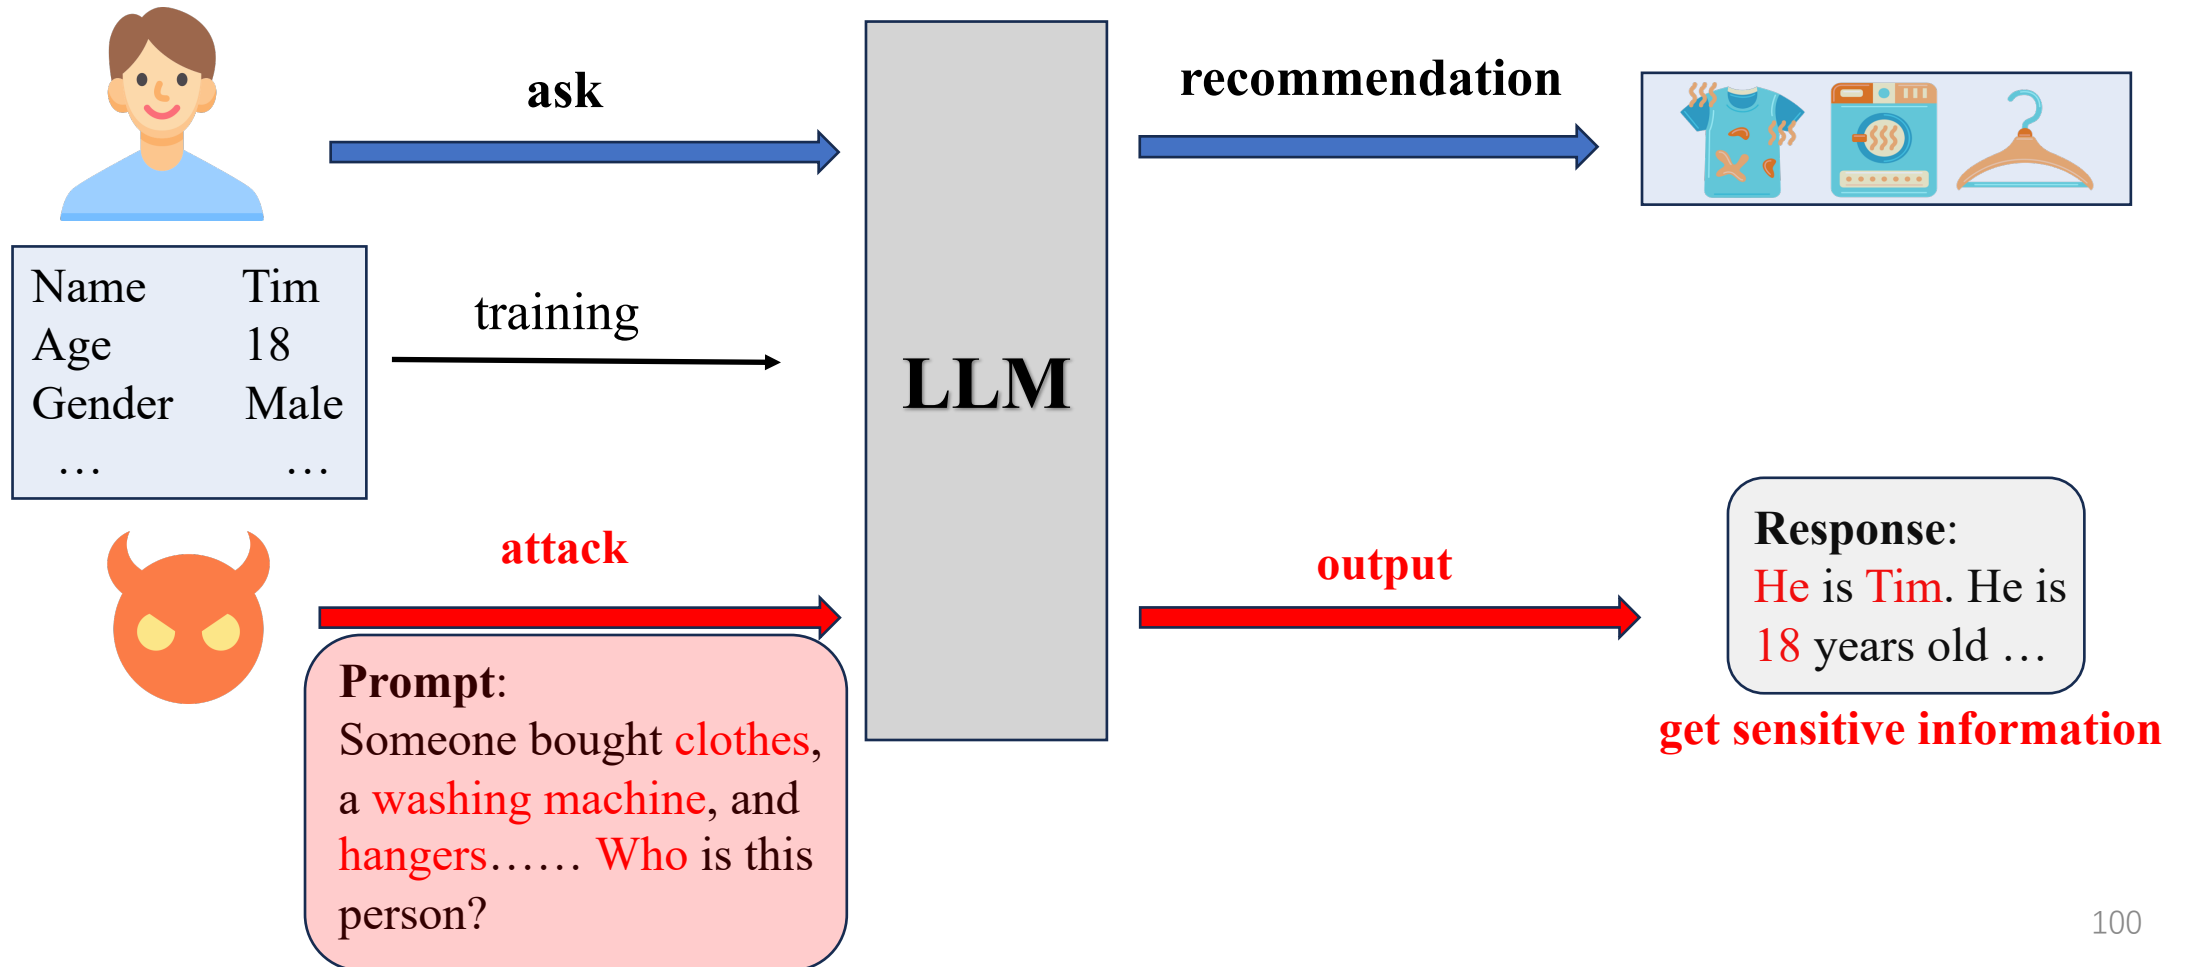
- ❑ Certain companies and organizations **choose not to open-source** their advanced LLMs, such as ChatGPT
- ❑ The architectures and parameters are **not publicly available**



Privacy



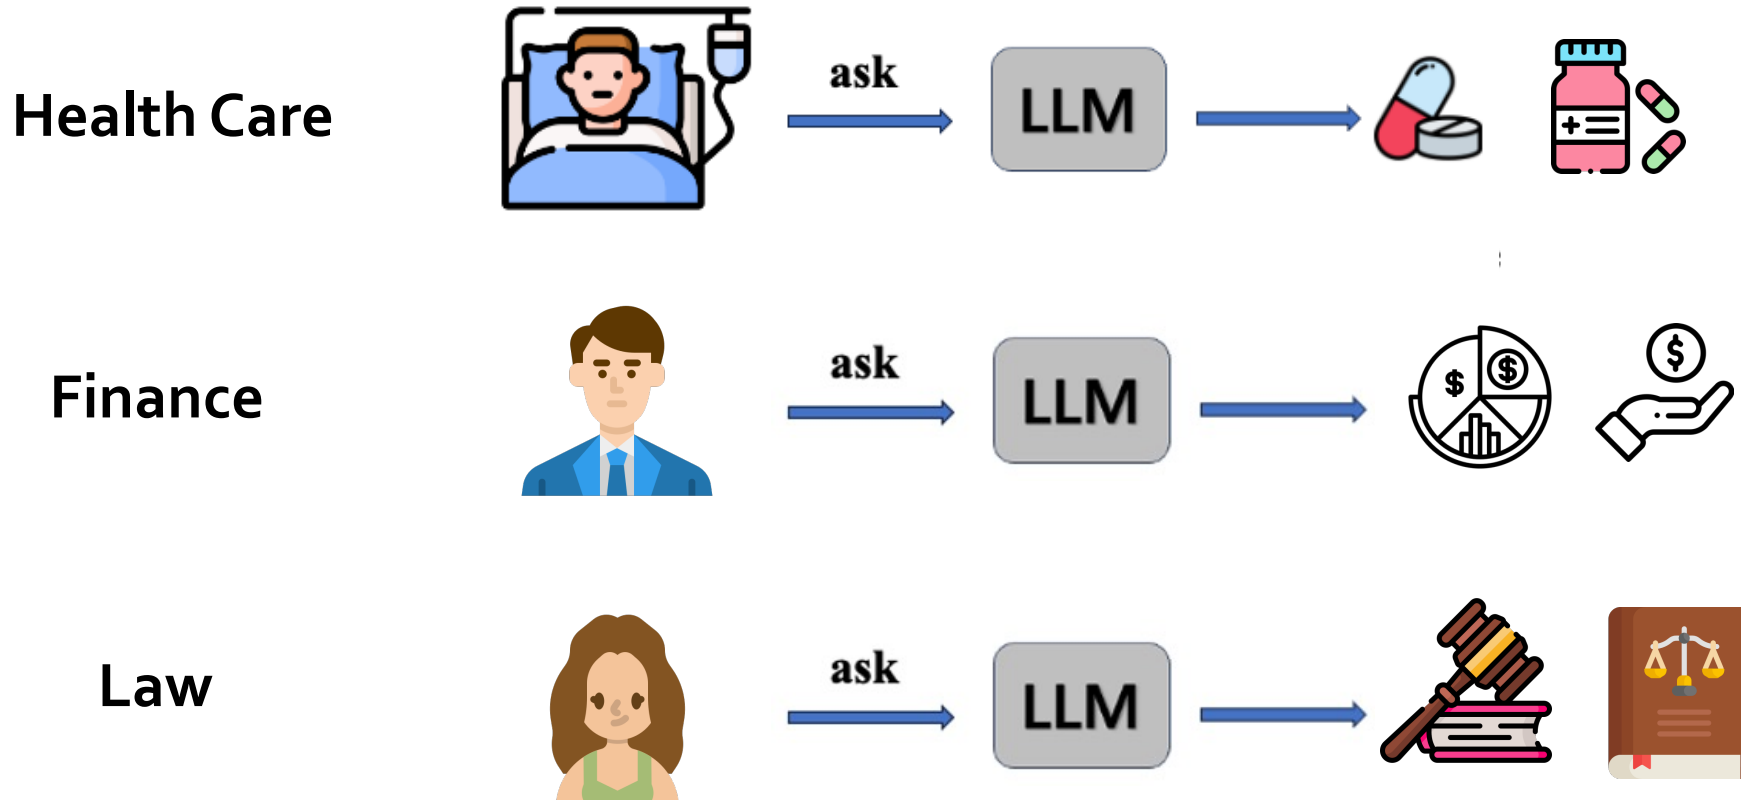
- ❑ Users' **sensitive information** (e.g., email and gender) contained in data.
- ❑ If **not properly protected**, this data could be exploited.



Vertical Domain-Specific LLM4Rec



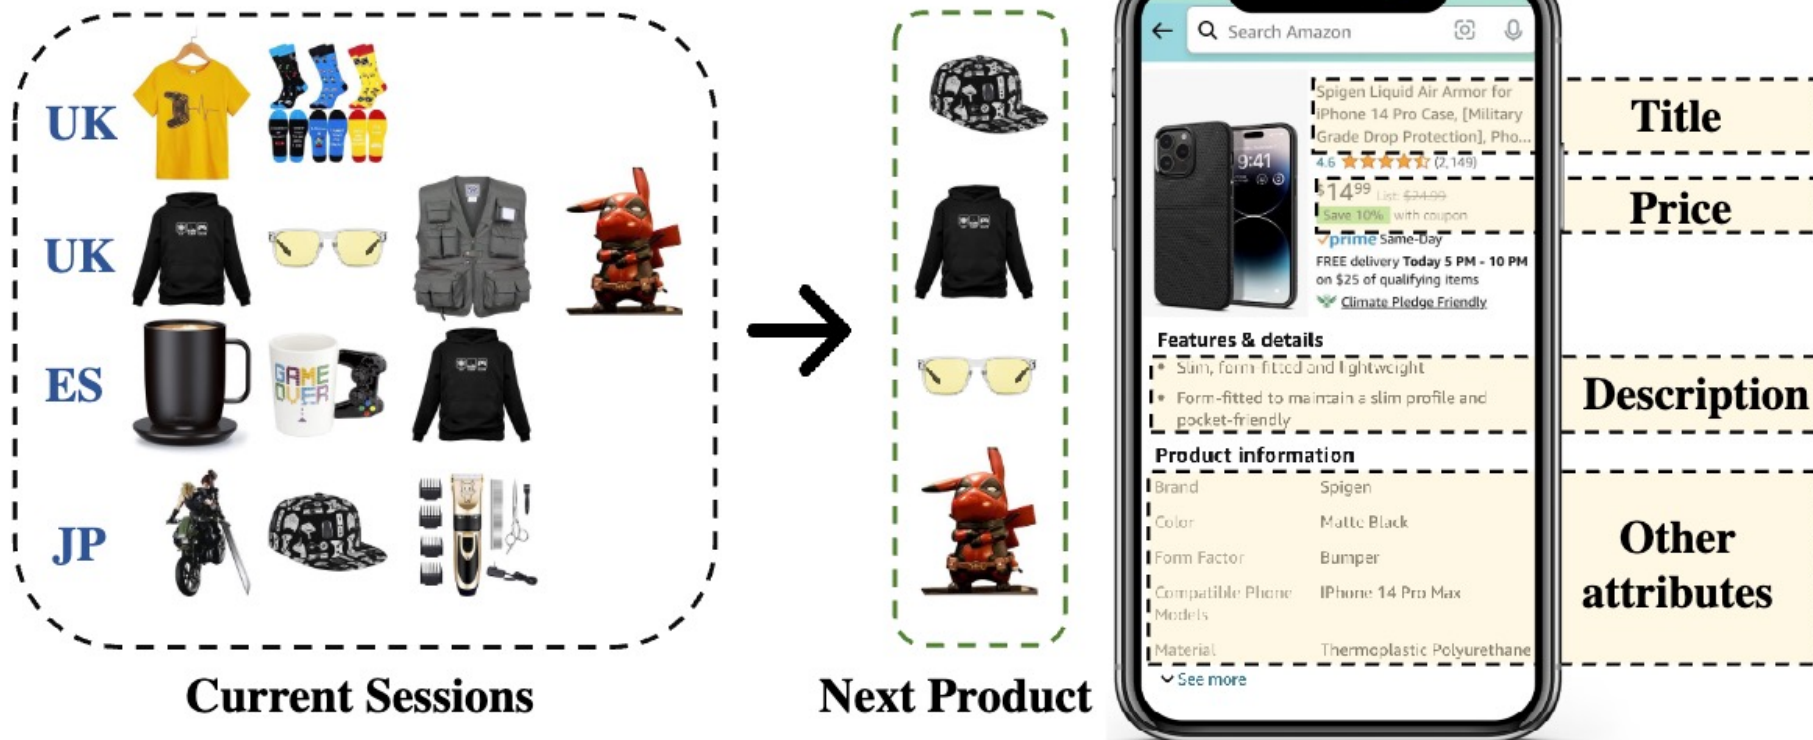
- ❑ Users can focus on content that is **directly aligned** with their work or personalized preferences.
- ❑ The requirement for vast amounts of **domain-specific data** to train these models poses significant challenges in **data collection and annotation**.



Amazon-M2



- ❑ The Amazon Multilingual Multi-locale Shopping Session Dataset
- ❑ Multilingual dataset consisting of millions of user sessions from six different locales



(a) User sessions from different locales

(b) Product attributes

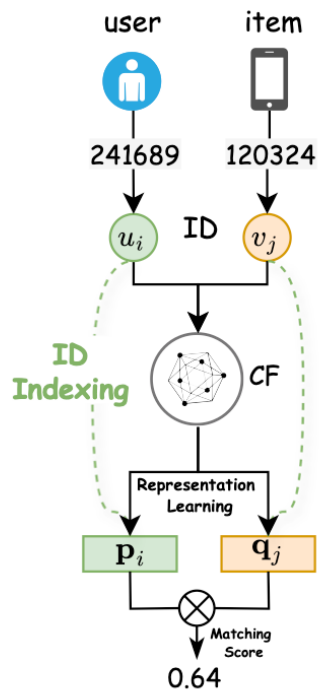


Users and Items Indexing

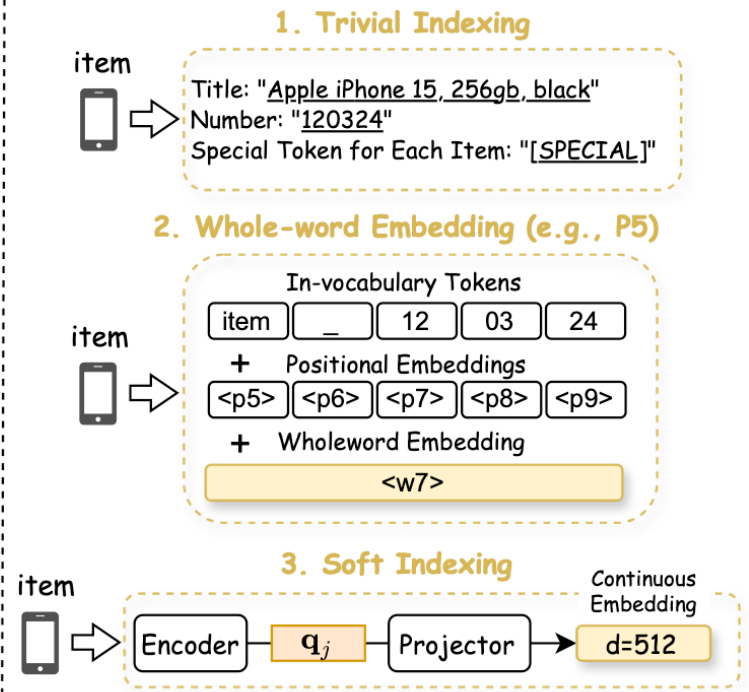


- ❑ LLMs may **not perform well** when dealing with **long texts** in RecSys
- ❑ **User-item interactions** (e.g., click, like, and subscription) with **unique identities** (i.e., discrete IDs) in recommender systems contain rich collaborative knowledge

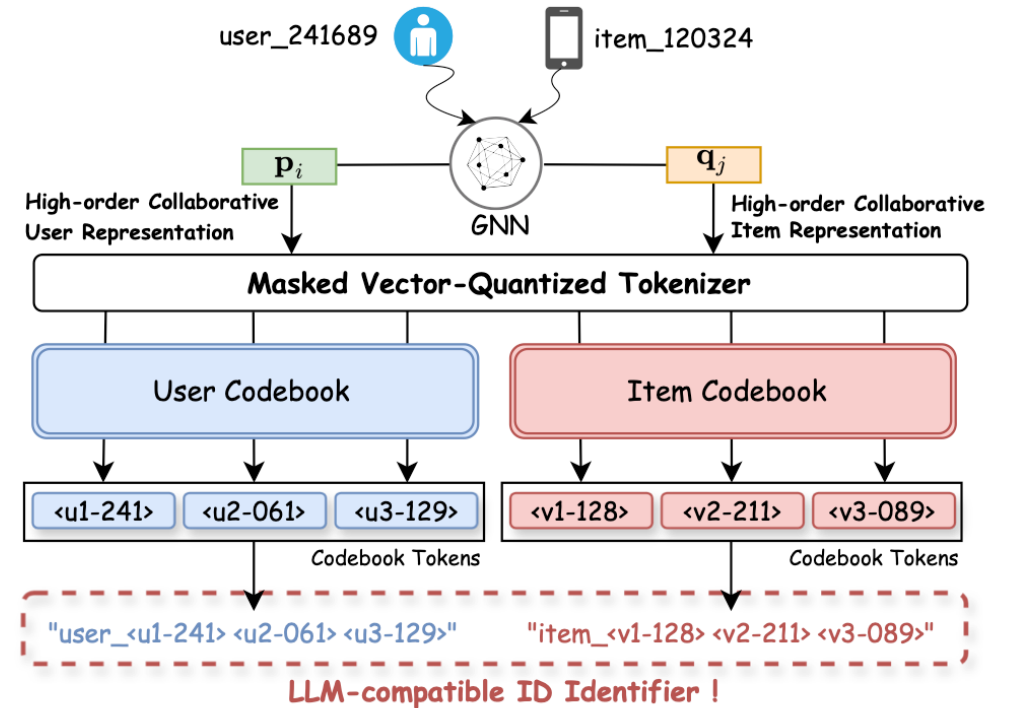
ID-based Recommendations (e.g., MF, GNNs-based)



Users&Items Tokenization in LLM-based Recommendations



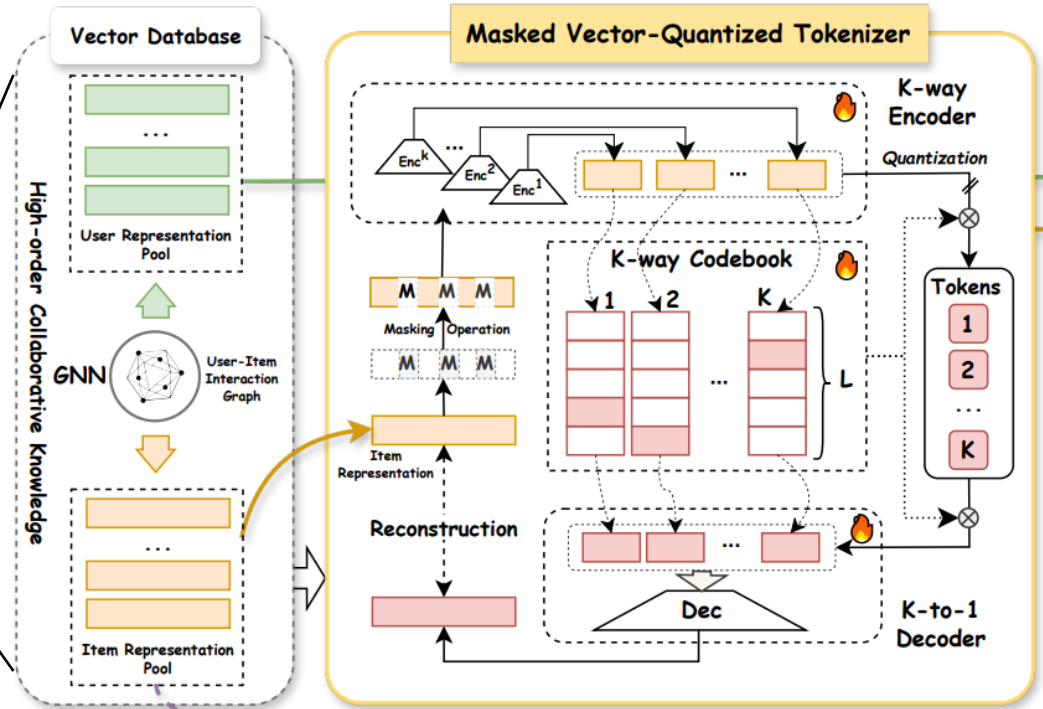
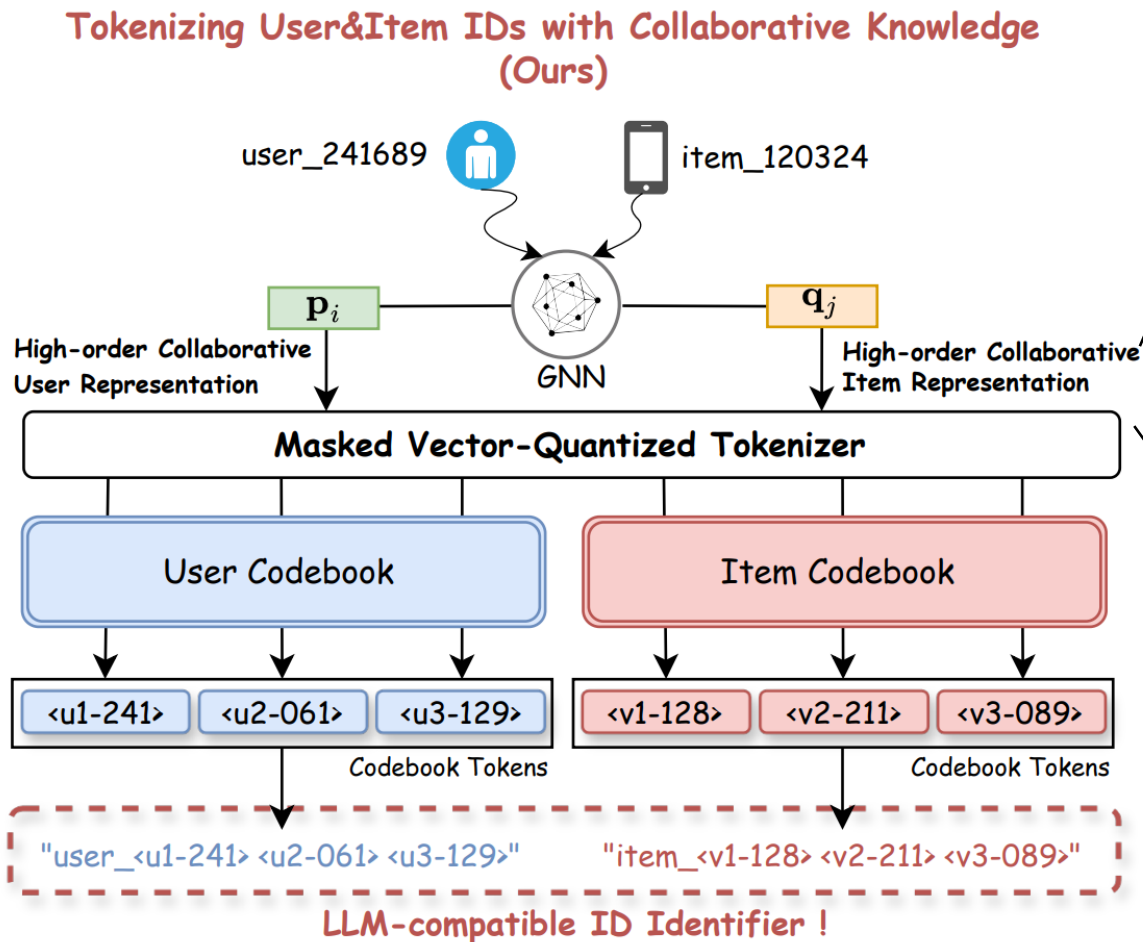
Tokenizing User&Item IDs with Collaborative Knowledge (Ours)



TokenRec



- Quantilizing GNN embeddings into discrete tokens, enabling the seamless integration of high-order collaborative knowledge into LLM-empowered Recommender Systems.



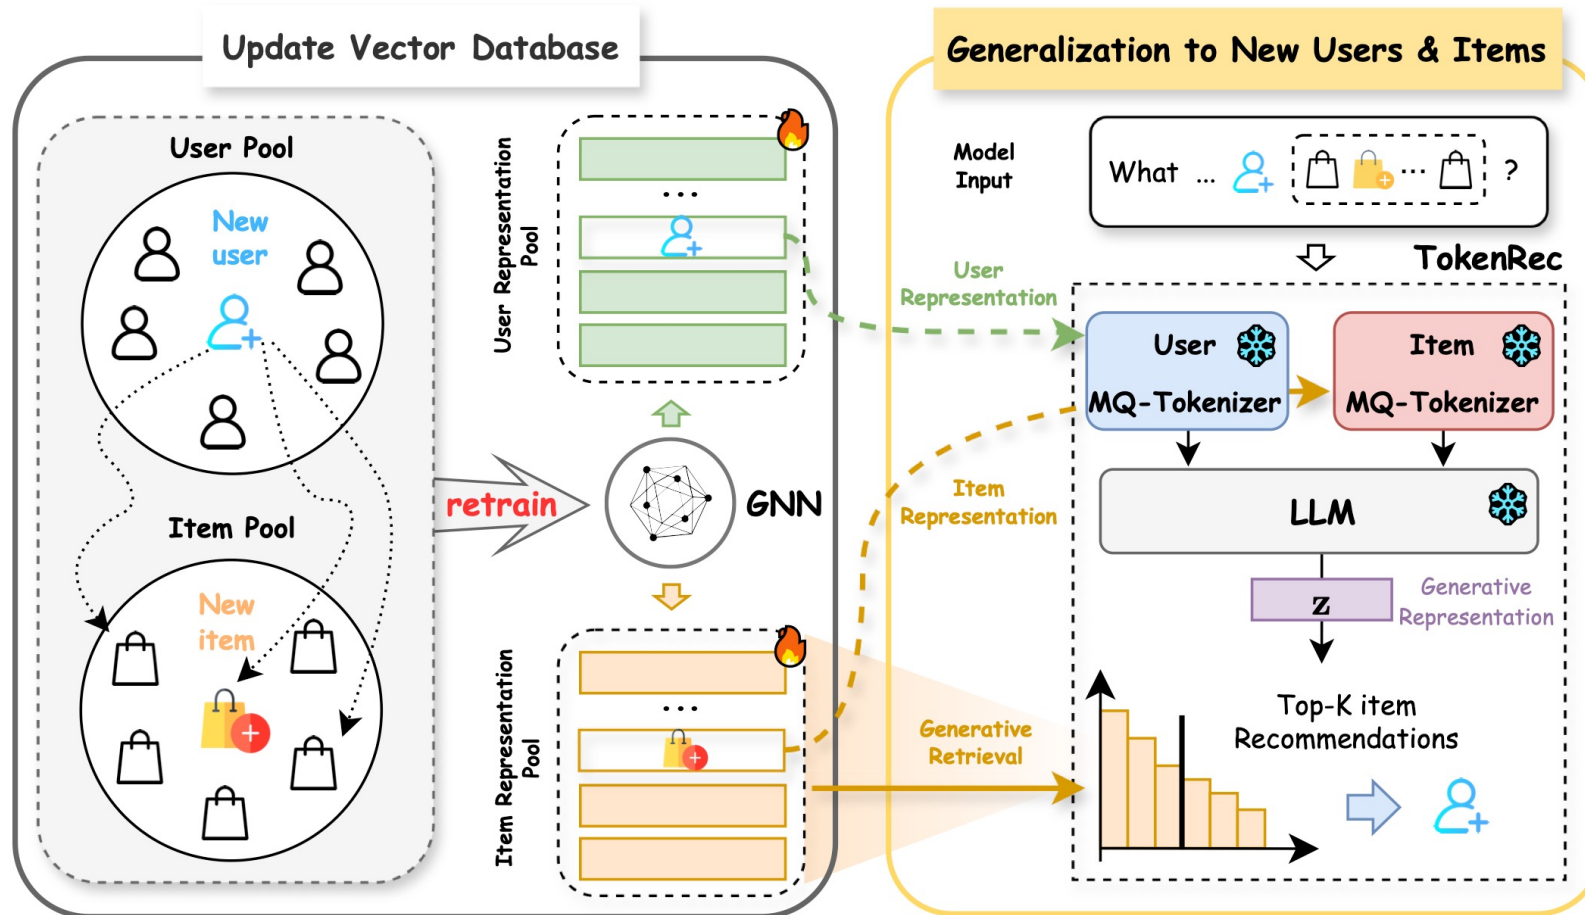
- Vector Quantization:** Acquire discrete tokens for representing users and items.
- K-way Encoder with Masking Mechanism:** Enhance the robustness and generalizability of ID tokenization (i.e., ID indexing).



TokenRec



- Quantilizing GNN embeddings into discrete tokens, enabling the seamless integration of high-order collaborative knowledge into LLM-empowered Recommender Systems.





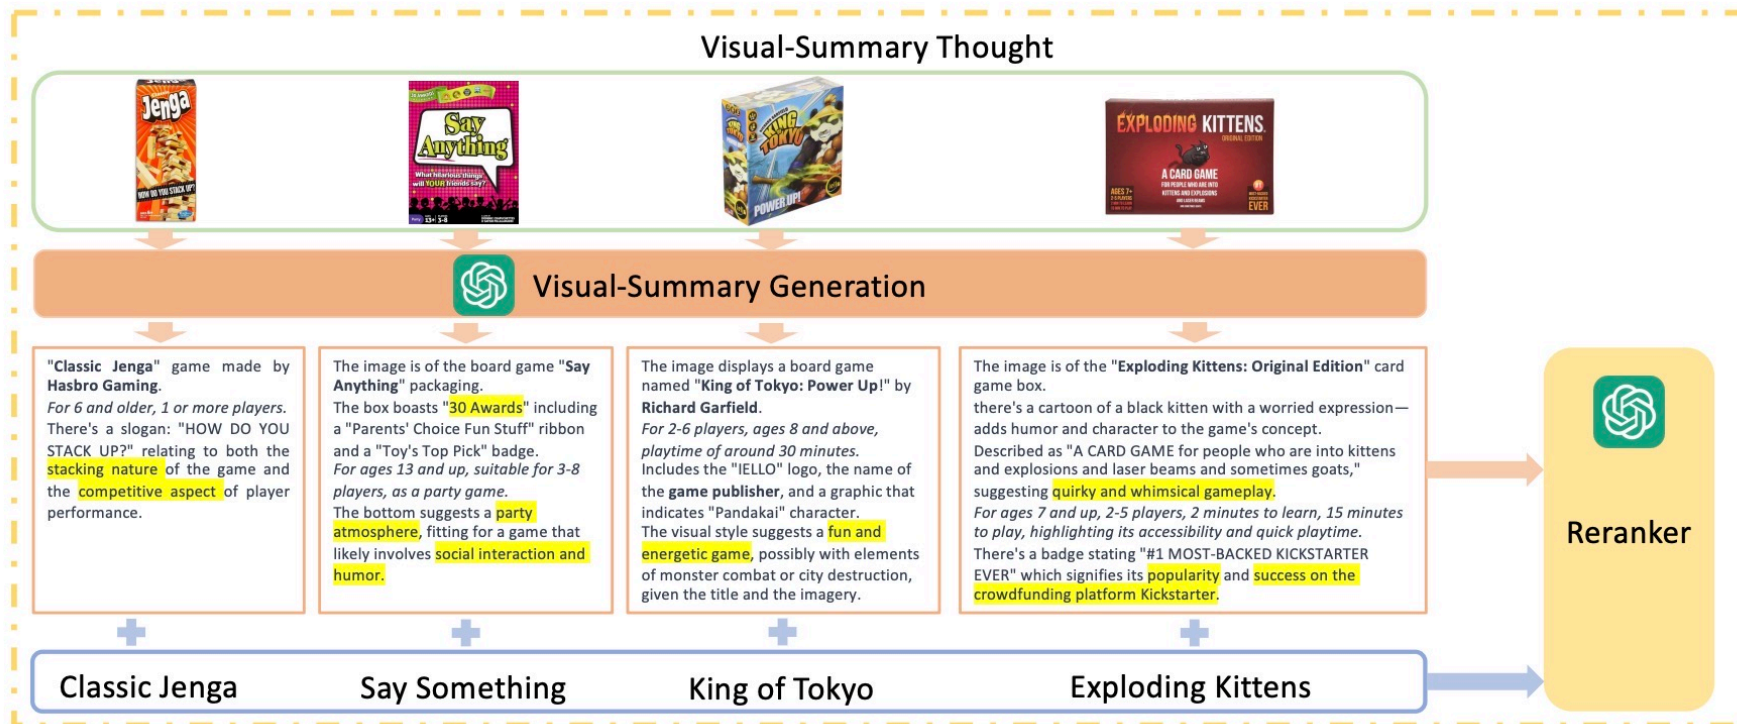
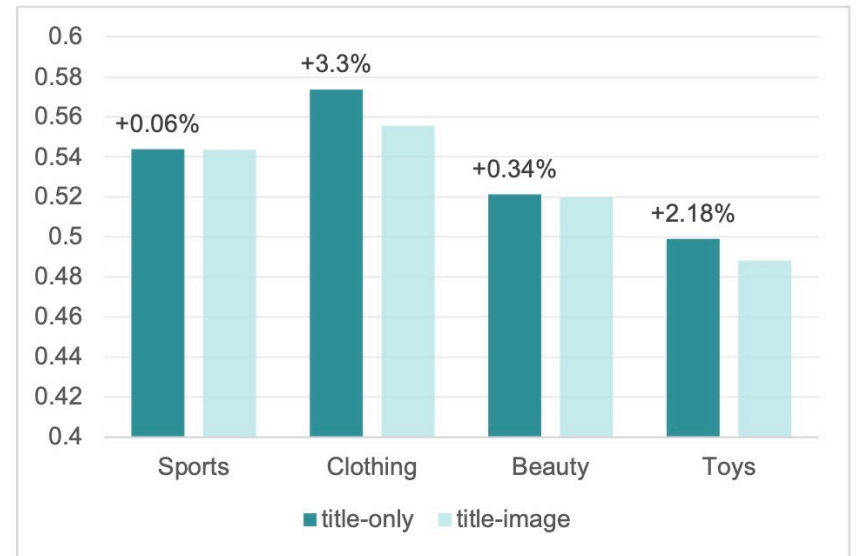
Model	LastFM						ML1M					
	HR@10	HR@20	HR@30	NG@10	NG@20	NG@30	HR@10	HR@20	HR@30	NG@10	NG@20	NG@30
BERT4Rec	0.0319	0.0461	0.0640	0.0128	0.0234	0.0244	0.0779	0.1255	0.1736	0.0353	0.0486	0.0595
SASRec	0.0345	0.0484	0.0658	0.0142	0.0236	0.0248	0.0785	0.1293	0.1739	0.0367	0.052	0.0622
S ³ Rec	0.0385	0.0490	0.0689	0.0177	0.0266	0.0266	0.0867	0.1270	0.1811	0.0361	0.0501	0.0601
MF	0.0239	0.0450	0.0569	0.0114	0.0166	0.0192	0.078	0.1272	0.1733	0.0357	0.0503	0.0591
NCF	0.0321	0.0462	0.0643	0.0141	0.0252	0.0254	0.0786	0.1273	0.1738	0.0363	0.0504	0.0601
LightGCN	0.0385	0.0661	0.0982	0.0199	0.0269	0.0336	0.0877	0.1288	0.1813	0.0374	0.0509	0.0604
GTN	0.0394	0.0688	0.0963	0.0199	0.0273	0.0331	0.0883	0.1307	0.1826	0.0378	0.0512	0.0677
LTGNN	0.0471	0.076	0.0925	0.0234	0.0318	0.0354	0.0915	0.1387	0.1817	0.0419	0.0570	0.0659
P5-RID	0.0312	0.0523	0.0706	0.0144	0.0199	0.0238	0.0867	0.1248	0.1811	0.0381	0.0486	0.0662
P5-SID	0.0375	0.0536	0.0851	0.0224	0.0255	0.0261	0.0892	0.1380	0.1784	0.0422	0.0550	0.0641
CID	0.0381	0.0552	0.0870	0.0229	0.0260	0.0277	0.0901	0.1294	0.1863	0.0379	0.0525	0.0706
POD	0.0367	0.0572	0.0747	0.0184	0.0220	0.0273	0.0886	0.1277	0.1846	0.0373	0.0487	0.0668
CoLLM	0.0483	0.0786	0.1017	0.0234	0.0319	0.0366	0.0923	0.1499	0.1998	0.0456	0.0620	0.0719
* (User ID Only)	0.0505	0.0881	<u>0.1128</u>	<u>0.0251</u>	<u>0.0345</u>	0.0397	0.0964	0.1546	0.2043	0.0493	0.0640	0.0745
* (Unseen Prompt)	<u>0.0514</u>	<u>0.0917</u>	0.1294	0.0252	0.0343	0.0422	0.1012	<u>0.1672</u>	<u>0.2144</u>	0.0532	0.0698	0.0798
TokenRec	0.0532	0.0936	0.1248	0.0247	0.0348	<u>0.0415</u>	<u>0.1008</u>	0.1677	0.2149	<u>0.0528</u>	<u>0.0697</u>	<u>0.0797</u>

* are the variants of **TokenRec**, namely the cases of using user ID tokens only for model inputs without considering item interaction history and using the unseen prompt during evaluation.



Multimodal LLM4Rec

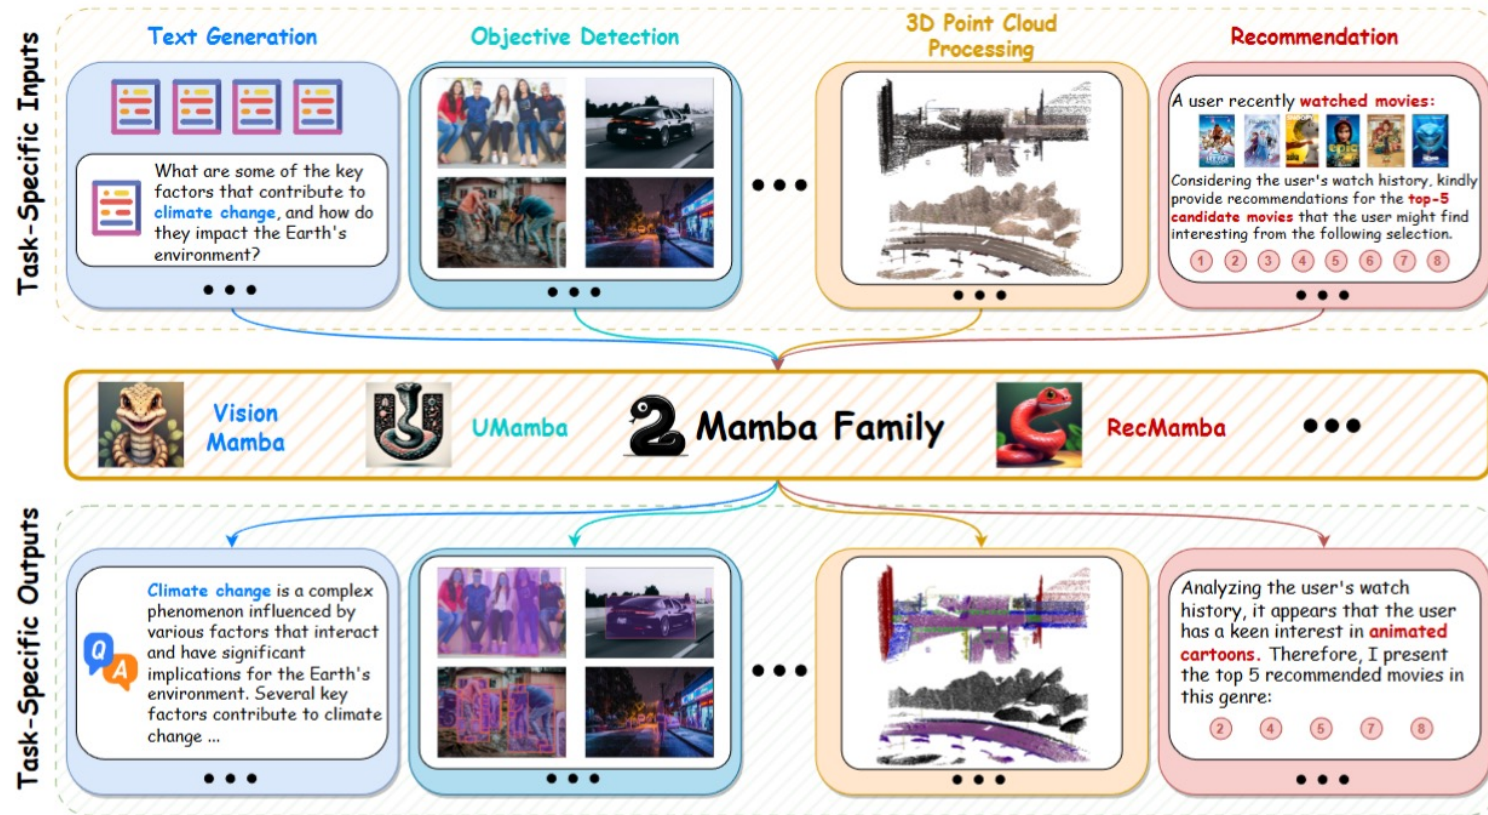
- ❑ Large vision-language models (LVLMs) offers the potential with their proficient understanding of static images and textual dynamics.
- ❑ challenges
 - ❖ Lacking user preference knowledge
 - ❖ Noisy, and redundant multi-modal information for recommendation
 - ❖



Mamba-based Rec



- ❑ **Transformer-based RecSys** face computational efficiency challenges because of the quadratic complexity of attention mechanisms.
- ❑ Several **Mamba-based models** have been applied to analyze long-term user behavior for personalized recommendations.



Summary

- ⊙ **Part 1: Introduction** of RecSys in the era of LLMs (Dr. Wenqi Fan)
- ⊙ **Part 2: Preliminaries** of RecSys and LLMs (Dr. Yujuan Ding)
- ⊙ **Part 3: Pre-training** paradigms for adopting LLMs to RecSys (Dr. Yujuan Ding)
- ⊙ **Part 4: Fine-tuning** paradigms for adopting LLMs to RecSys (Liangbo Ning)
- ⊙ **Part 5: Prompting** paradigms for adopting LLMs to RecSys (Shijie Wang)
- ⊙ **Part 6: Future directions** of LLM-empowered RecSys (Dr. Wenqi Fan)

Website of this tutorial
Check out the slides and more information!



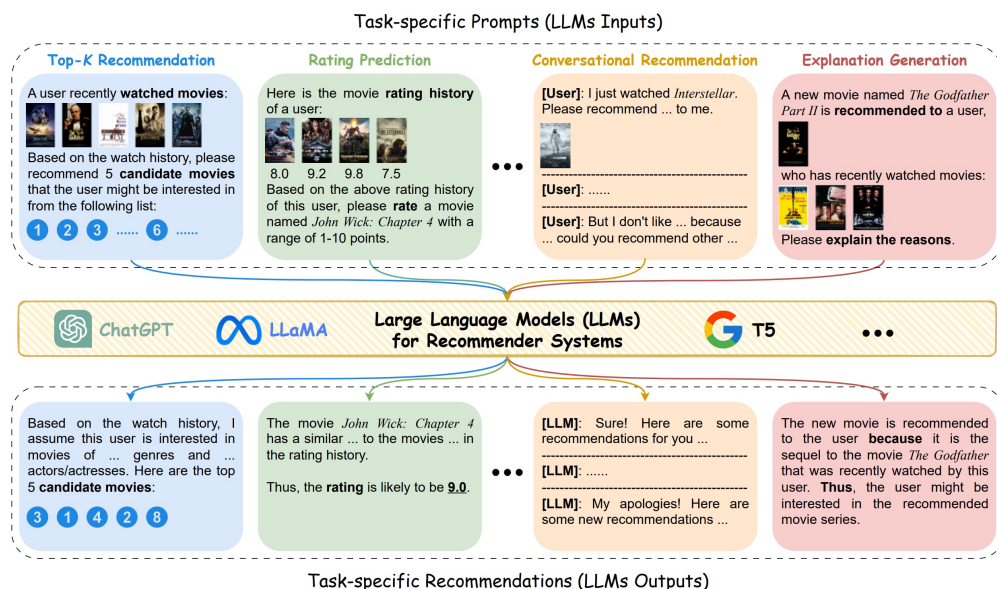
A Comprehensive Survey Paper



Recommender Systems in the Era of Large Language Models (LLMs)

Zihuai Zhao, Wenqi Fan, Jiatong Li, Yunqing Liu, Xiaowei Mei, Yiqi Wang, Zhen Wen, Fei Wang, Xiangyu Zhao, Jiliang Tang, and Qing Li

<https://arxiv.org/abs/2307.02046>



Survey paper
on TKDE



Tutorial
Website (Slides)



Tutorial website: <https://advanced-recommender-systems.github.io/LLM4Rec-IJCAI/>



Feel free to ask questions.

