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

  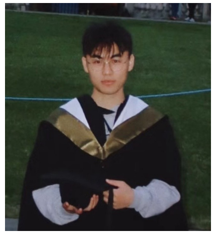  

Jiatong Li[1], Zihuai Zhao[1], Yunqing Liu[1], Yiqi Wang[2], Wenqi Fan[1]

[1]The Hong Kong Polytechnic University
[2]National University of Defense Technology

10:00 AM – 12:00 AM (UTC+8) Room 6, ICDM
Zoom ID: 91649466943, **Password:** 202312

Website QR Code

Website (Slides): https://advanced-recommender-systems.github.io/llms_rec_tutorial/
Survey Paper: "Recommender systems in the era of large language models (llms)." arXiv:2307.02046 (2023).

Tutorial Outline

Introduction → Preliminaries → Pre-training → Fine-tuning → Prompting → Future Directions

- **Introduction** of RecSys in the era of LLMs (Dr. Wenqi Fan)
- **Preliminaries** of RecSys and LLMs (Yunqing Liu)
- **Pre-training** paradigms for adopting LLMs to RecSys (Jiatong Li)
- **Fine-tuning** paradigms for adopting LLMs to RecSys (Jiatong Li)
- **Prompting** paradigms for adopting LLMs to RecSys (Zihuai Zhao)
- **Future directions** of LLM-empowered RecSys (Dr. Yiqi Wang)

Zoom ID:
91649466943
Password:
202312

Website QR Code

Recommender Systems (RecSys)

Age of Information Explosion

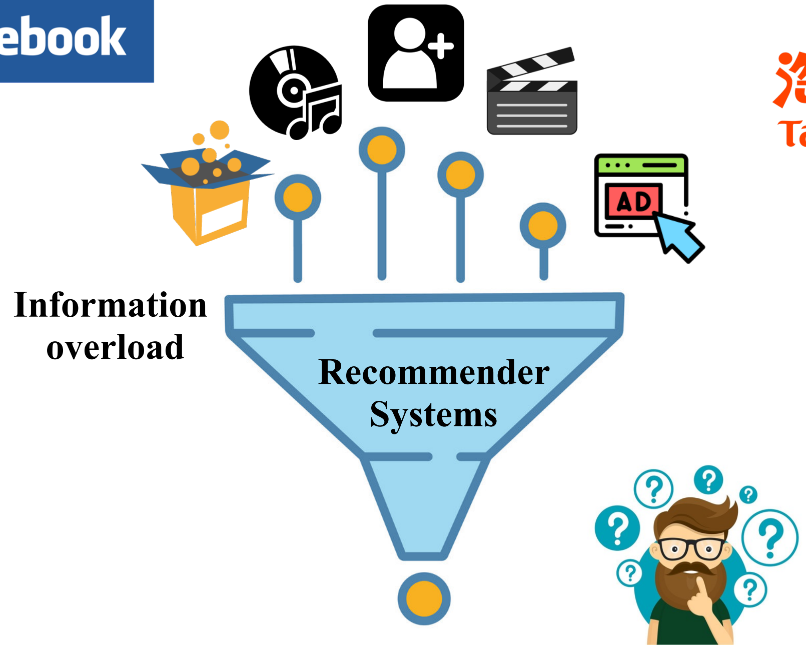

Recommend item X to user

Items can be: Products, Friends, News, Movies, Videos, etc.

Recommender Systems (RecSys)

- Recommendation has been widely applied in online services:
  - E-commerce, Content Sharing, Social Networking ...

amazon ebay 淘宝网 Taobao.com JD.COM

Product Recommendation

Frequently bought together

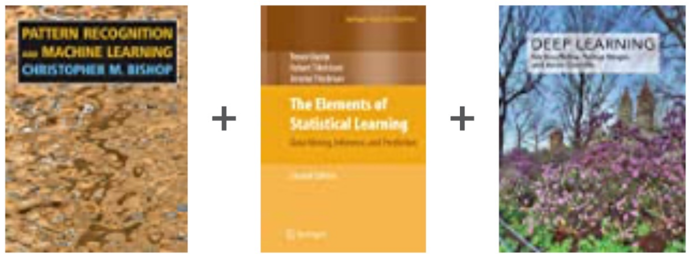

A B C

Total price: $208.9

Add all three to Cart

Add all three to List

Amazon's recommendation algorithm drives **35%** of its sales [from McKinsey, 2012]

Recommender Systems (RecSys)

- Recommendation has been widely applied in online services:
  - E-commerce, Content Sharing, Social Networking ...

News/Video/Image Recommendation

TikTok's recommendation algorithm

Top 10 Global Breakthrough Technologies in 2021

MIT Technology Review

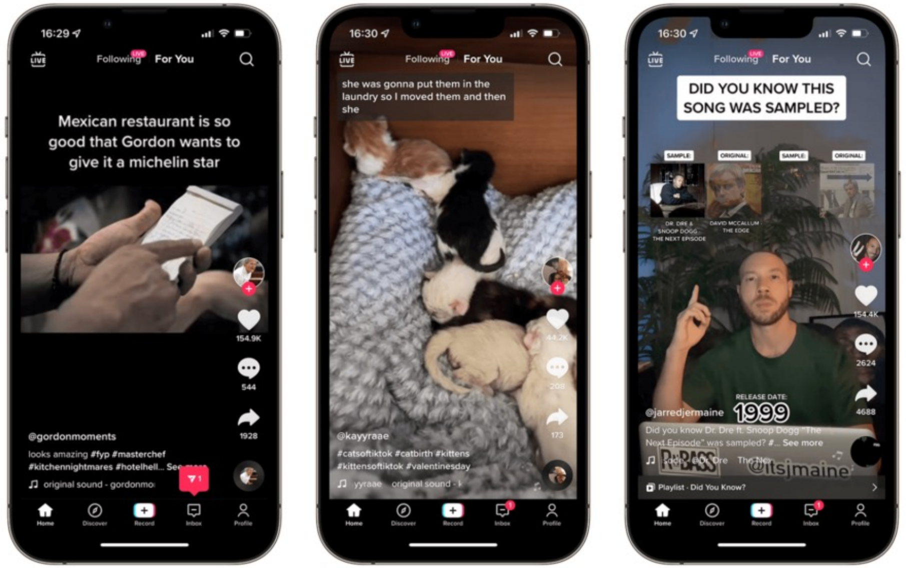

Large Language Models (LLMs) are Changing Our Lives

OpenAI | Meta | Google AI T5 | ......

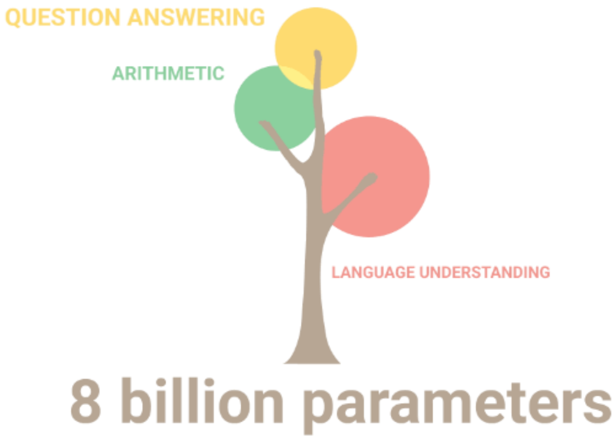

https://github.com/Hannibal046/Awesome-LLM/tree/main

LLMs in Natural Language Processing

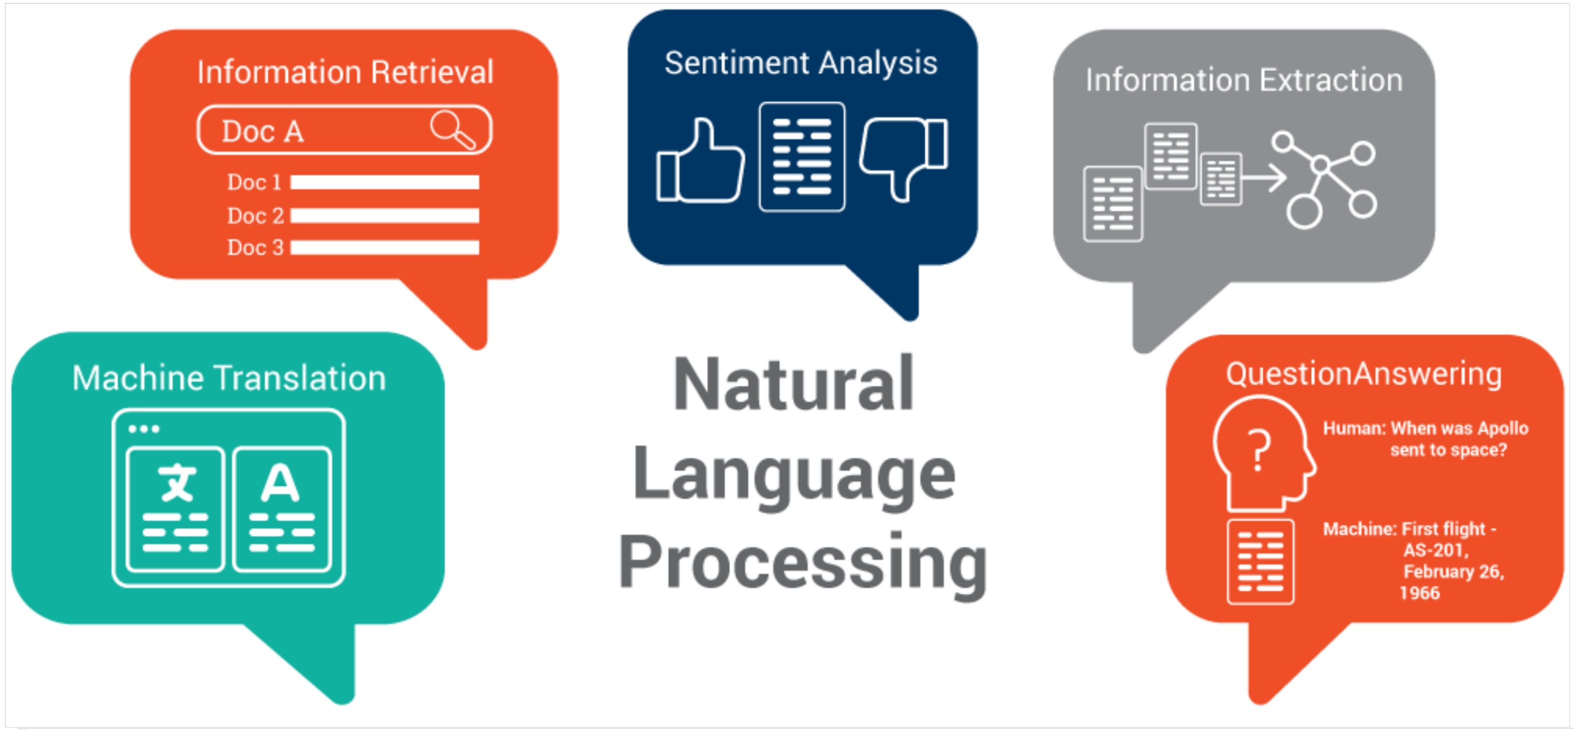

Large Language Models (LLMs)

LLMs in Downstream Domains

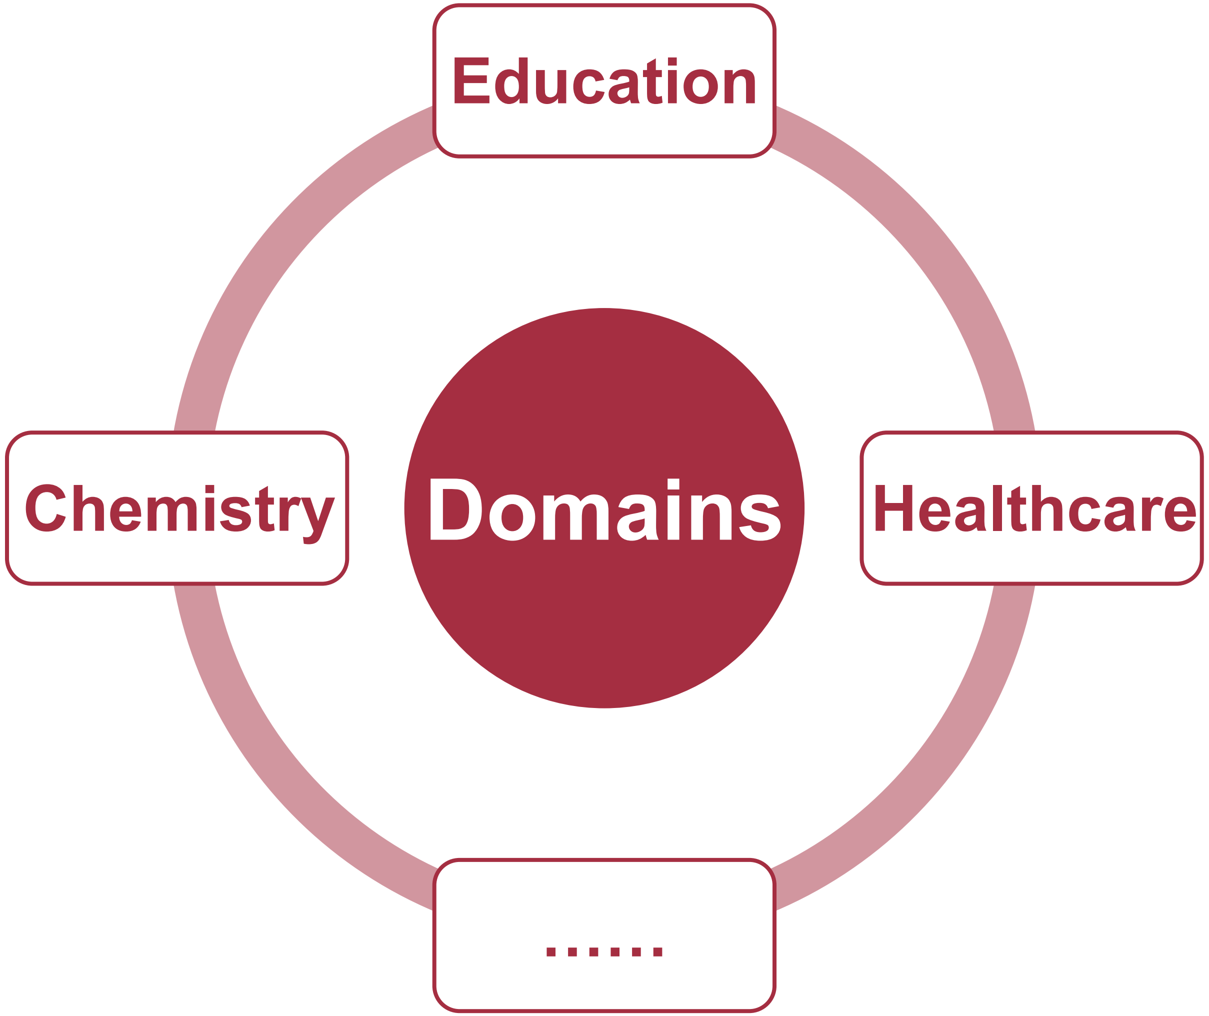

❑ **Molecule discovery**, etc.

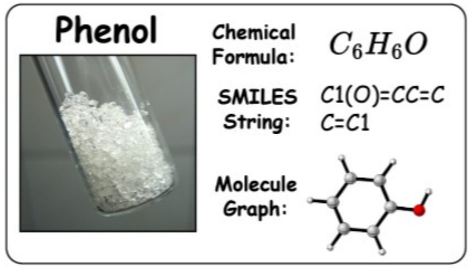

(a) Molecule Representations.

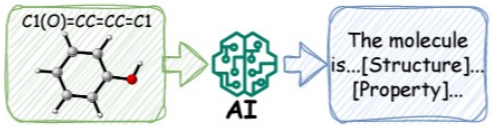

(b) Molecule Captioning.

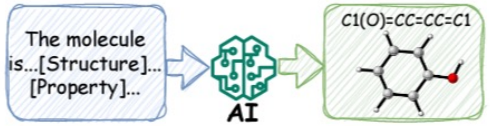

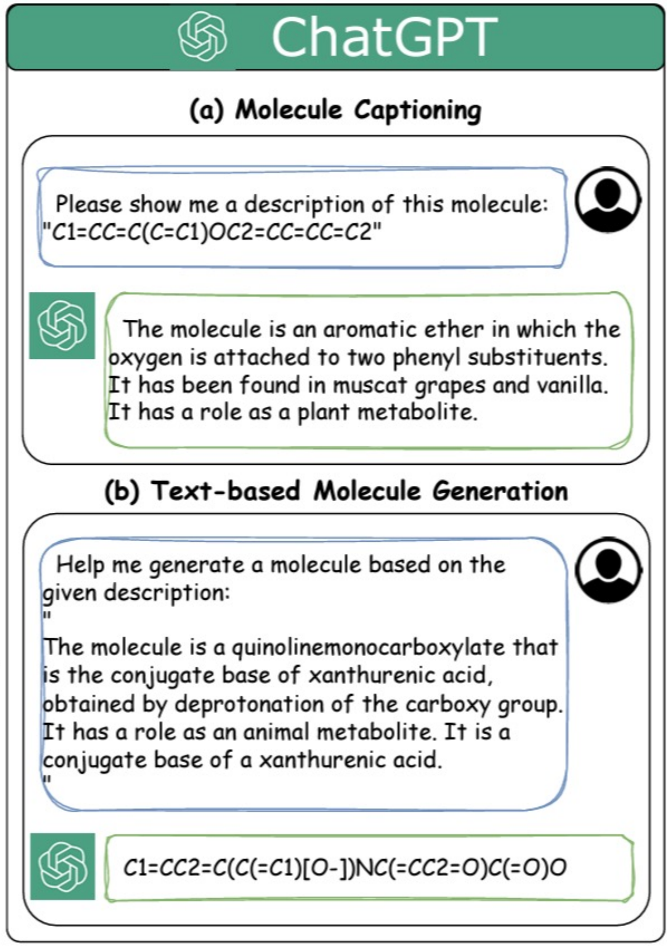

"Empowering Molecule Discovery for Molecule-Caption Translation with Large Language Models: A ChatGPT Perspective." arXiv preprint arXiv:2306.06615 (2023).

LLMs in Downstream Domains

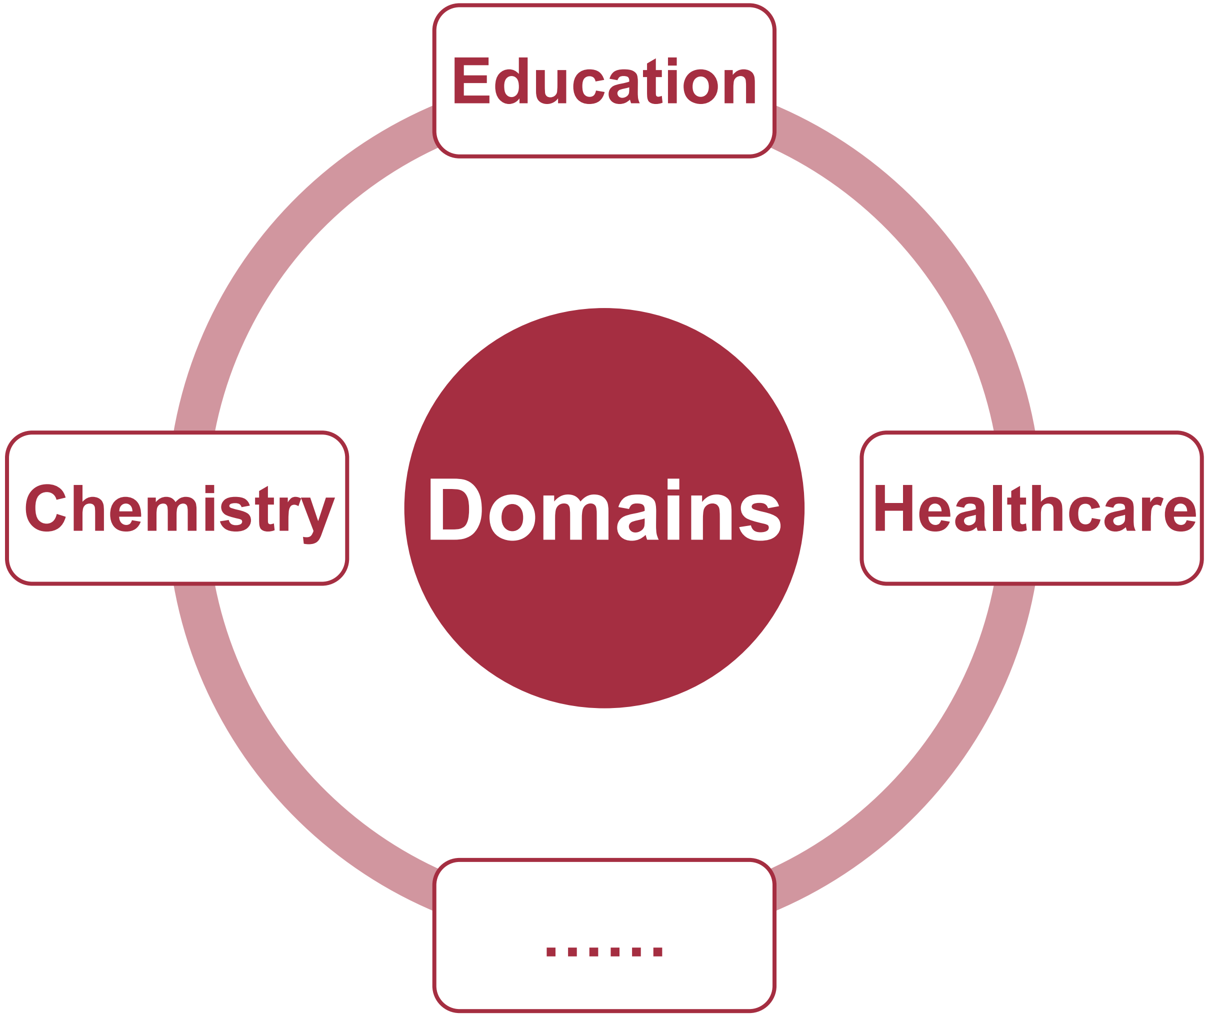

❑ **Curriculum & Teaching**, etc.

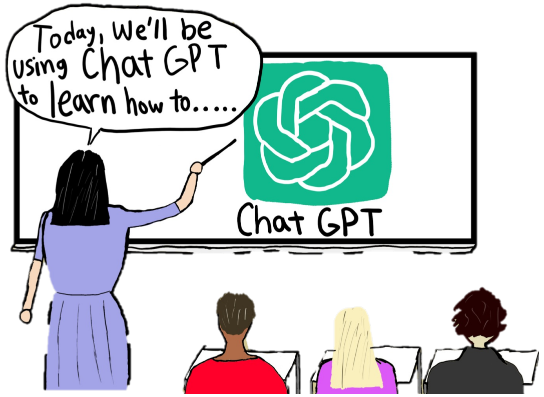

❑ **Medical consultation**, etc.

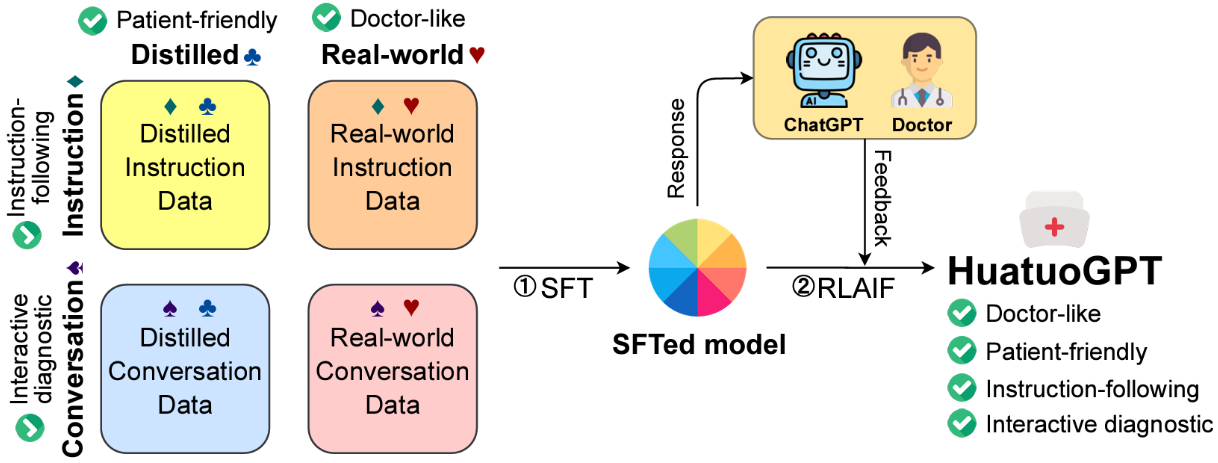

"HuatuoGPT, towards Taming Language Model to Be a Doctor." arXiv preprint arXiv:2305.15075 (2023).

LLMs in RecSys

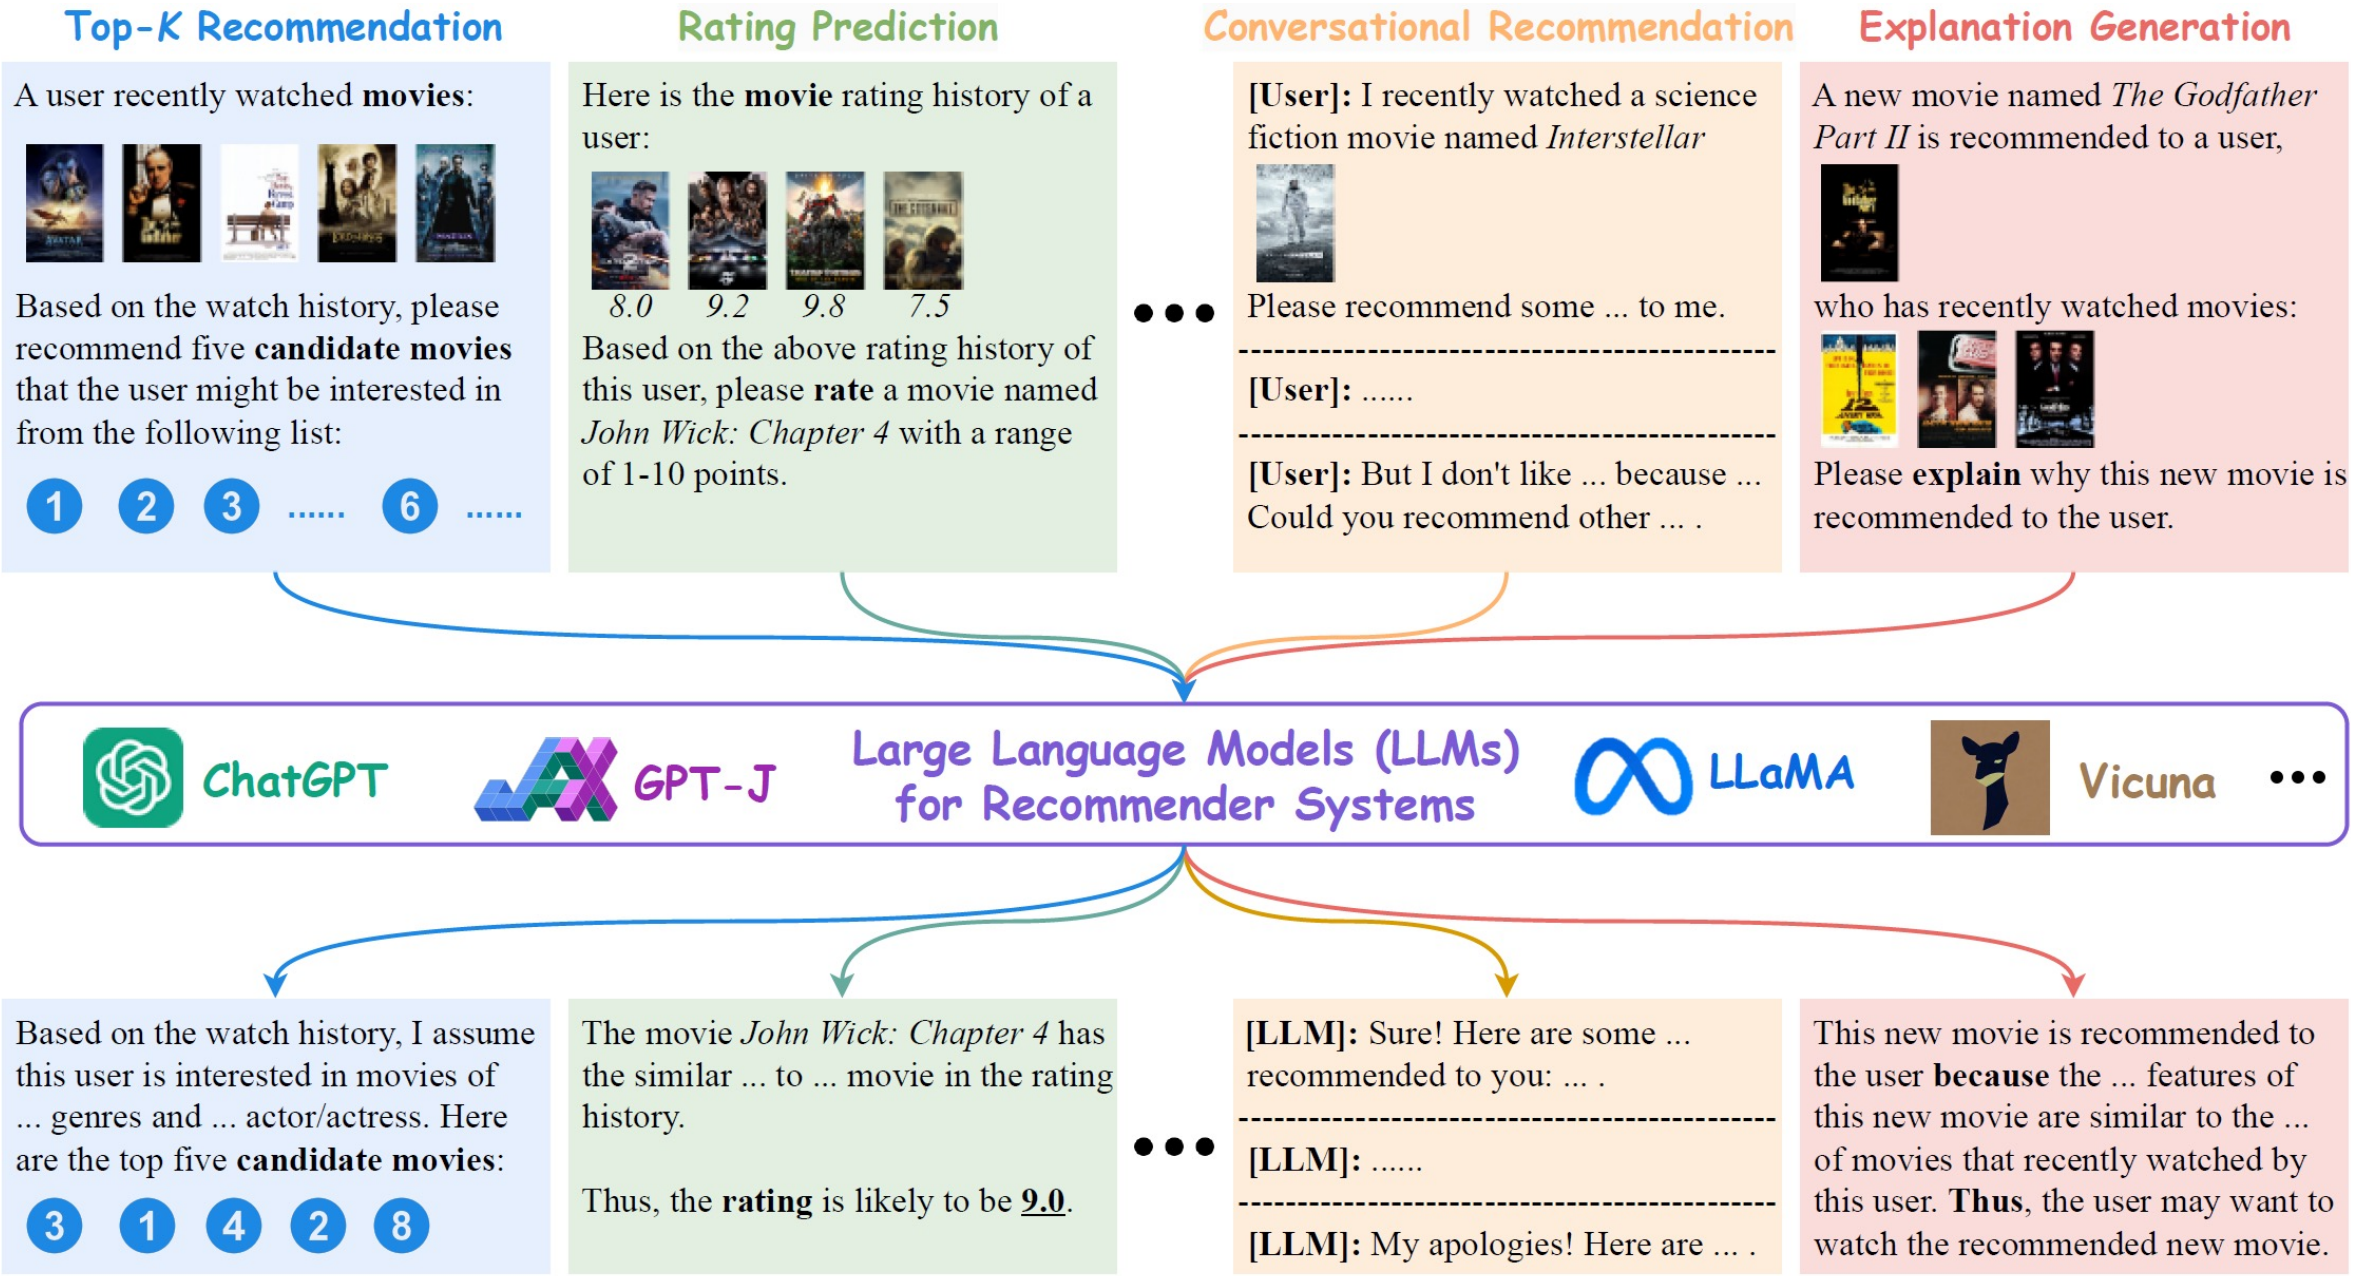

Potentials of LLMs in RecSys

As the **parameter size** of LLMs continues to **scale up** with a larger **training corpus** …

❑ Language understanding and generation ability

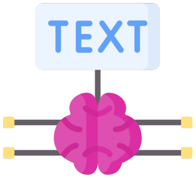

❖ LLMs can comprehend **human intentions** and generate **language responses** that are more human-like in nature.

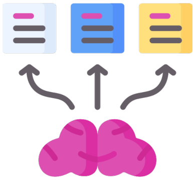

❑ Generalization capability

❖ LLMs can apply their learned knowledge to **fit various downstream tasks**, even **without being fine-tuned** on specific tasks.

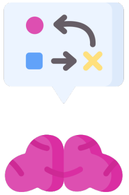

❑ Reasoning capability

❖ LLMs can generate the outputs with **step-by-step reasonings** to support complex **decision-making processes**.

Language Understanding & Generation

- Sufficiently capture **textual knowledge** about users and items
  - Rich **textual side information** about users and items in RecSys
  - Diverse **open-world knowledge** encoded in LLMs

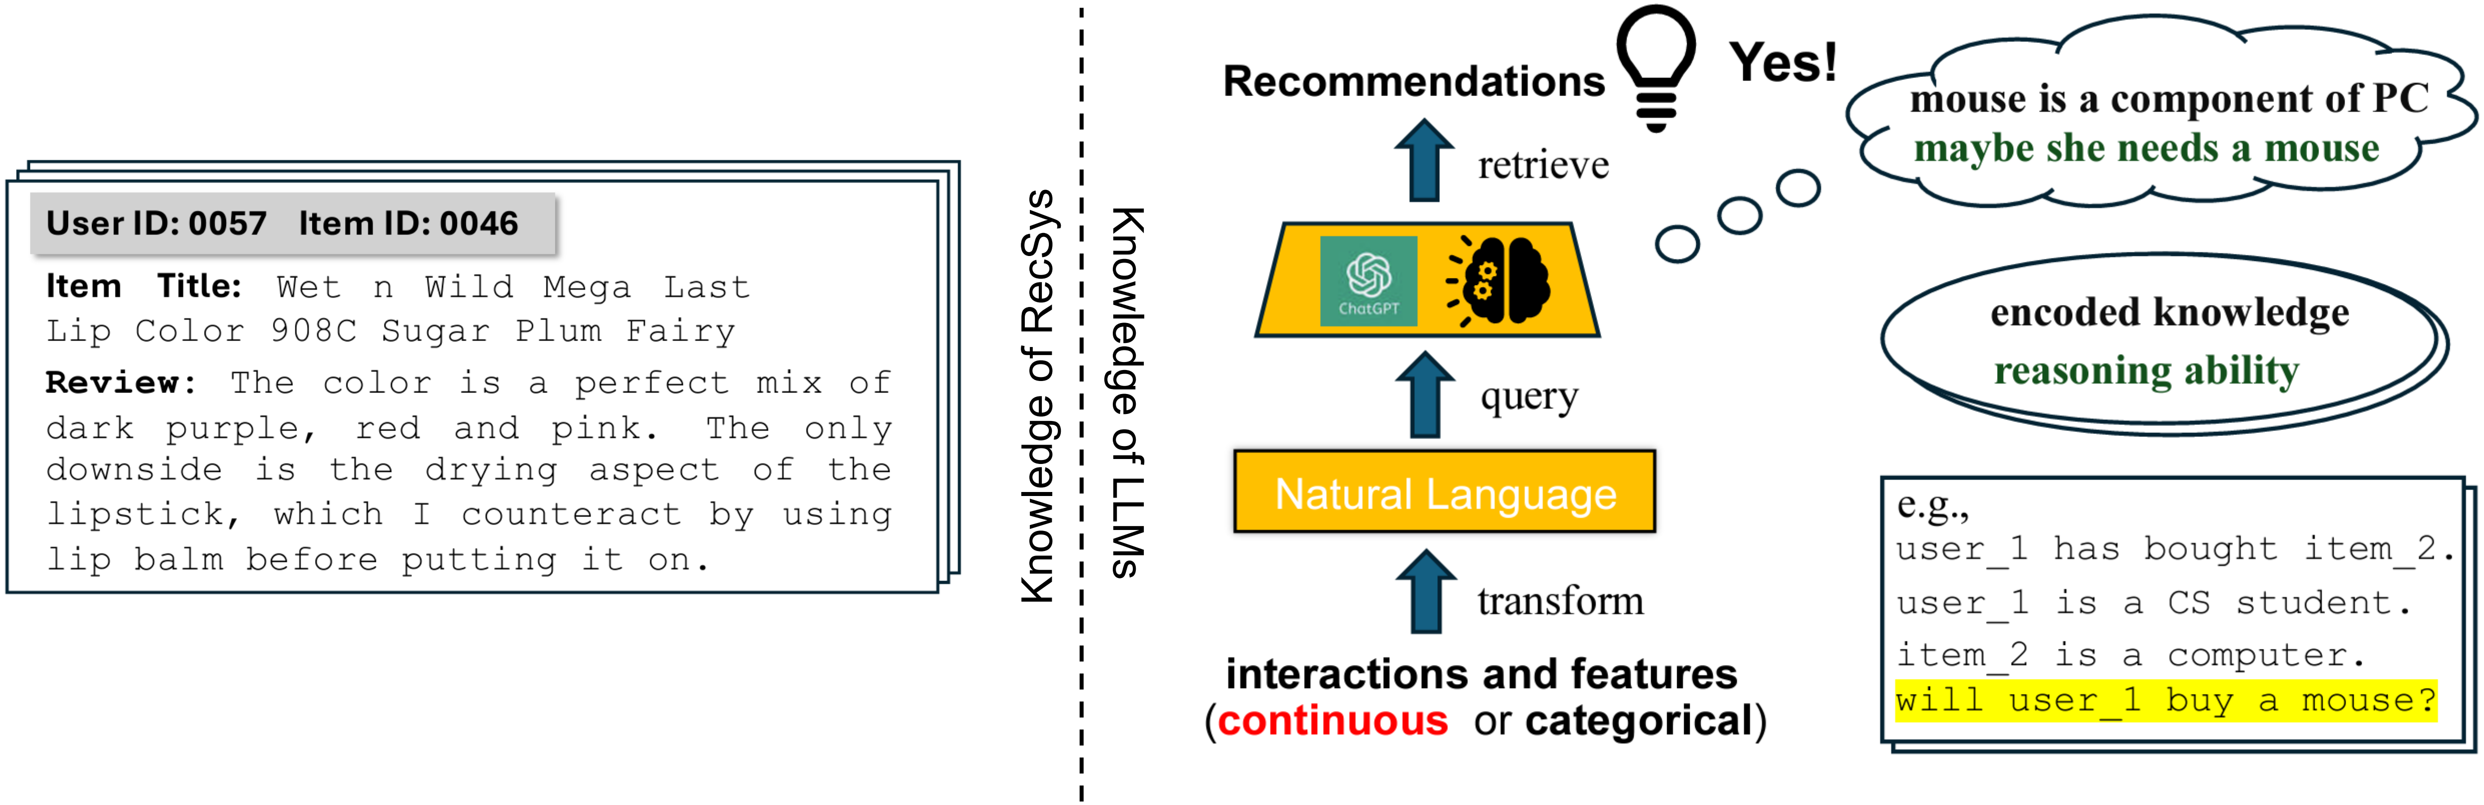

"Collaborative Large Language Model for Recommender Systems." arXiv preprint arXiv:2311.01343 (2023).

Generalization

- ❑ Adapt to **various recommendation tasks** even without being fine-tuned
  - ❖ LLMs can apply their **learned knowledge** to address recommendation objectives
  - ❖ **Multi-task adaption** by providing appropriate task instructions or a few task demonstrations

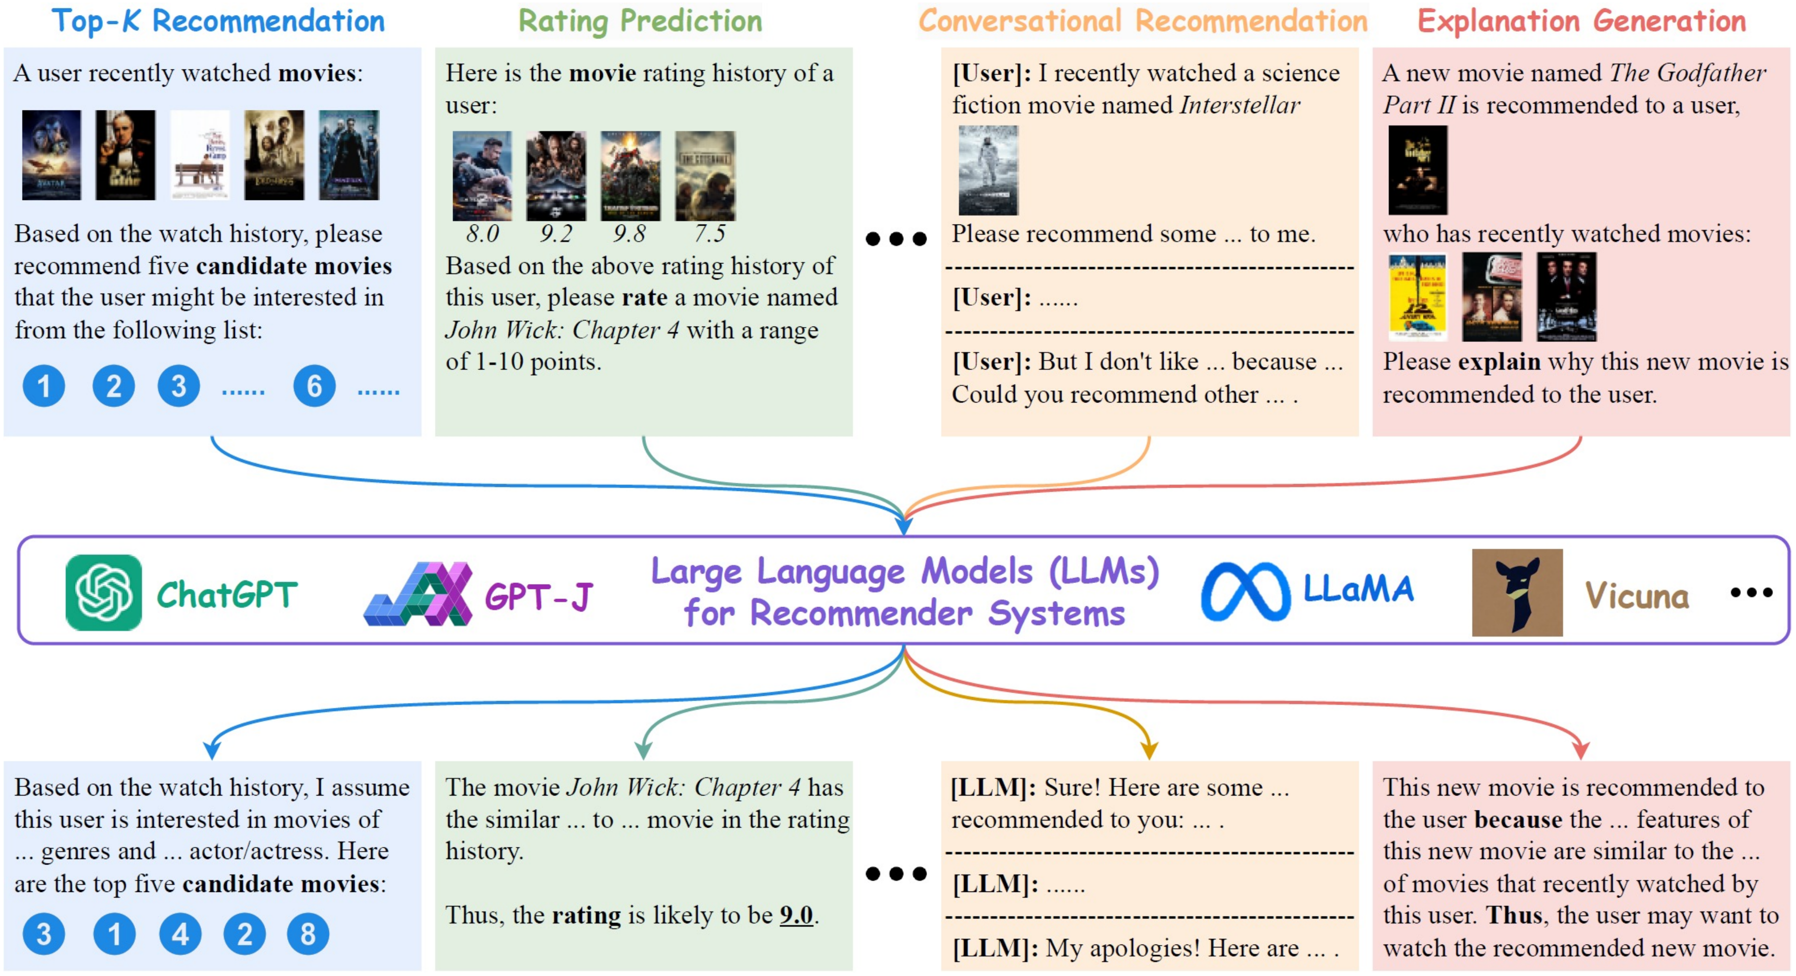

Reasoning

- Support complex **decision-making processes** in RecSys
  - Retrieve information from **large contexts** and control **multi-step** recommendation tasks
  - Generate outputs with **step-by-step reasoning** empowered by chain-of-thought prompting

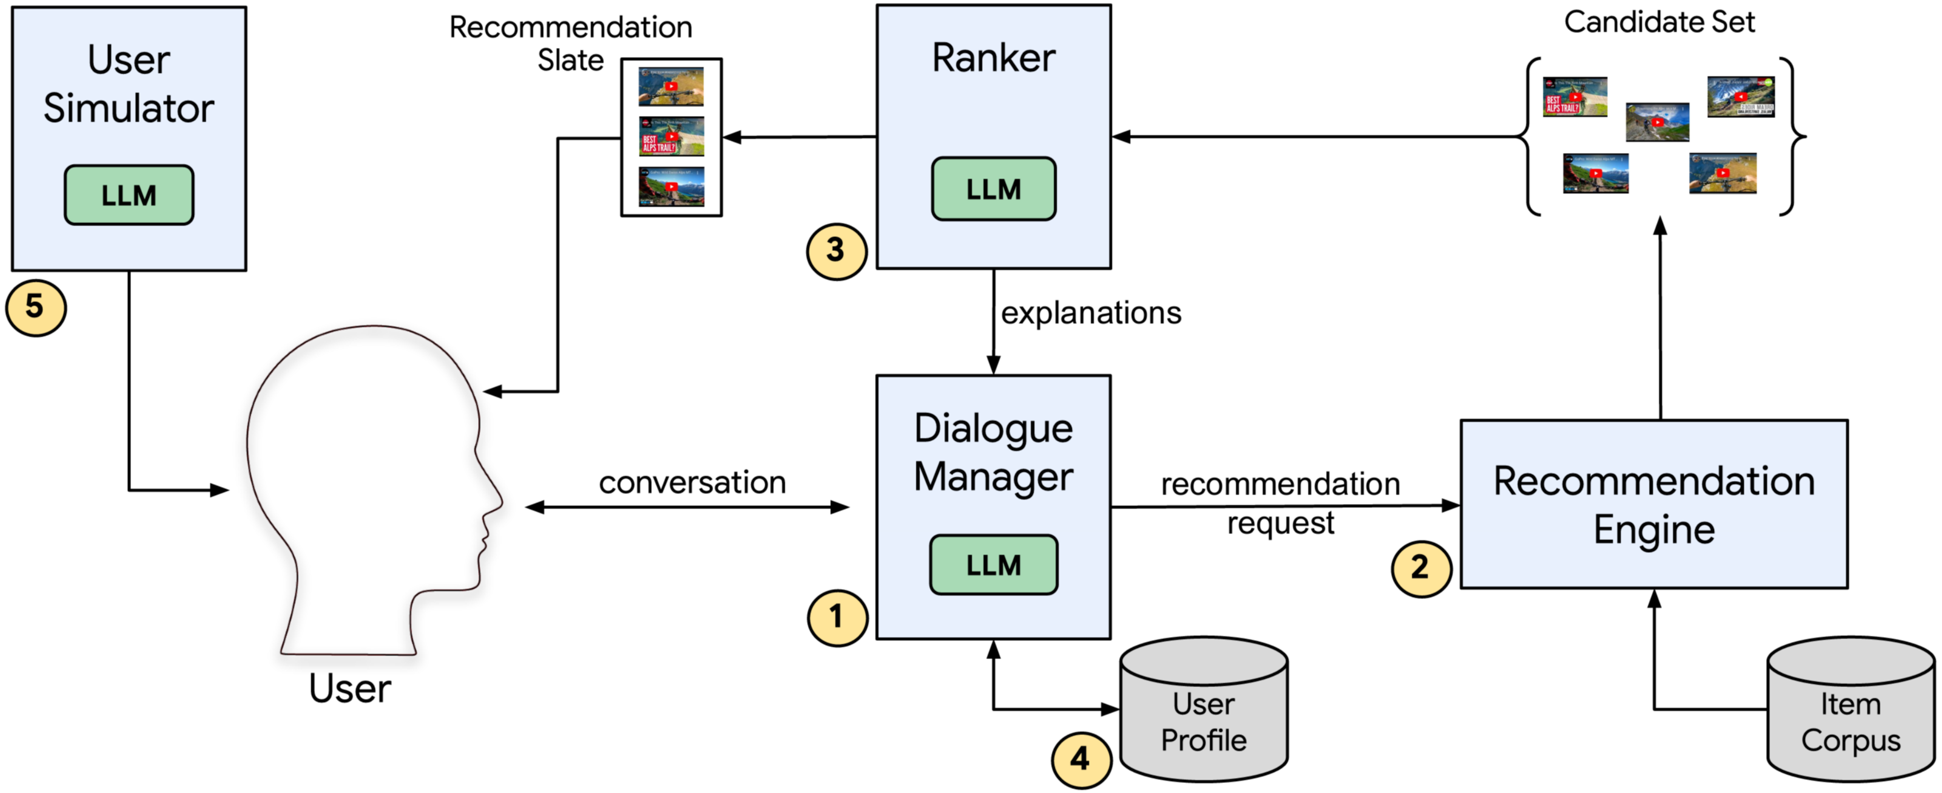

"Leveraging Large Language Models in Conversational Recommender Systems." arXiv preprint arXiv:2305.07961 (2023).

A Comprehensive Survey Paper

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

Wenqi Fan, Zihuai Zhao, 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

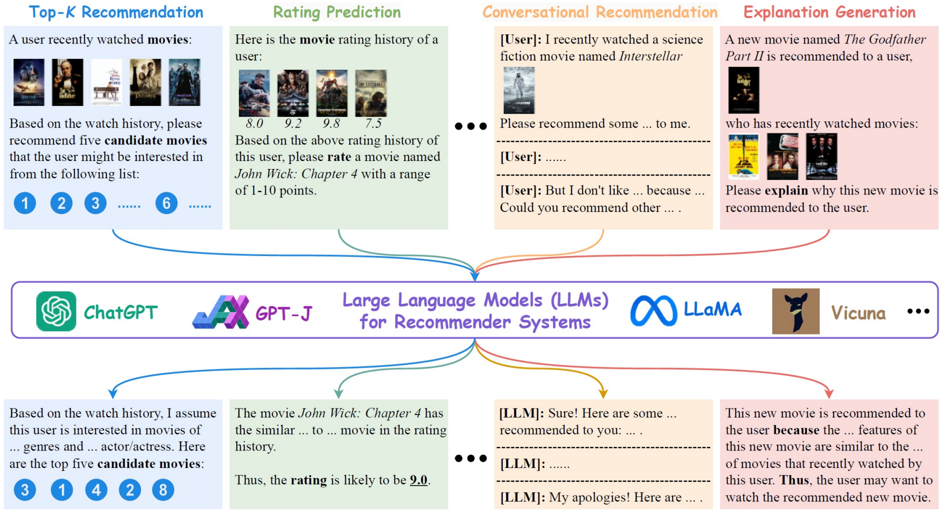

ICDM'2023 Tutorial Website (Slides)

Zoom ID: 91649466943
Password: 202312

Tutorial website: https://advanced-recommender-systems.github.io/llms_rec_tutorial/

Recruitment

- Our research group are actively recruiting self-motivated **Postdoc, Ph.D. students, and Research Assistants**, etc. **Visiting scholars, interns, and self-funded students** are also welcome. Send me an email if you are interested.

  - Research areas: machine learning (ML), data mining (DM), artificial intelligence (AI), deep learning (DNNs), graph neural networks (GNNs), computer vision (CV), natural language processing (NLP), etc.

  - Position Details: **https://wenqifan03.github.io/openings.html**

Tutorial Outline

Introduction | **Preliminaries** | Pre-training | Fine-tuning | Prompting | Future Directions

- **Introduction** of RecSys in the era of LLMs (Dr. Wenqi Fan)
- **Preliminaries of RecSys and LLMs (Yunqing Liu)**
- **Pre-training** paradigms for adopting LLMs to RecSys (Jiatong Li)
- **Fine-tuning** paradigms for adopting LLMs to RecSys (Jiatong Li)
- **Prompting** paradigms for adopting LLMs to RecSys (Zihuai Zhao)
- **Future directions** of LLM-empowered RecSys (Dr. Yiqi Wang)

Zoom ID:
91649466943
Password:
202312

Website QR Code

Overview

Presenter:
Yunqing Liu

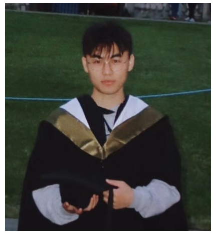

RecSys

- Collaborative Filtering (CF)
- Content-based Recommendation
- Deep Recommender Systems

LLMs

- Encoder-Only
- Decoder-Only
- Encoder-Decoder

Zoom ID:
91649466943
Password:
202312

Website QR Code

Recommender Systems

INPUT

Historical user-item interactions or additional side information (e.g., social relations, item's knowledge, etc.)

OUTPUT

Predict how likely a user would interact with a target Item (e.g., click, view, or purchase)

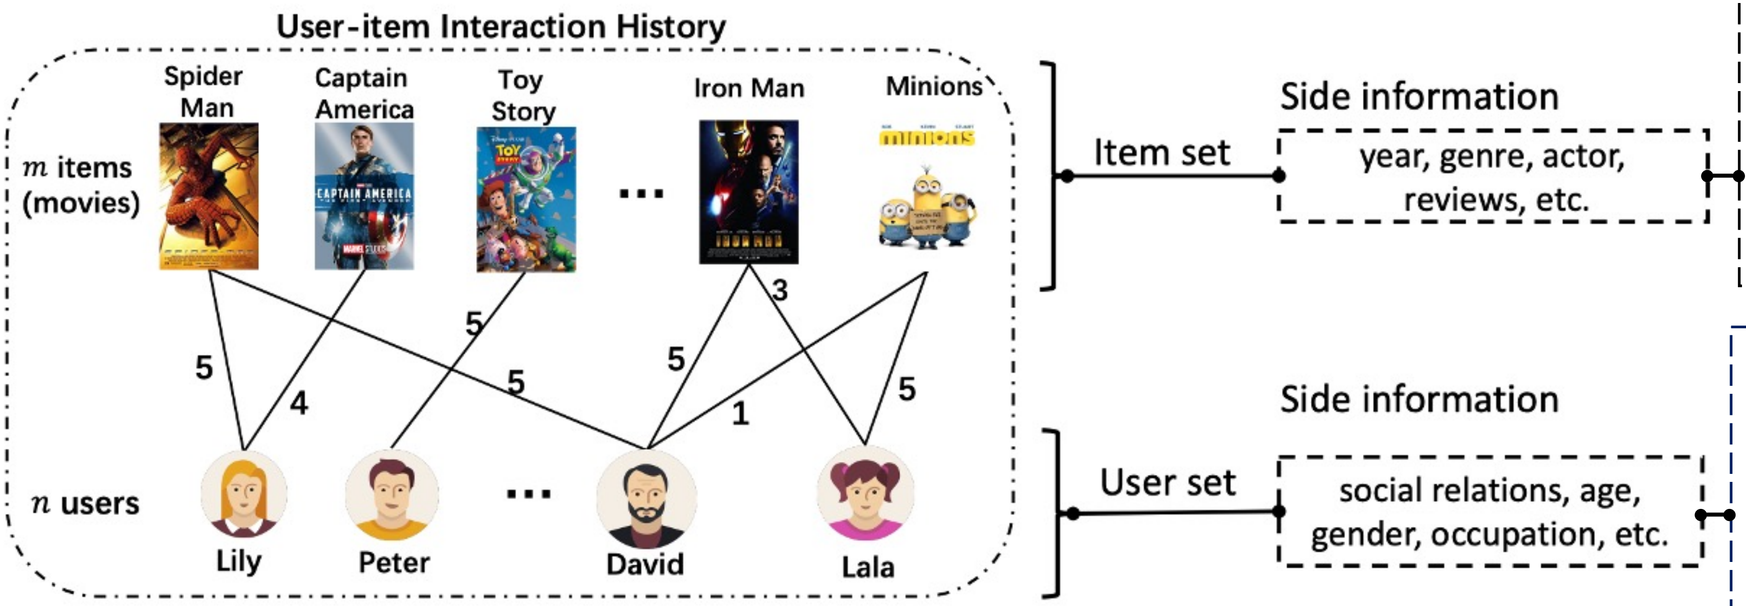

Title: Spider Man (2002)
Genre: Action · Adventure · Sci-Fi
Actor: Tobey Maguire, ……
Reviews: 1. Considered as one of the most successful superhero movies ever made ……
……

Name: Peter
Social Relations: David (Friend),…
Age: 18
Gender: Male
Occupation: Student
……

Collaborative Filtering (CF)

- The most well-known technique for recommendation
  - Similar users (with respect to their historical interactions) have similar preferences.
  - Modelling user's preferences on items based on their past interactions (e.g., ratings and clicks).
- Learning representations of users and items is the key to CF

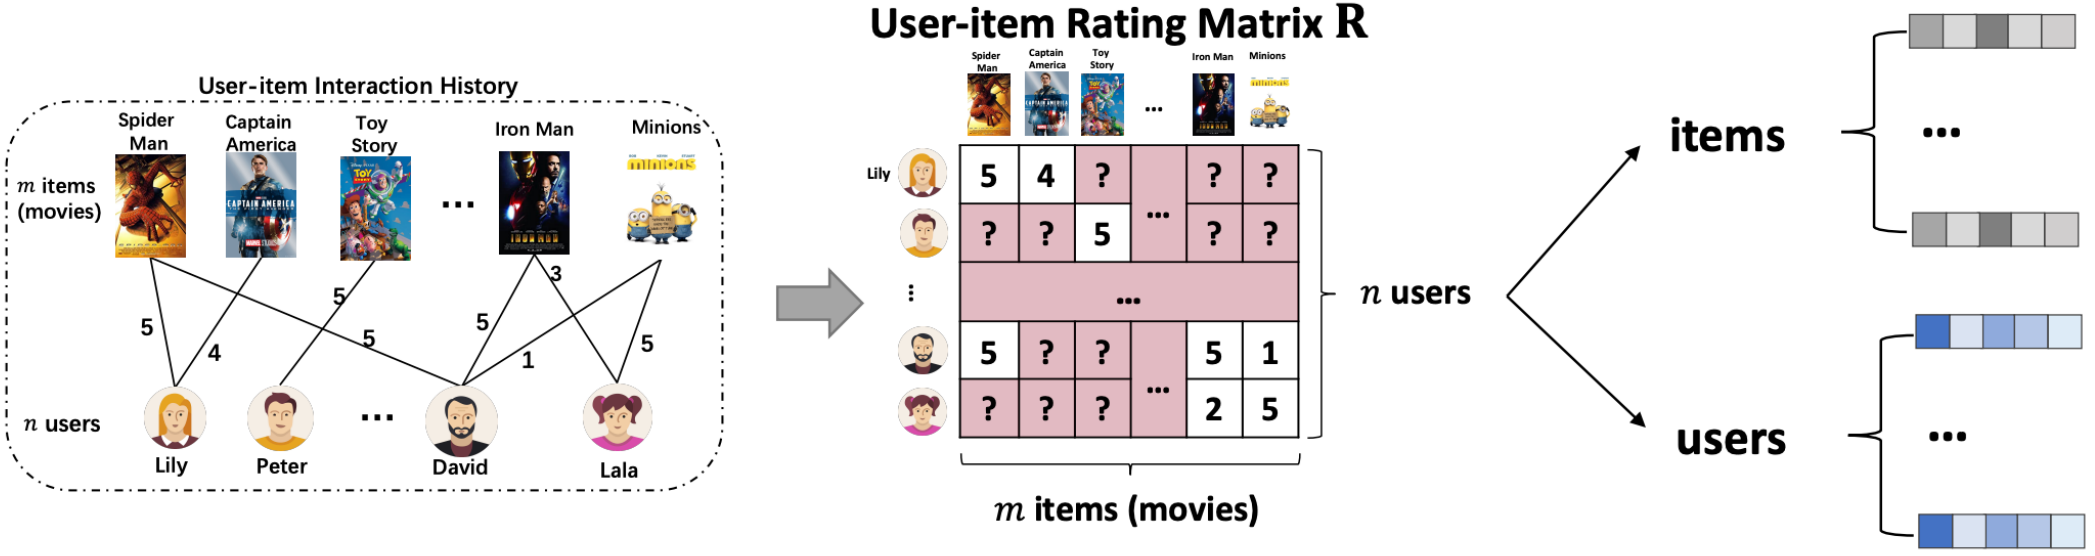

Task: predicting missing movie ratings in Netflix.

Content-based Recommendation

- Taking advantage of additional knowledge about users or items
- Enhancing user and item representations for improving recommendation performance

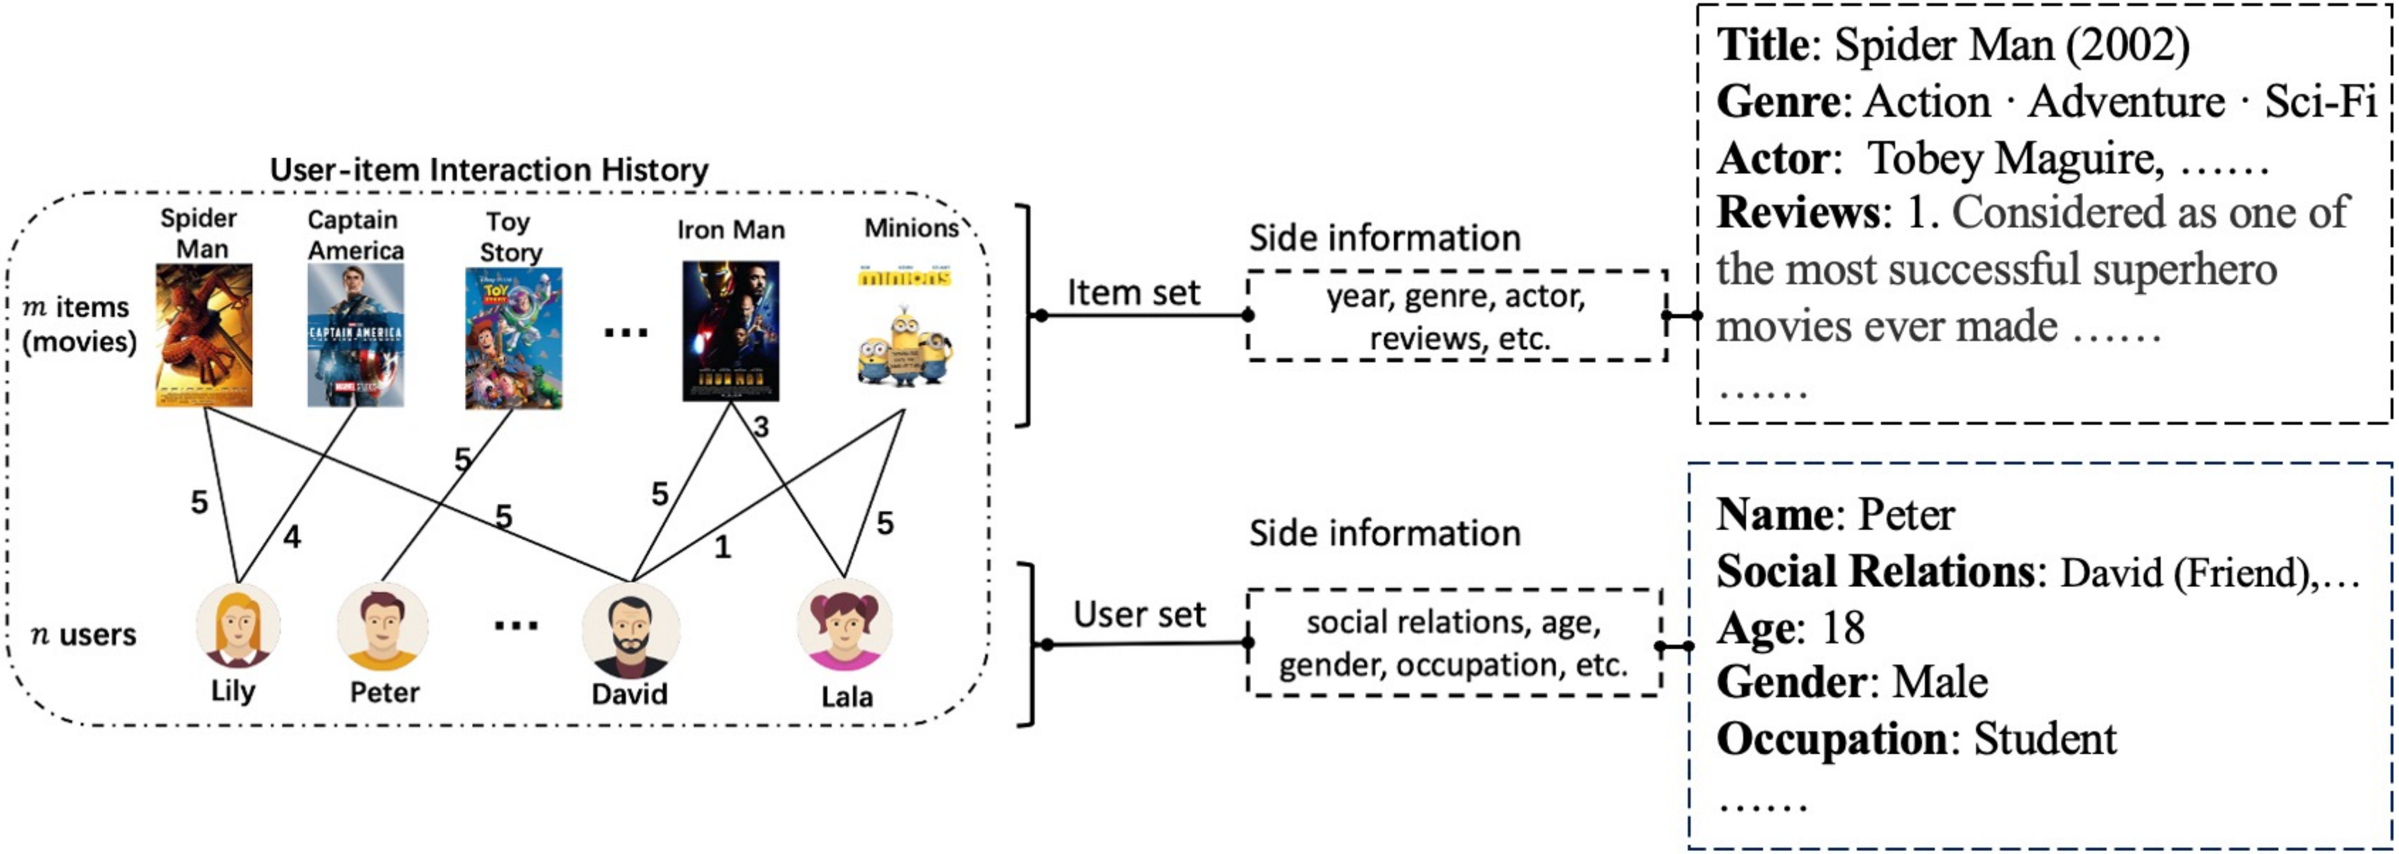

Deep Recommender Systems

- Deep learning techniques have been effectively applied to develop recommender systems
- Remarkable representation learning capabilities

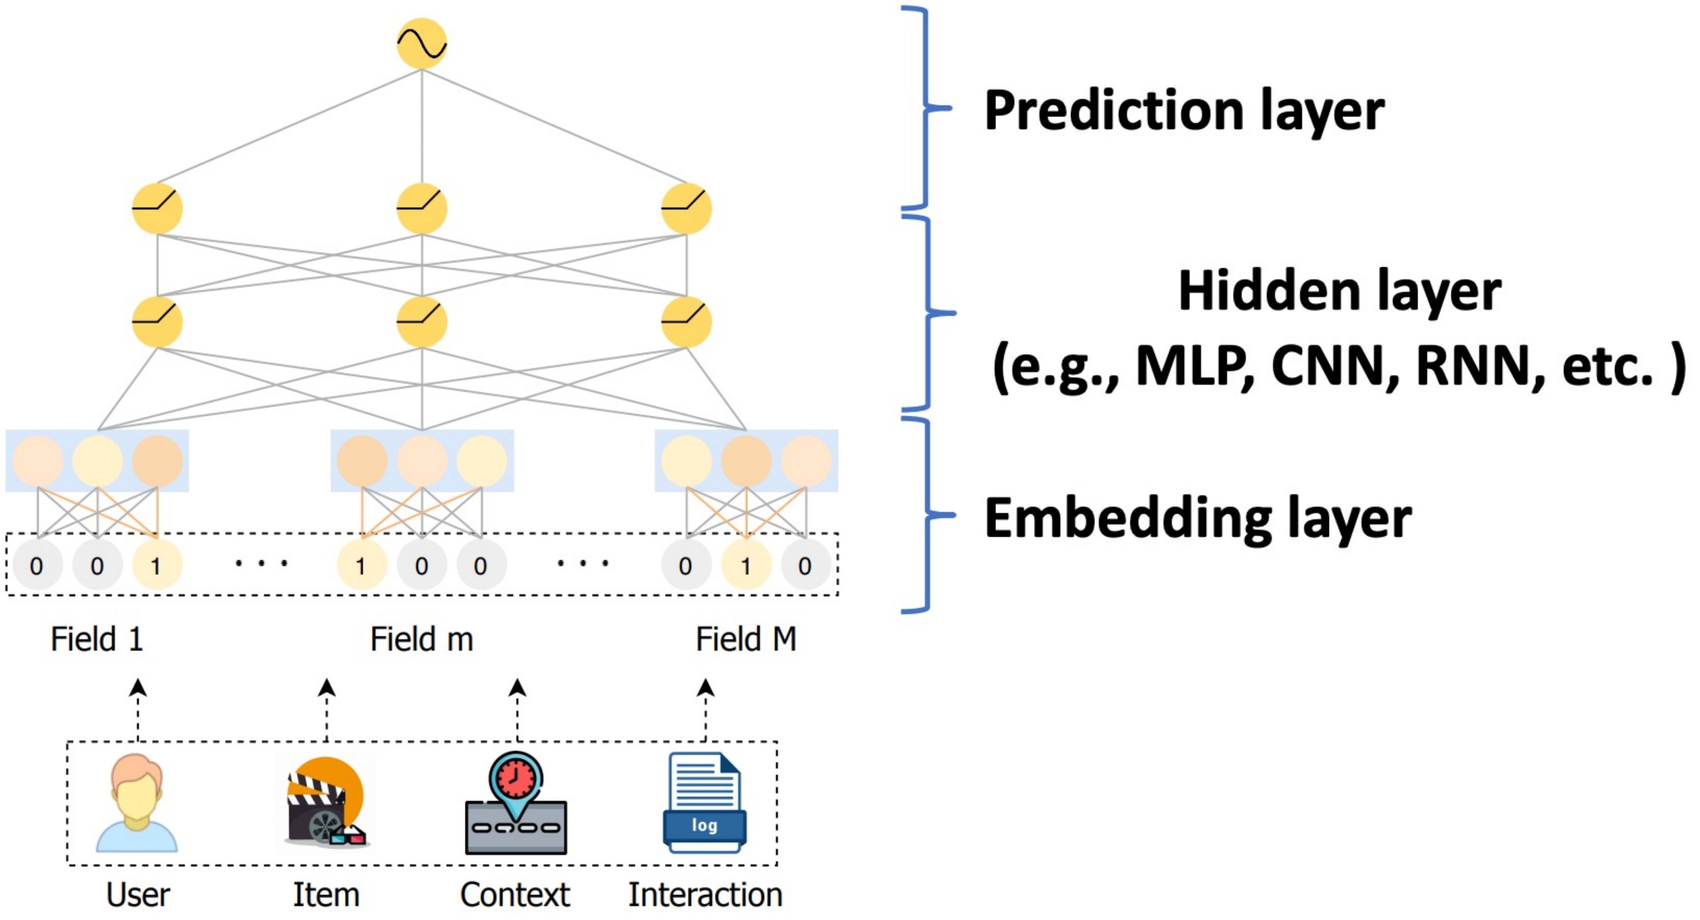

NeuMF

- **NeuMF** unifies the strengths of MF and MLP in modelling user-item interactions
  - **MF** uses an inner product as the interaction function
  - **MLP** is more sufficient to capture the complex structure of user interaction data

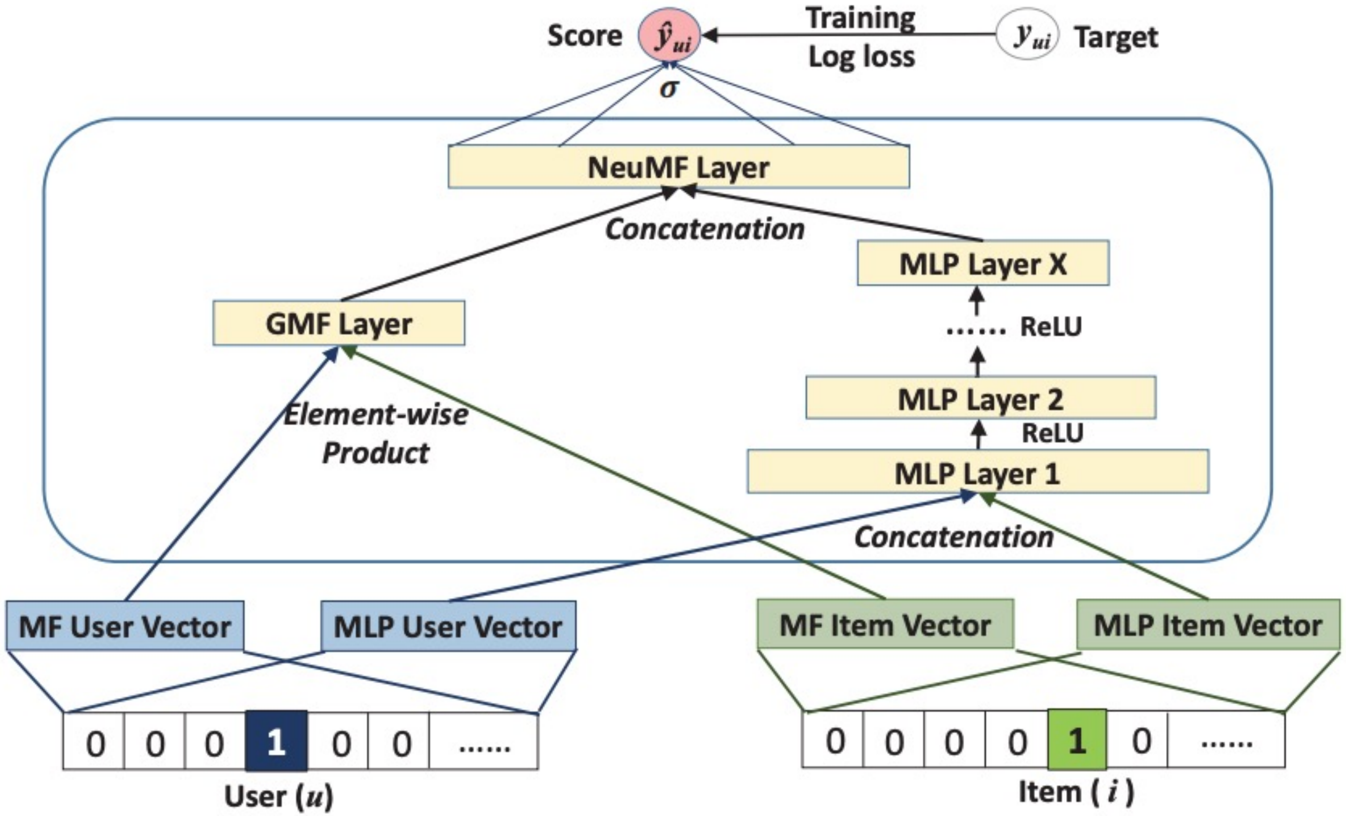

"Neural collaborative filtering." Proceedings of the 26th international conference on world wide web (2017).

Caser

- ❑ **Top-N sequential recommendation** models each user as a **sequence of items** interacted in the past and aims to **predict top-N ranked items**
- ❑ Convolutional Sequence Embedding Recommendation Model

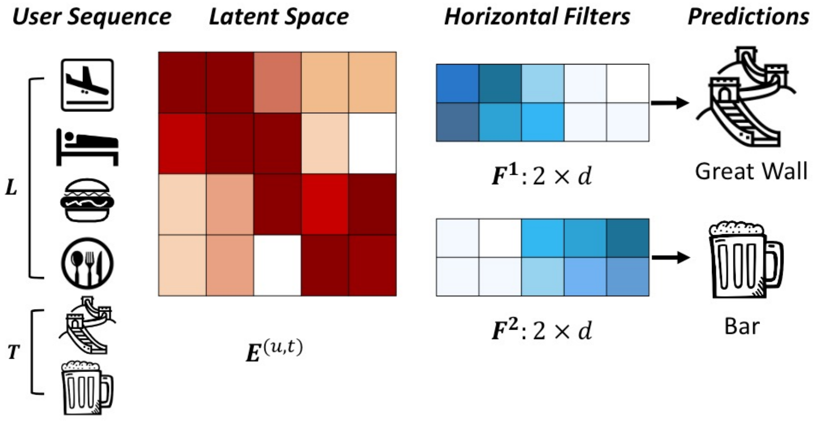

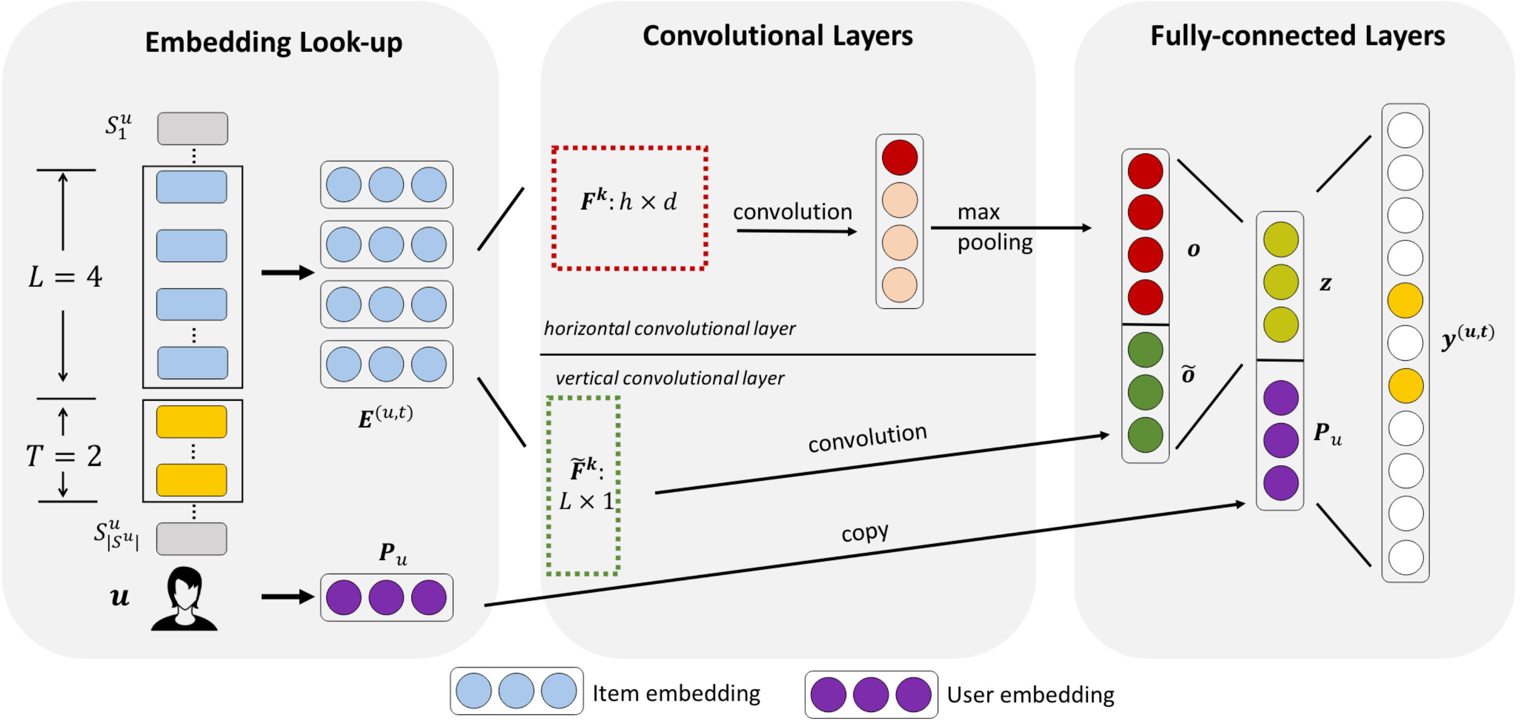

"Personalized top-n sequential recommendation via convolutional sequence embedding." WSDM (2018).

GRU4Rec

- ❑ **Session-based** Recommendations with Recurrent Neural Networks (**RNN**)
- ❑ Introducing **session-parallel mini-batches**, mini-batch based output sampling and ranking loss function

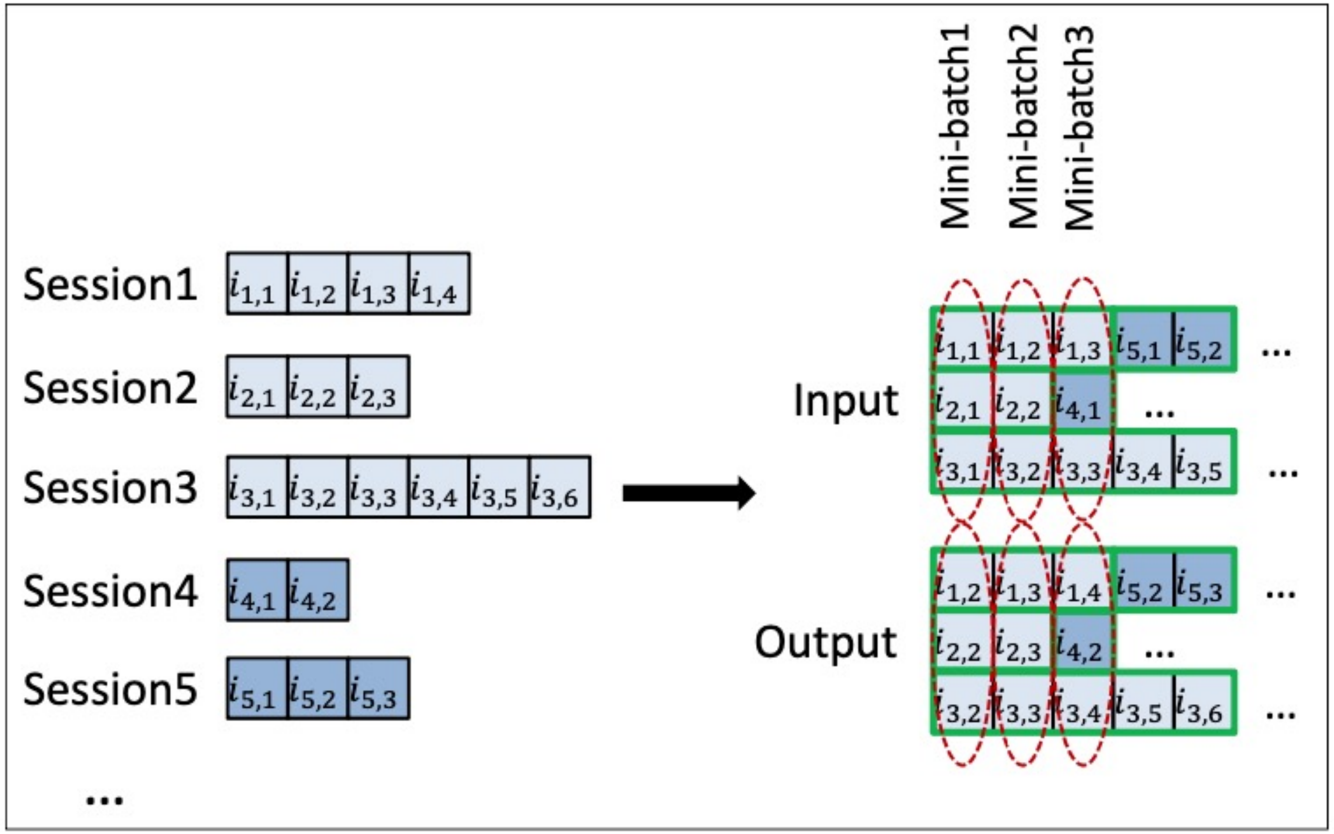

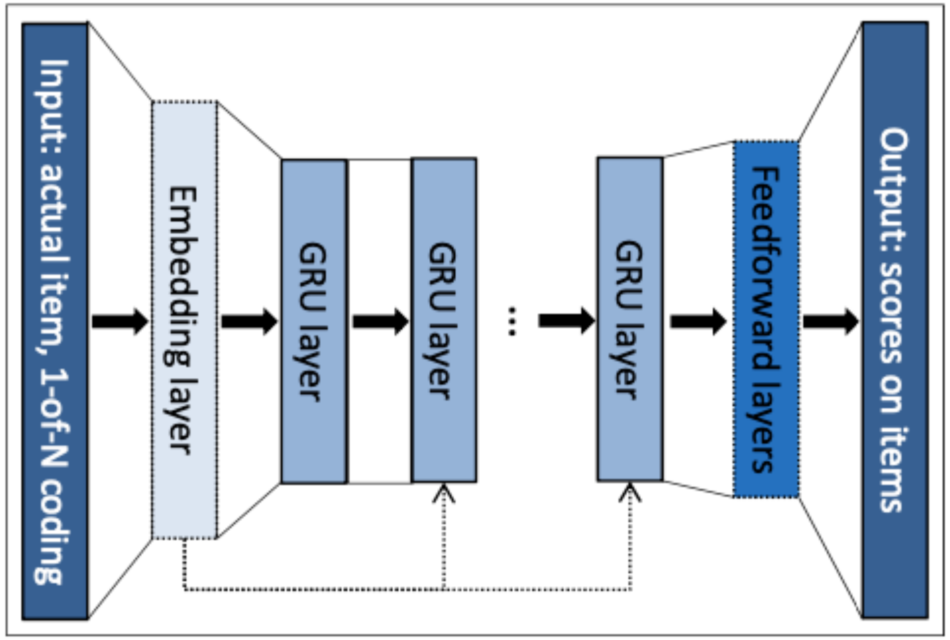

"Session-based recommendations with recurrent neural networks." ICLR (2016).

SASRec

- Self-Attentive Sequential Recommendation
- Using an **attention** mechanism to capture **long-term semantics** and makes its **predictions** based on relatively **few actions**

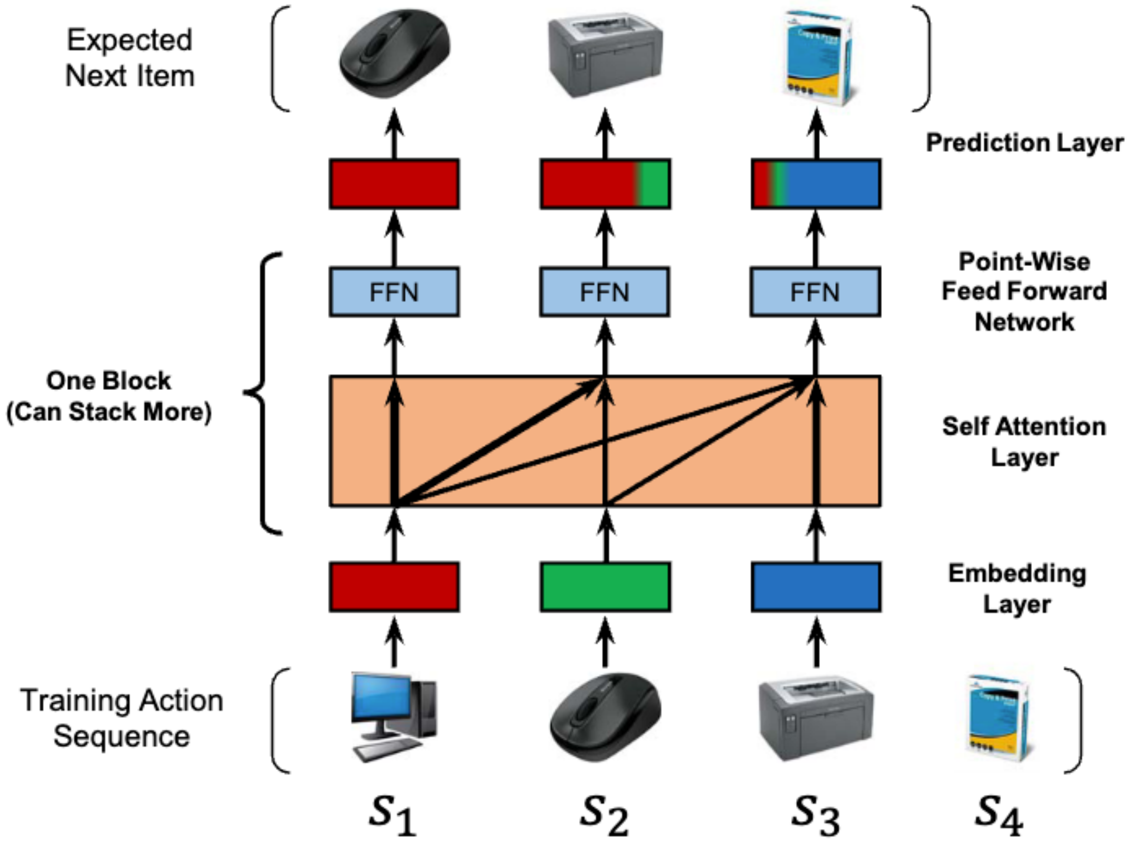

"Self-attentive sequential recommendation." ICDM (2018).

S3-Rec

- ❑ Utilizing the intrinsic data correlation to derive **self-supervision** signals
- ❑ Enhancing the data representations via **pre-training** methods

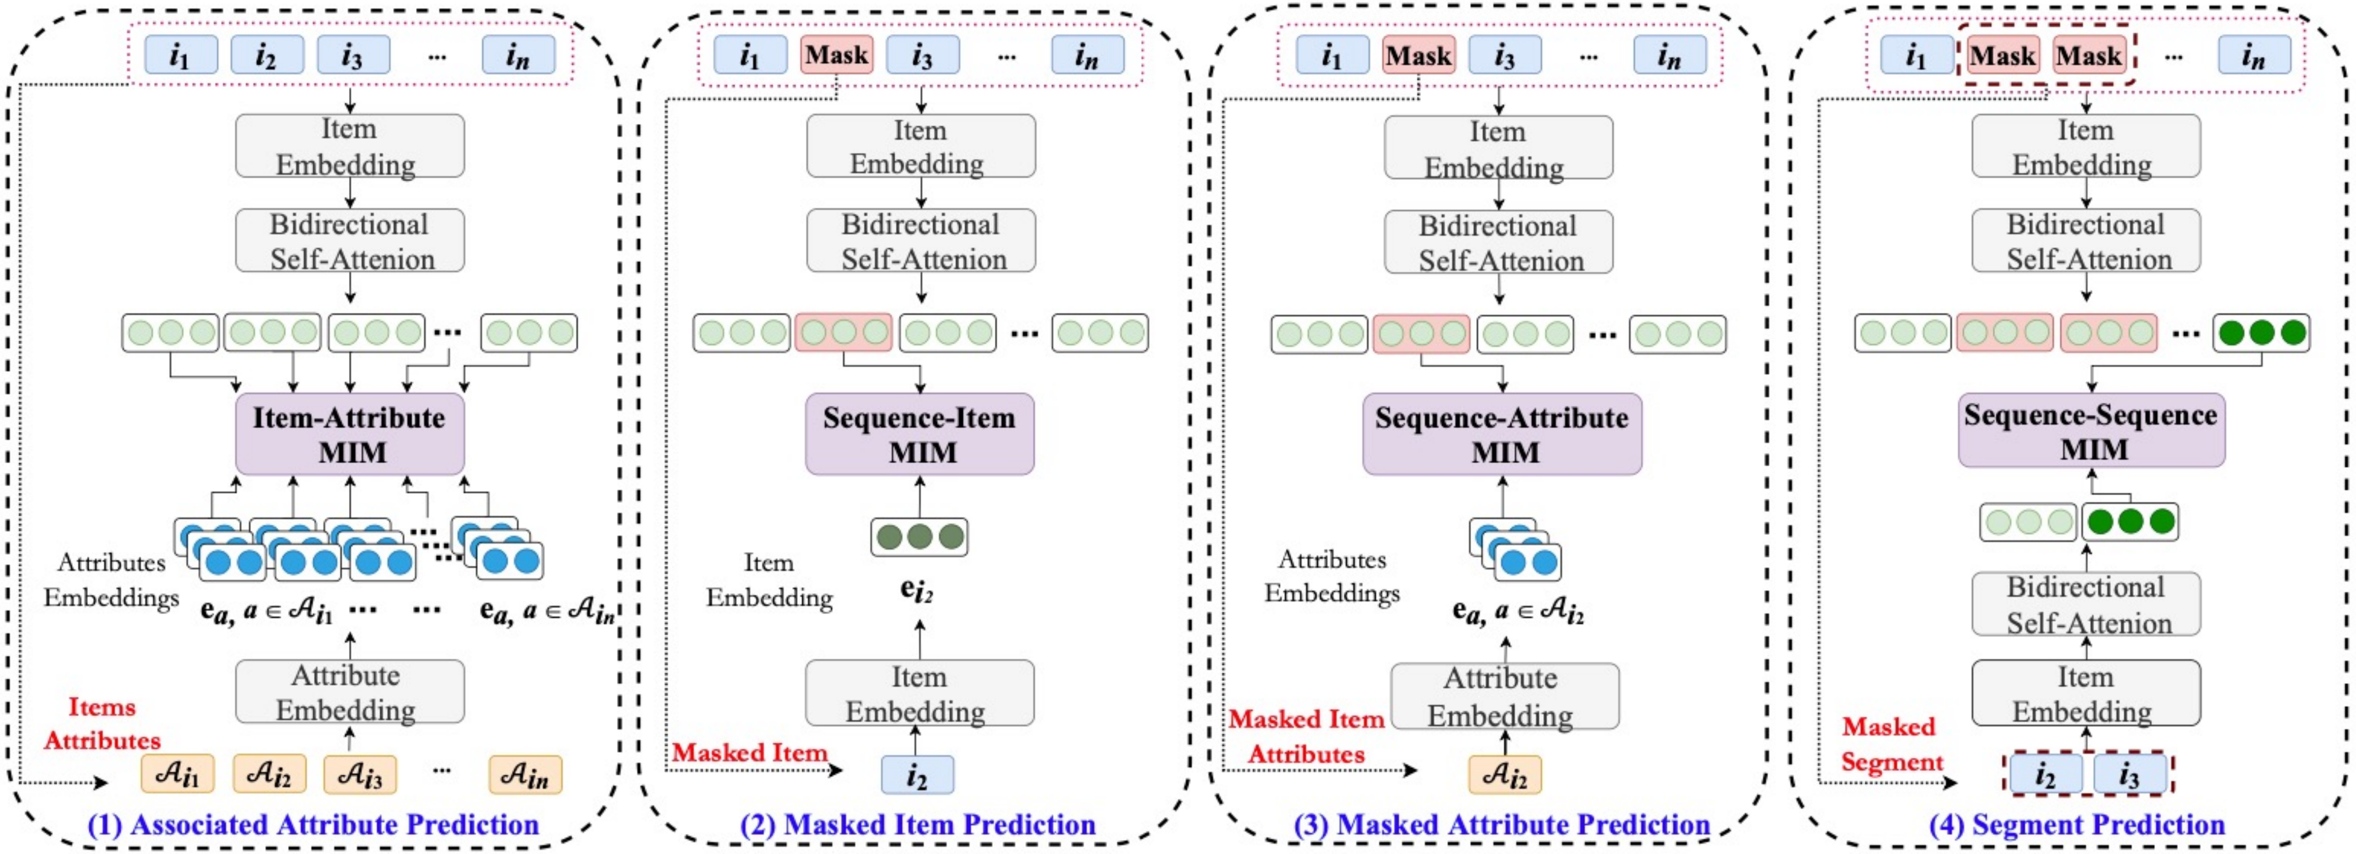

"S3-rec: Self-supervised learning for sequential recommendation with mutual information maximization." CIKM (2020).

GraphRec

- Data in social recommender systems can be represented as **user-user** social graph and **user-item** graph

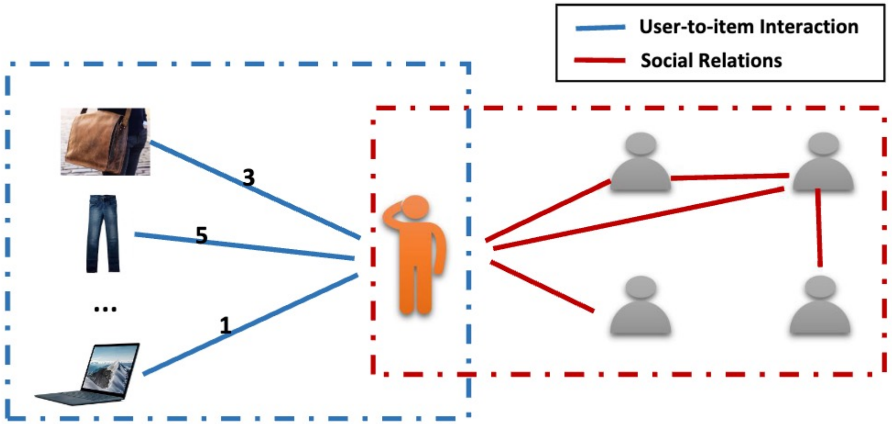

user-user graph

user-item graph

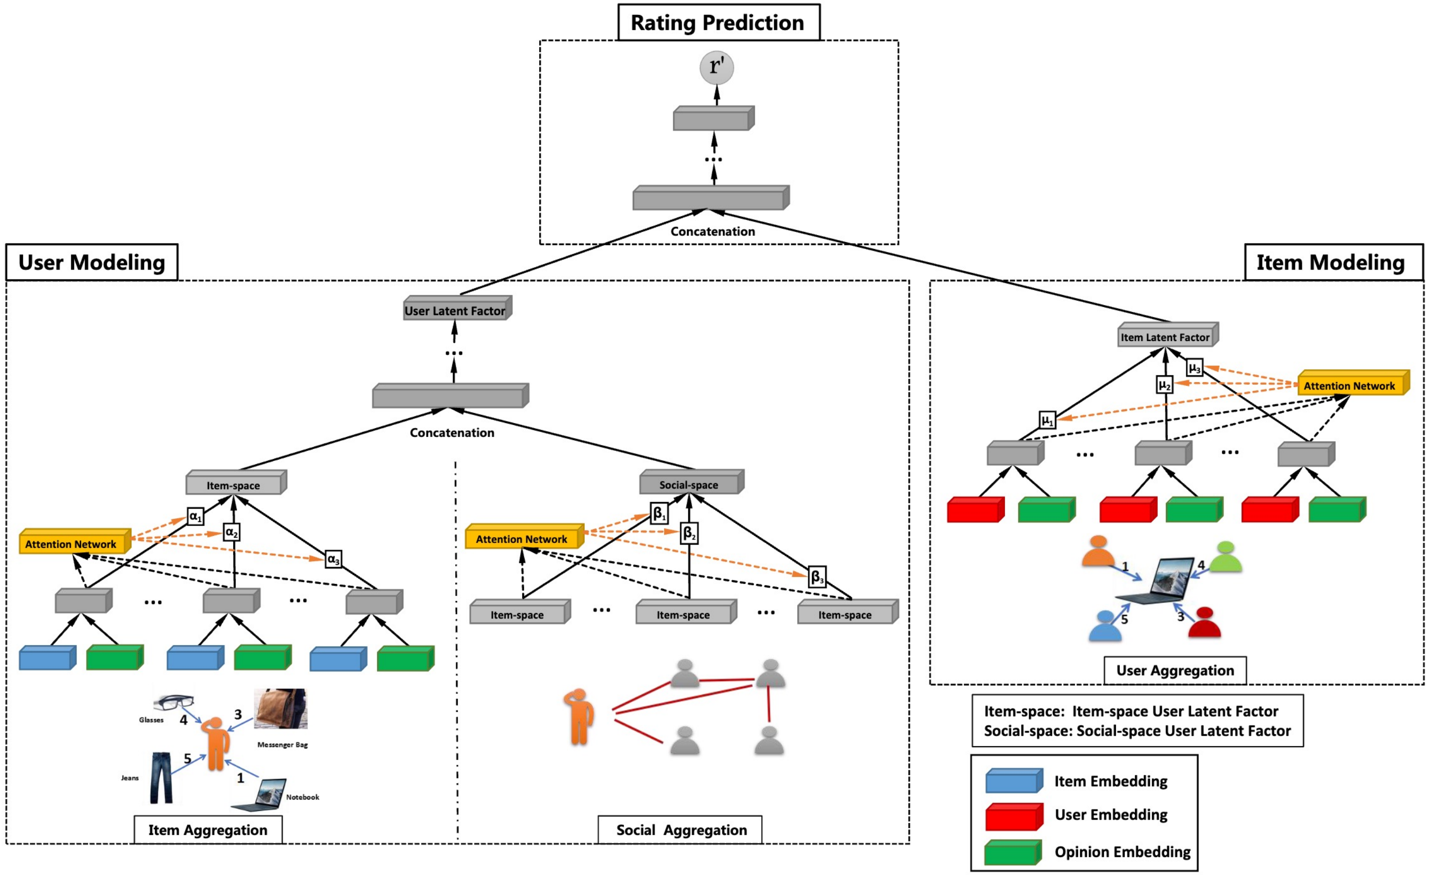

"Graph neural networks for social recommendation." WWW (2019).

Preliminaries | Pre-training | Fine-tuning | Prompting | Future Directions

Overview

Presenter:
Yunqing Liu

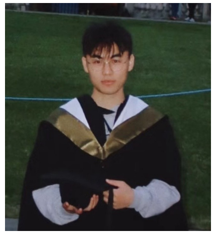

RecSys

- Collaborative Filtering (CF)
- Content-based Recommendation
- Deep Recommender Systems

LLMs

- Encoder-Only
- Decoder-Only
- Encoder-Decoder

Zoom ID:
91649466943
Password:
202312

Website QR Code

LLMs to RecSys

- ❑ LLMs can be used for a variety of tasks, such as **Image Generation**

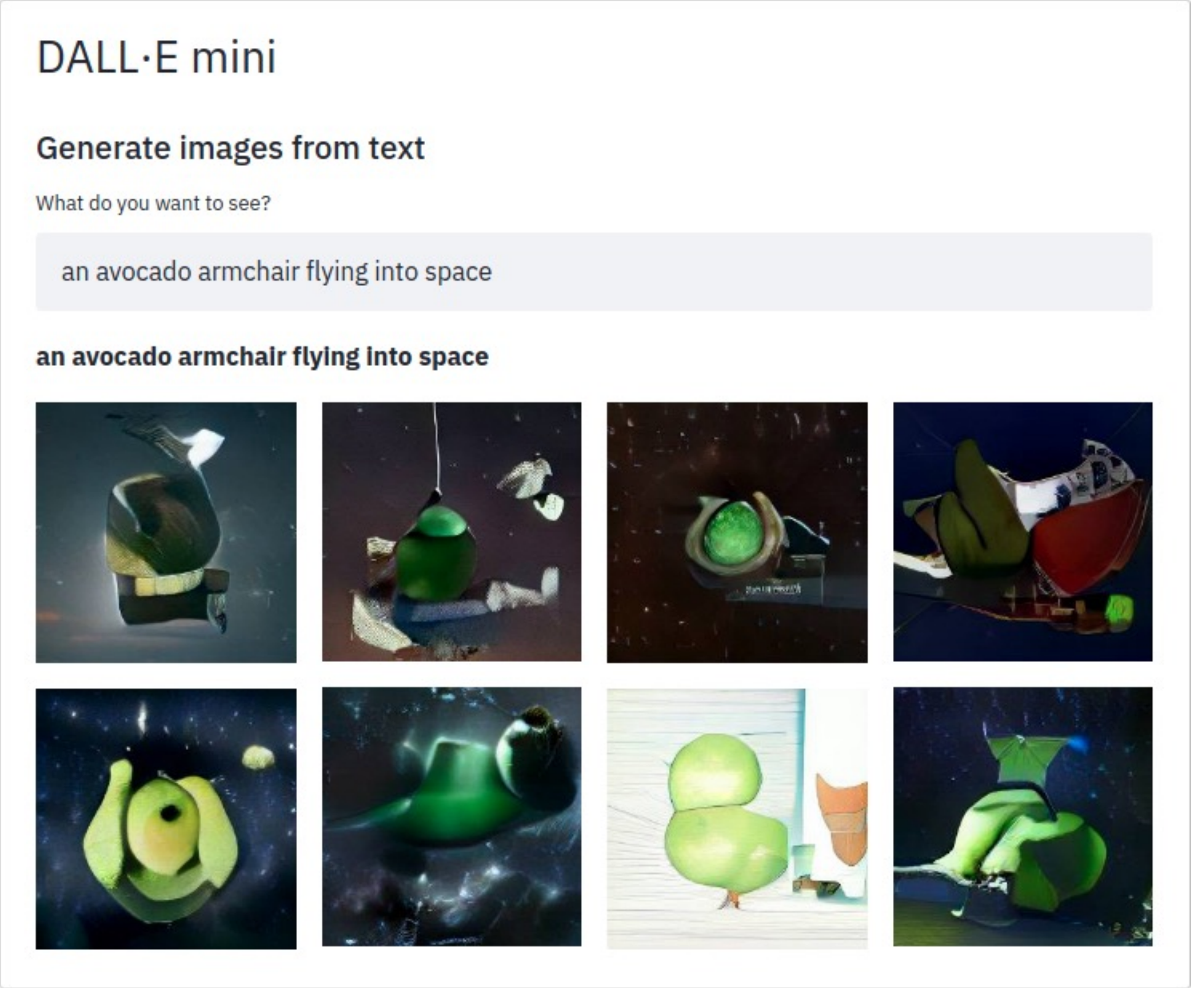

Text to Image

"Zero-shot text-to-image generation." *International Conference on Machine Learning*. PMLR, 2021.

LLMs to RecSys

- LLMs can be used for a variety of tasks, such as **Molecule Generation**

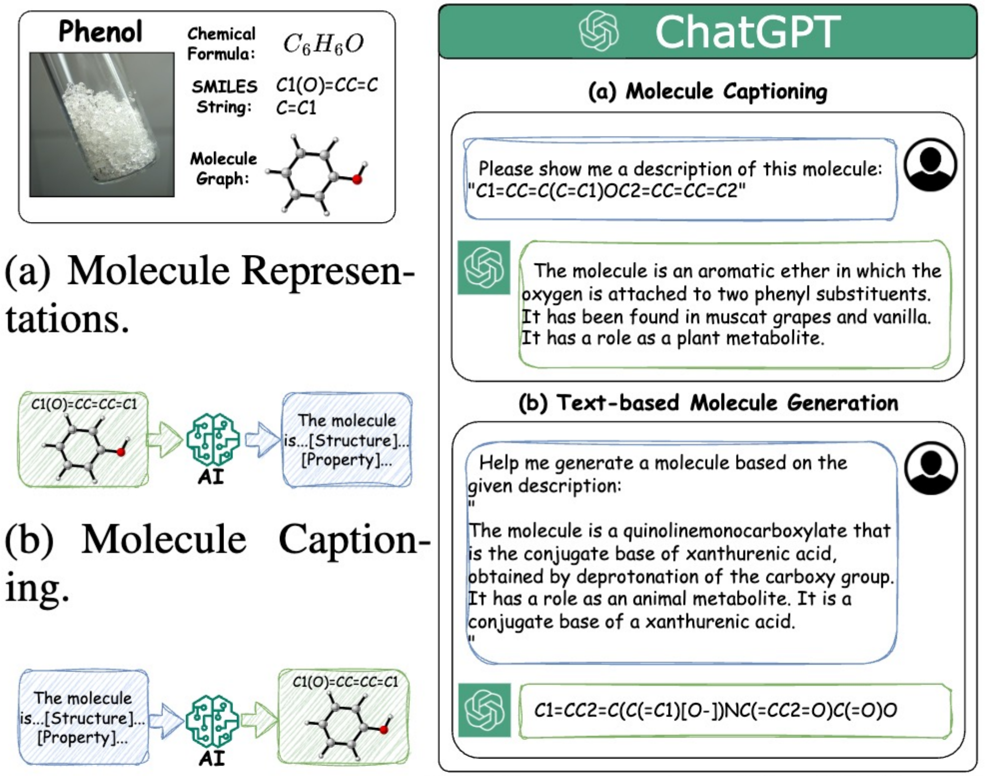

(a) Molecule Representations.

(b) Molecule Captioning.

(c) Text-based Molecule Generation.

(d) Empowering ChatGPT with molecule captioning and text-based molecule generation abilities.

Text to Molecule

"Empowering Molecule Discovery for Molecule-Caption Translation with Large Language Models: A ChatGPT Perspective." *arXiv* (2023).

LLMs to RecSys

- LLMs can be used for a variety of tasks, such as **Recommendation**

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&reg, 5.0 ; 4.Nail Art Jumbo Stamp Stamping Manicure Image Plate 6 Happy Holidays by Cheeky&reg, 5.0; Based on above rating history, please predict user's rating for the product: "SHANY Nail Art Set (24 Famouse Colors Nail Art Polish, Nail Art Decoration)", (1 being lowest and5 being highest,The output should be like: (x stars, xx%), do not explain the reason.)

Sequential Recommendation

zero-shot	Requirements: you must choose 10 items for recommendation and sort them in order of priority, from highest to lowest. Output format: a python list. Do not explain the reason or include any other words. The user has interacted with the following items in chronological order: ['Better Living Classic Two Chamber Dispenser, White', 'Andre Silhouettes Shampoo Cape, Metallic Black', ……, 'John Frieda JFHA5 Hot Air Brush, 1.5 inch'].Please recommend the next item that the user might interact with.
few-shot	Requirements: you must choose 10 items for recommendation and sort them in order of priority, from highest to lowest. Output format: a python list. Do not explain the reason or include any other words. Given the user's interaction history in chronological order: ['Avalon Biotin B-Complex Thickening Conditioner, 14 Ounce', 'Conair 1600 Watt Folding Handle Hair Dryer', ……, 'RoC Multi-Correxion 4-Zone Daily Moisturizer, SPF 30, 1.7 Ounce'], the next interacted item is ['Le Edge Full Body Exfoliator - Pink']. Now, if the interaction history is updated to ['Avalon Biotin B-Complex Thickening Conditioner, 14 Ounce', 'Conair 1600 Watt Folding Handle Hair Dryer',……, 'RoC Multi-Correxion 4-Zone Daily Moisturizer, SPF 30, 1.7 Ounce', 'Le Edge Full Body Exfoliator - Pink'] and the user is likely to interact again, recommend the next item.

Text to Recommendation

"Is ChatGPT a good recommender? a preliminary study." CIKM (2023).

What are Language Models?

- Narrow Sense

  - A **probabilistic model** that assigns a probability to every **finite sequence** (grammatical or not)

Sentence: "the cat sat on the mat"

$$P(\text{the cat sat on the mat}) = P(\text{the}) * P(\text{cat}|\text{the}) * P(\text{sat}|\text{the cat})$$
$$* P(\text{on}|\text{the cat sat}) * P(\text{the}|\text{the cat sat on})$$
$$* P(\text{mat}|\text{the cat sat on the})$$

Implicit order

- Broad Sense

  - Encoder-only models (BERT, RoBERTa, ELECTRA)
  - Decoder-only models (GPT-X, OPT, LLaMa, PaLM)
  - Encoder-decoder models (T5, BART)

Large Language Models

- ❑ Trained on more and more data – **Hundreds of Billions of Tokens**

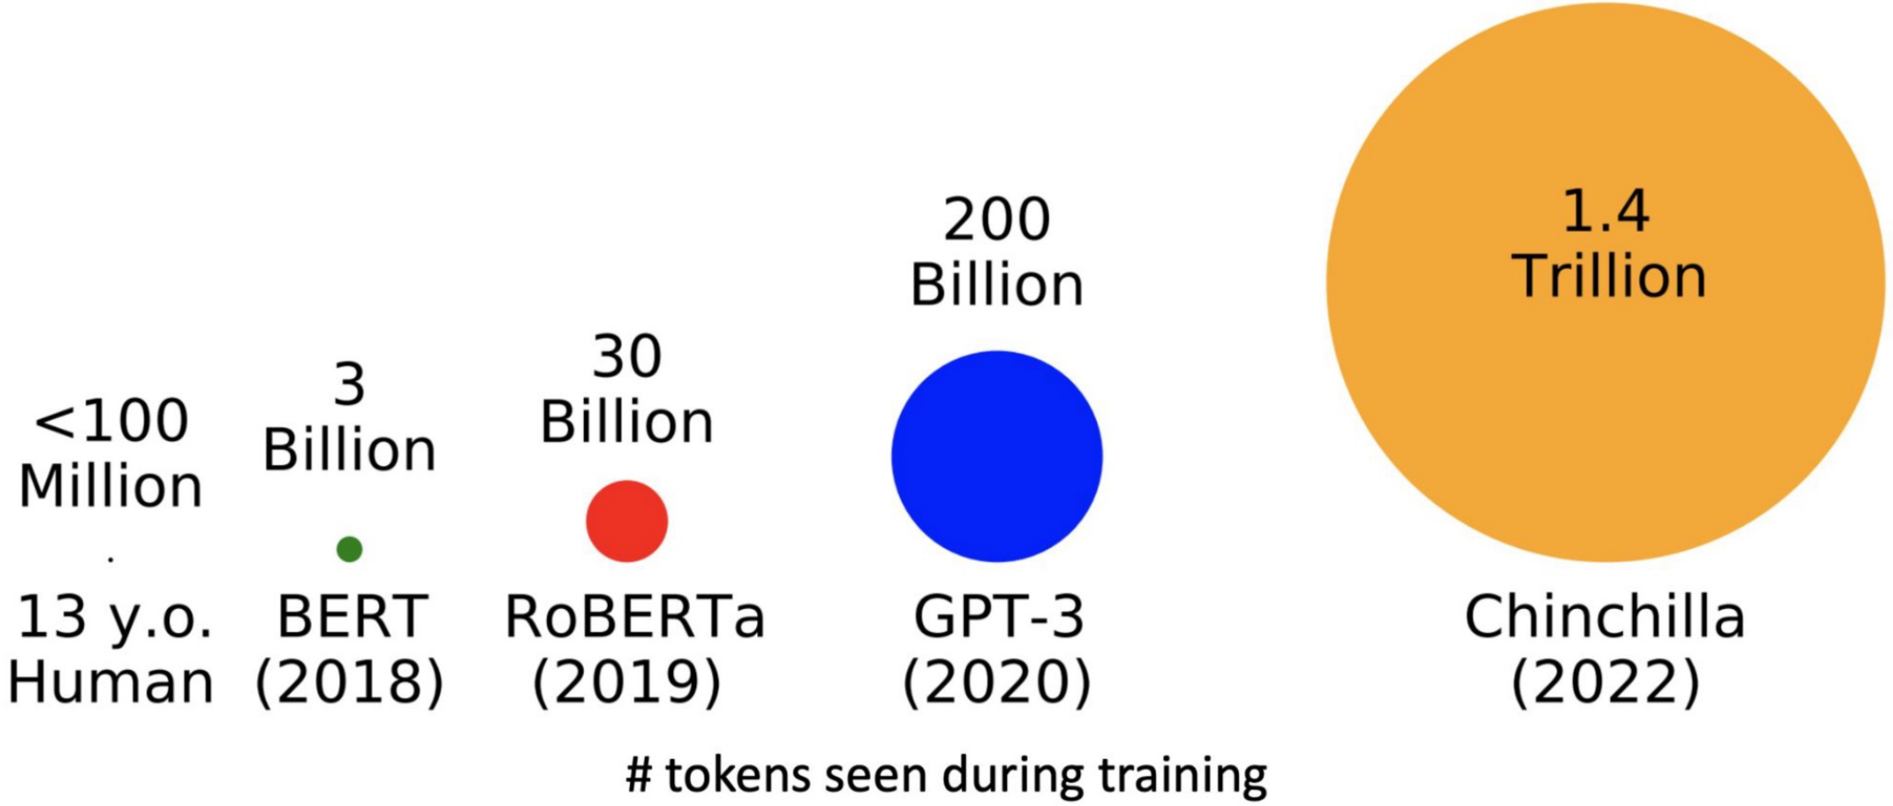

https://babylm.github.io/

Large Language Models

- Larger and larger models – **Billions of Parameters**

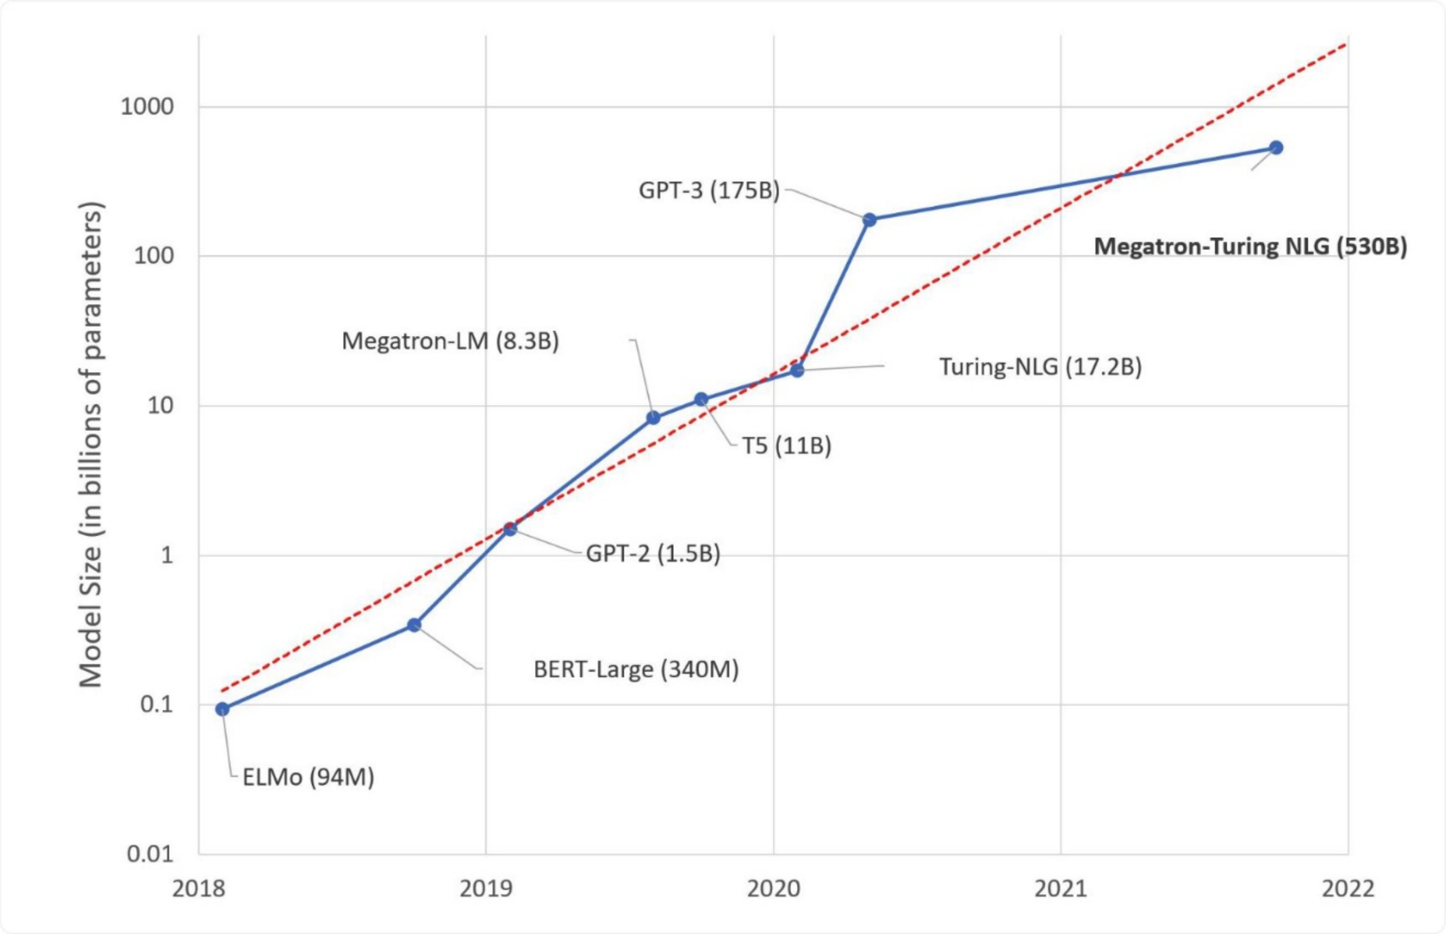

https://huggingface.co/blog/large-language-models

Why Large Language Models

Scaling Law for Neural Language Models

- Performance depends strongly on scale! We keep getting better performance as we scale the model, data, and compute up!

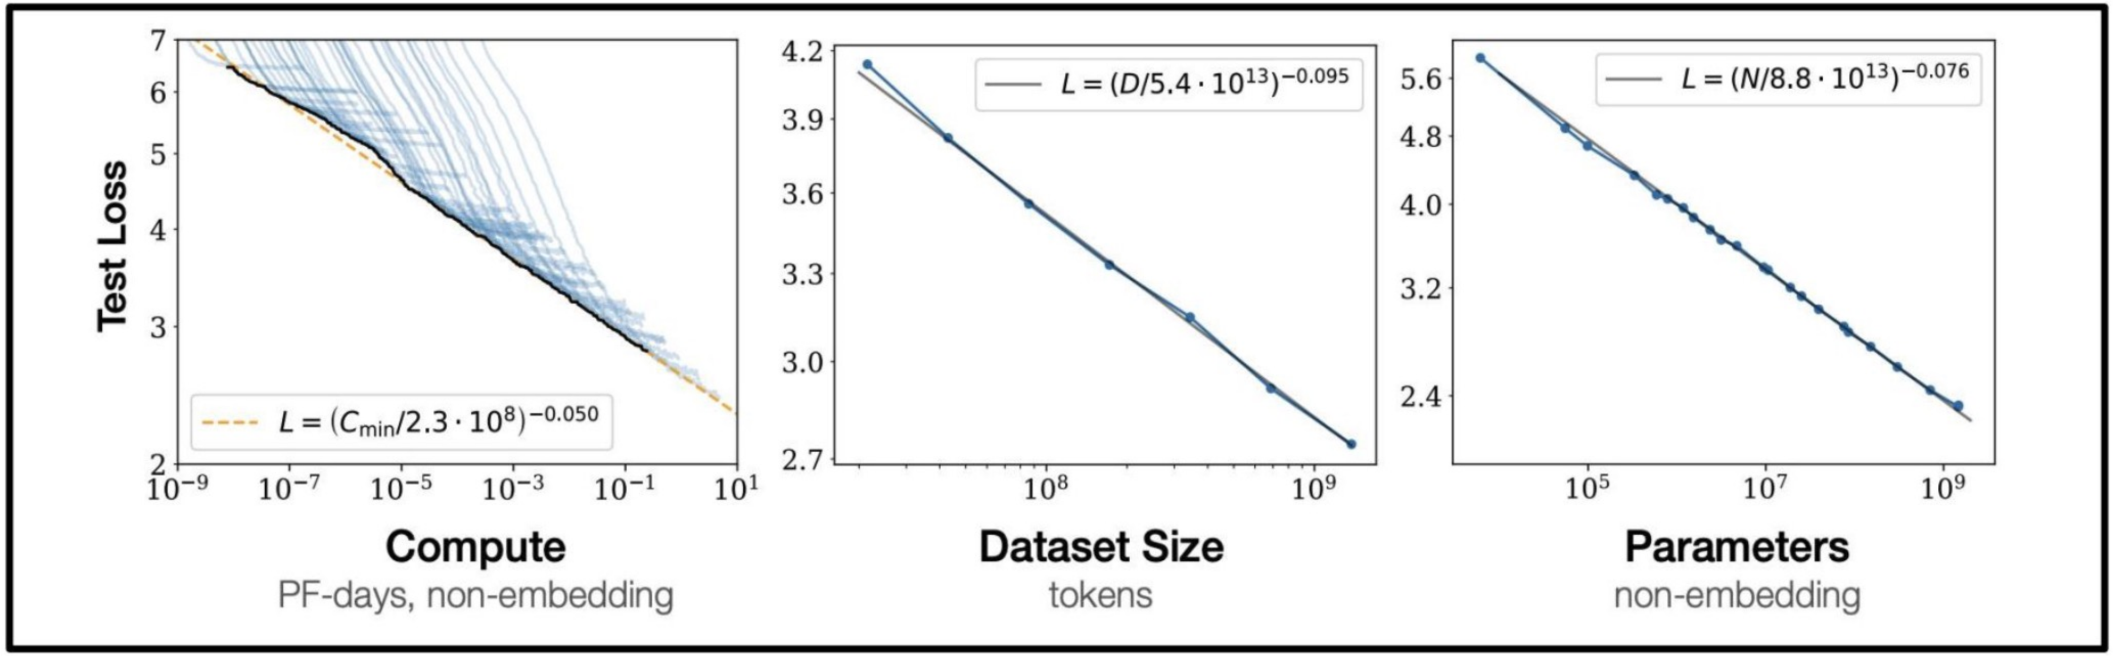

"Scaling laws for neural language models." arXiv preprint (2020).

Why Large Language Models?

- Generalization

    - We can now use one single model to solve many NLP tasks.

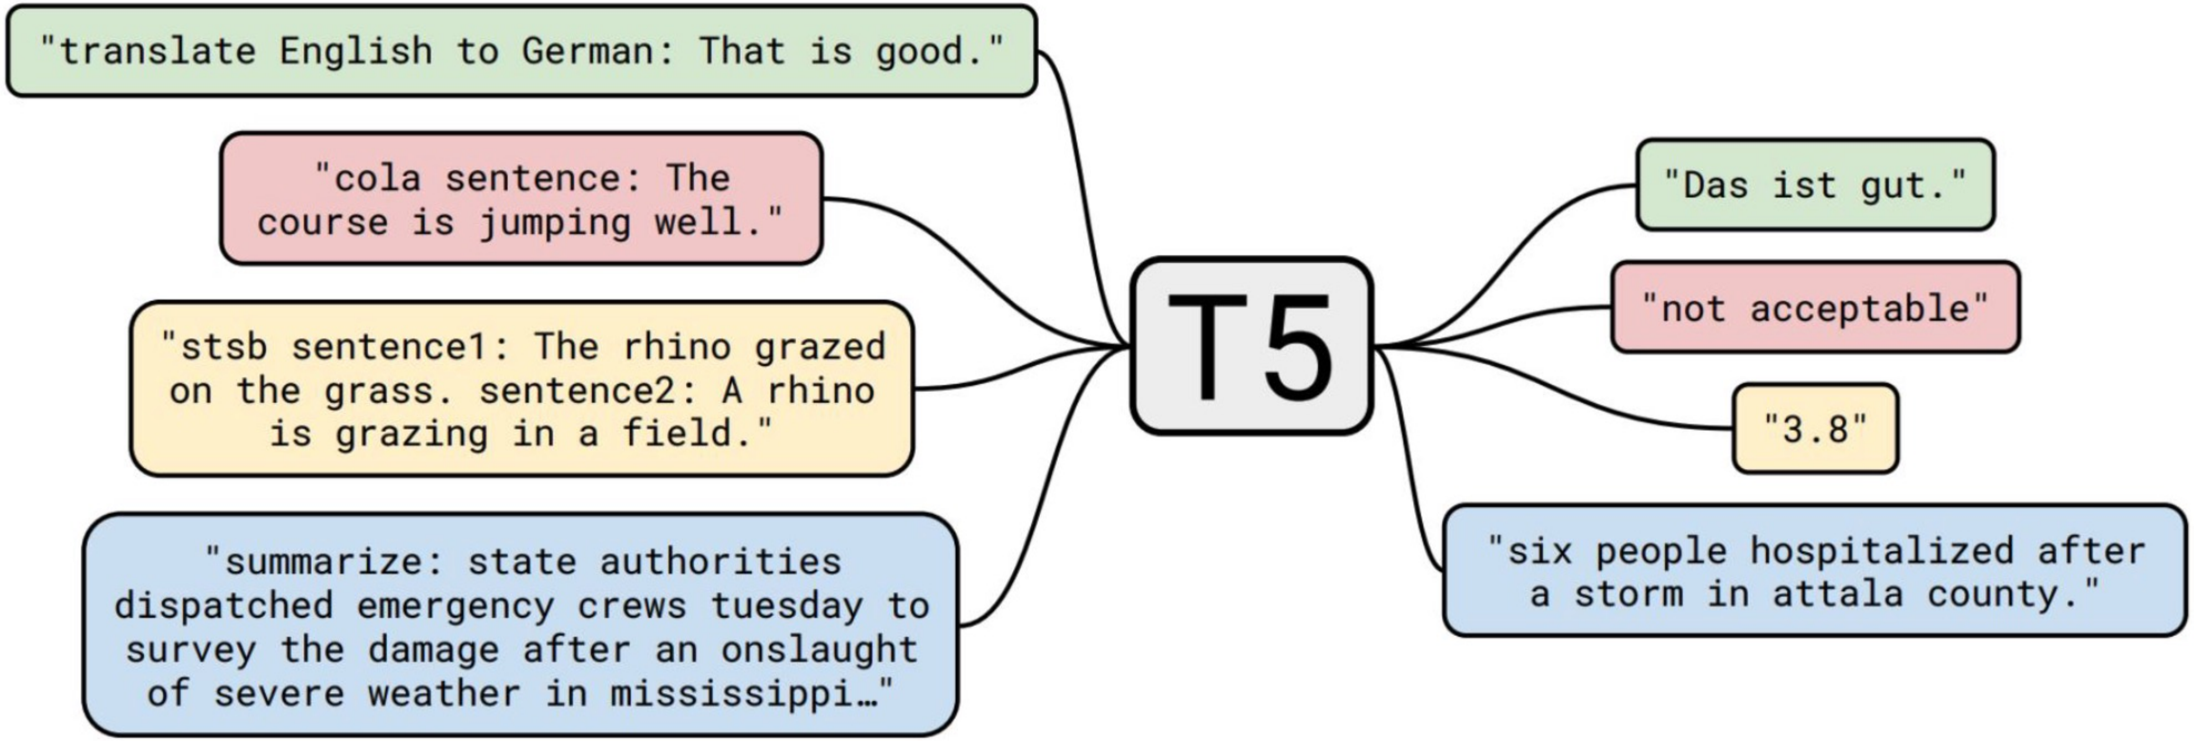

"Exploring the limits of transfer learning with a unified text-to-text transformer." The Journal of Machine Learning Research (2020).

Why Large Language Models?

- The Strong Zero-shot/Few-shot Ability

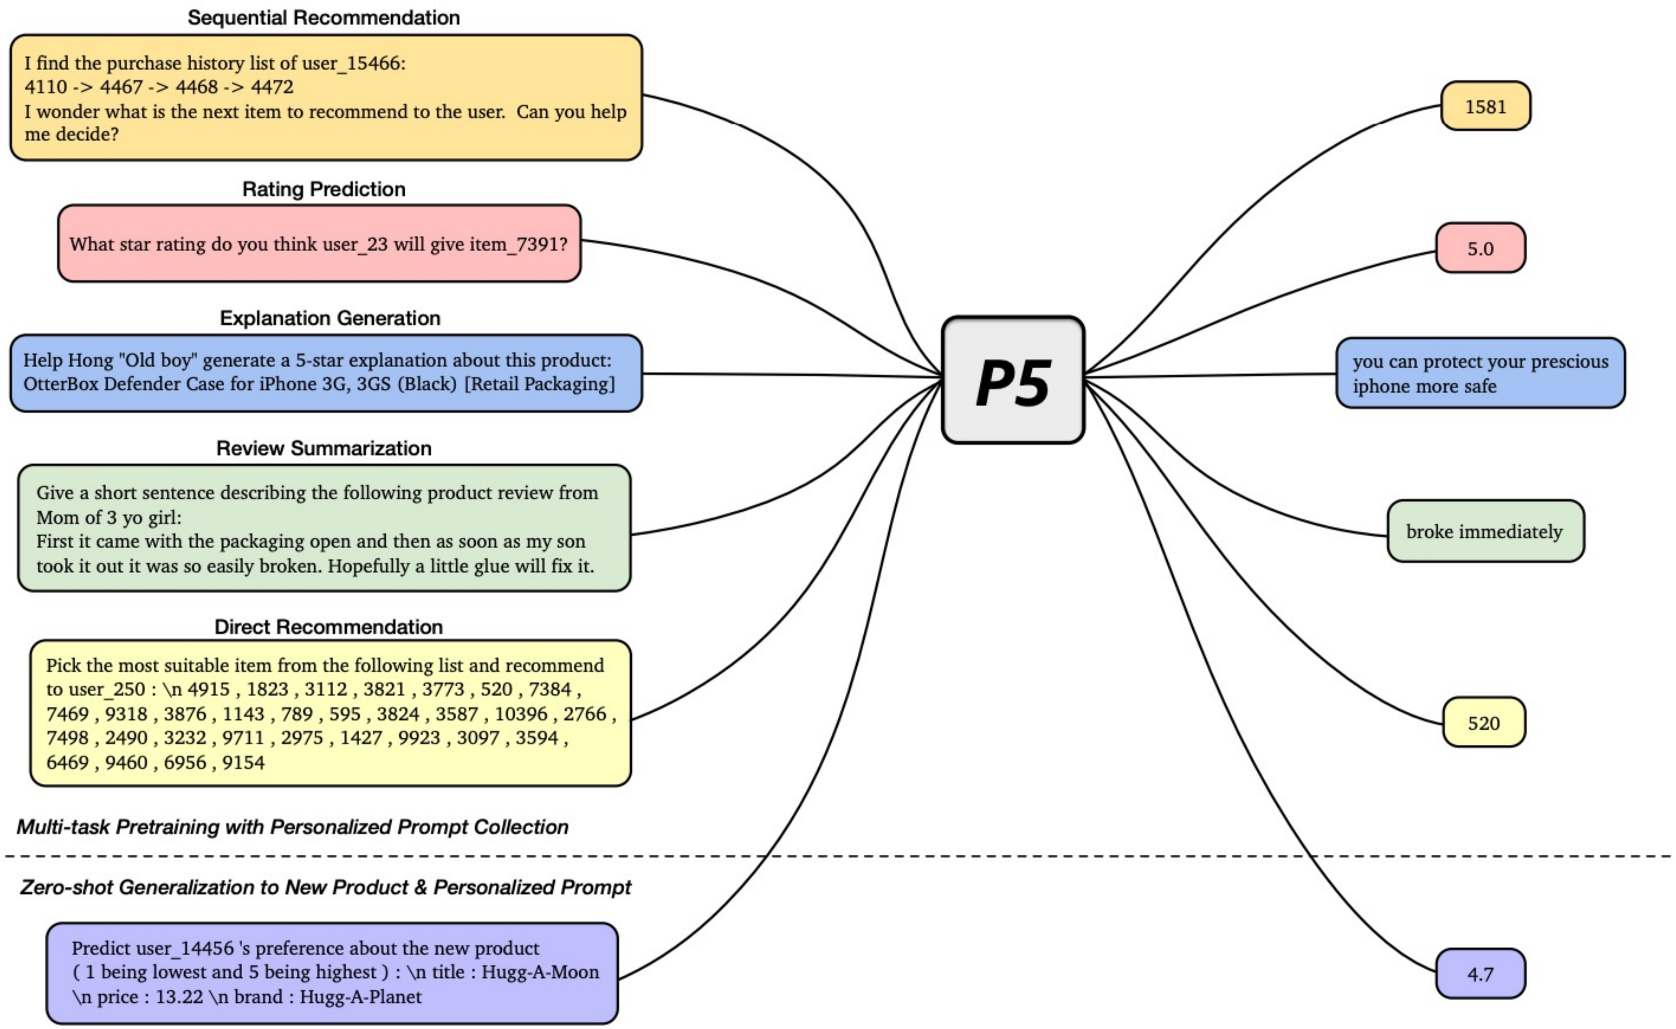

"Recommendation as language processing (rlp): A unified pretrain, personalized prompt & predict paradigm (p5)." RecSys (2022).

Large Language Models

- Encoder-Only Models
  - BERT, RoBERTa, ELECTRA
- Decoder-Only Models
  - GPT-X, OPT, LLaMa, PaLM
- Encoder-Decoder Models
  - T5, BART

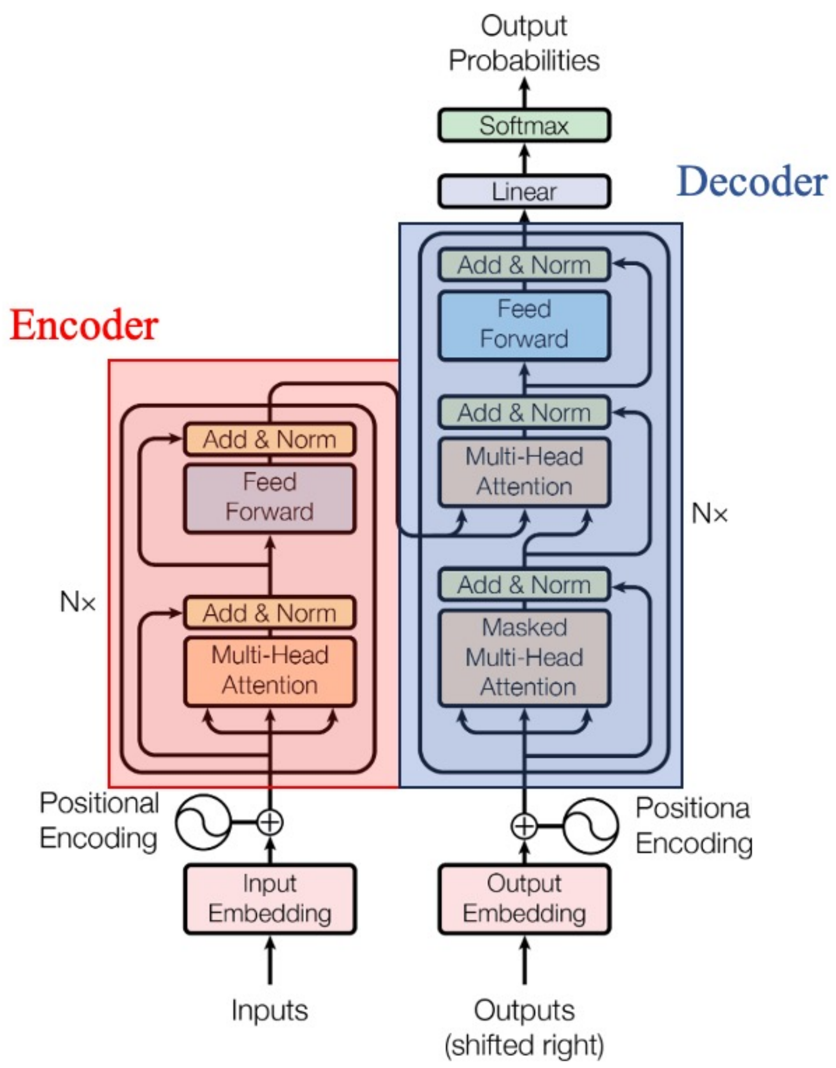

The Transformer – model architecture

Vaswani, Ashish, et al. "Attention is all you need." *Advances in neural information processing systems* 30 (2017).

Transformer

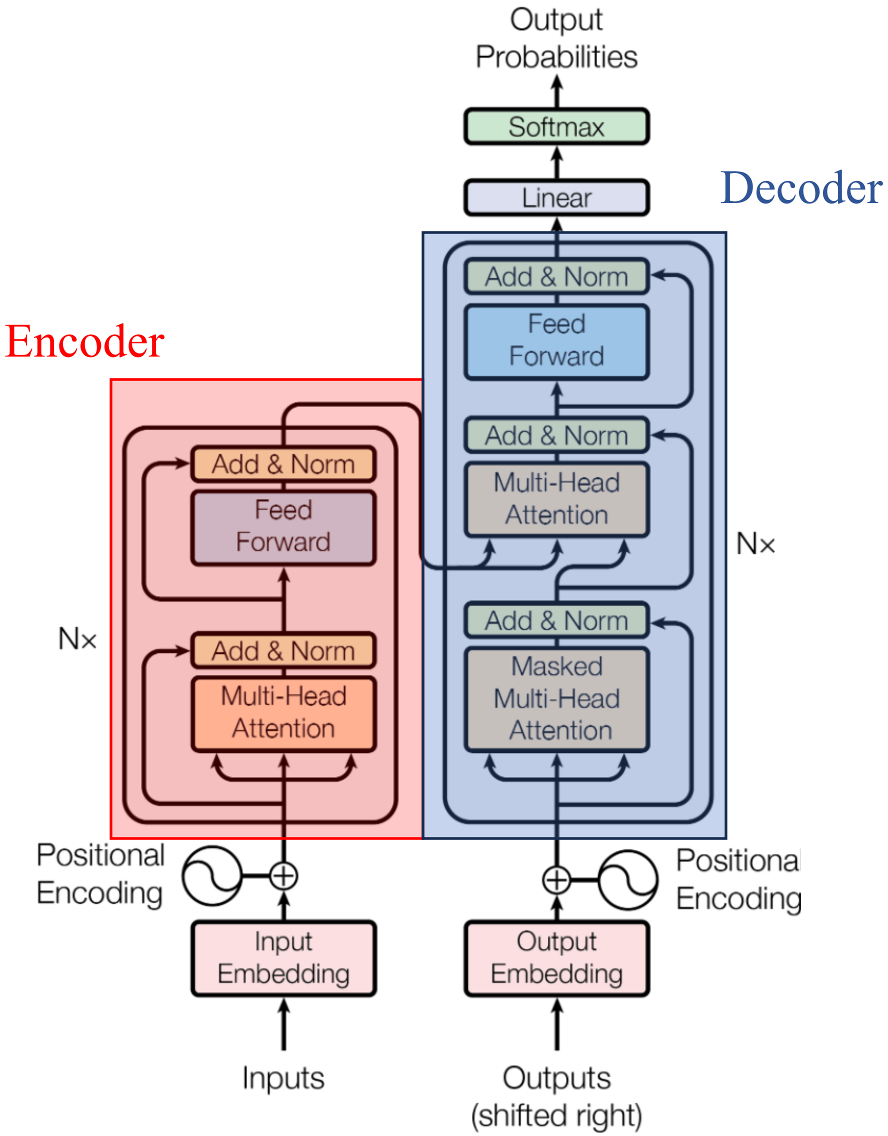

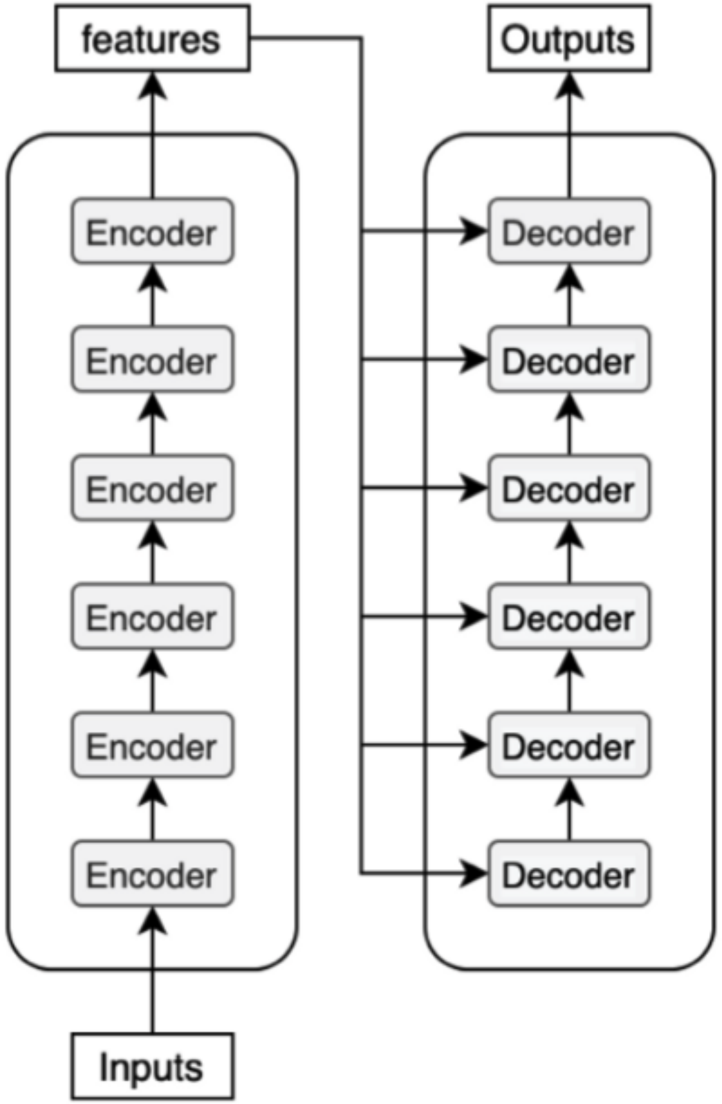

Vaswani, Ashish, et al. "Attention is all you need." *Advances in neural information processing systems* 30 (2017).

Large Language Models

- Encoder-Only
- Decoder-Only
- Encoder-Decoder

BERT

- BERT uses a bidirectional Transformer

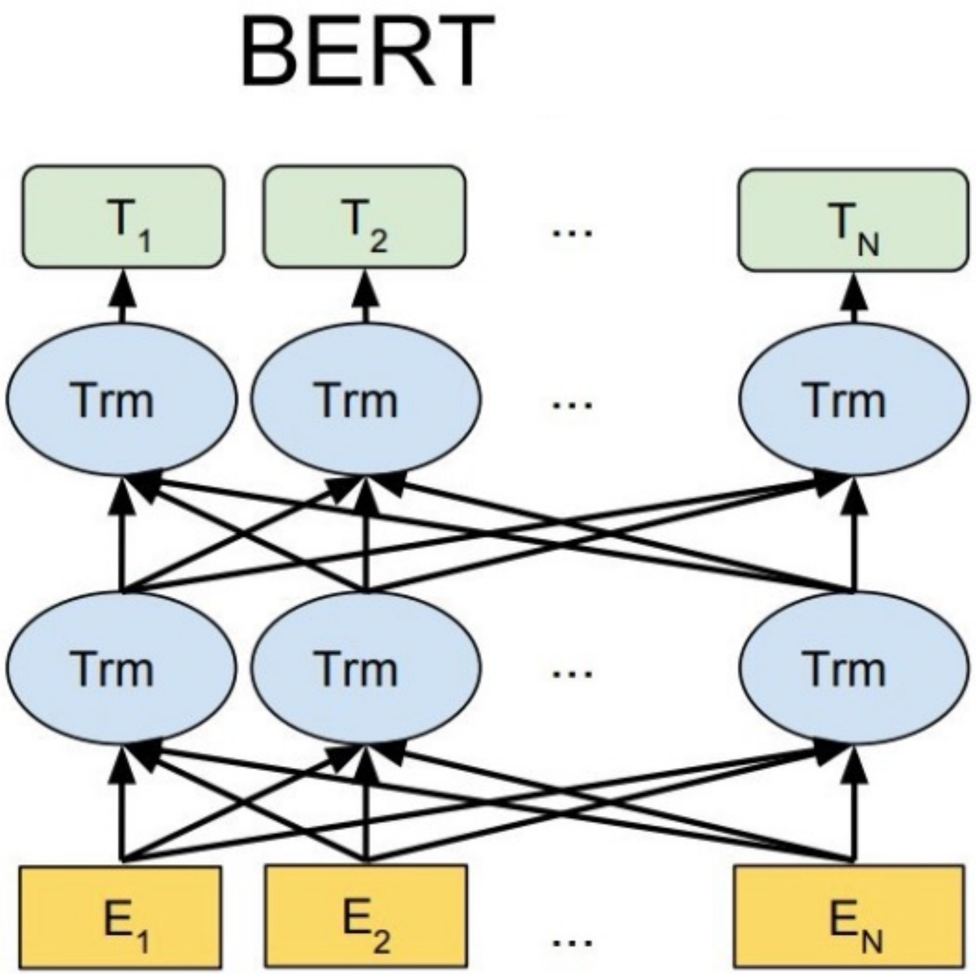

"Bert: Pre-training of deep bidirectional transformers for language understanding." arXiv preprint (2018).

BERT4Rec

- Adopt **Bidirectional Encoder Representations** from Transformers to model the sequential nature of user behaviors

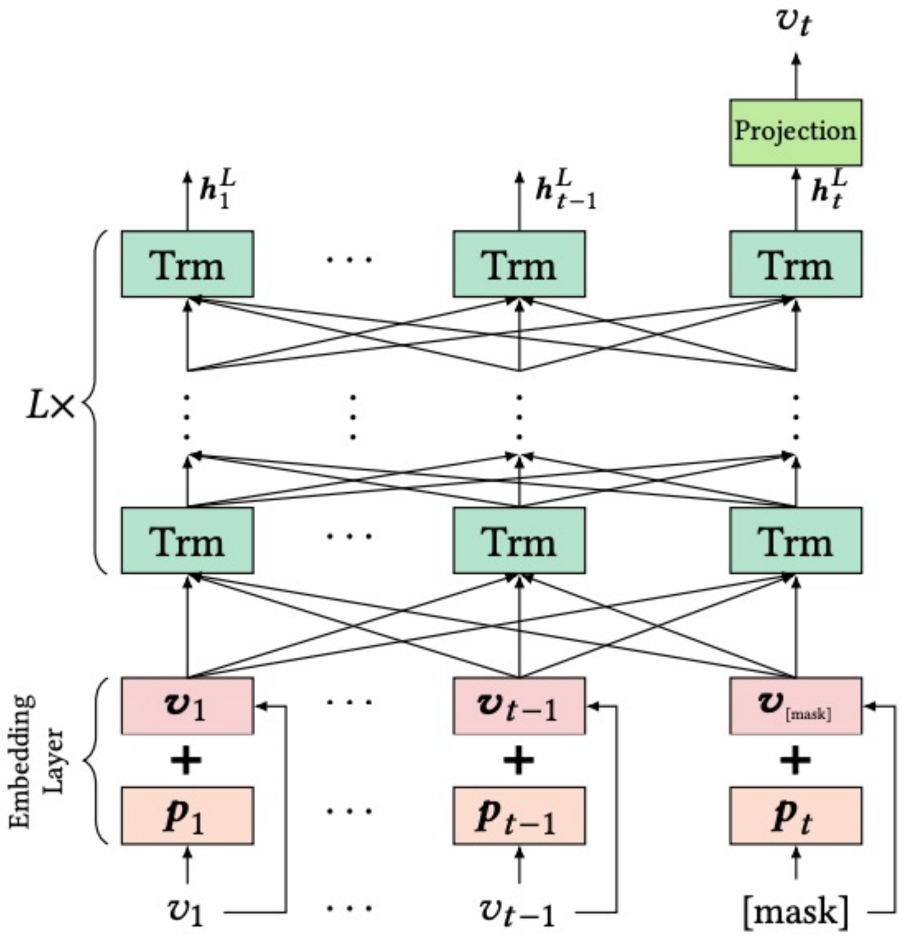

"BERT4Rec: Sequential recommendation with bidirectional encoder representations from transformer." CIKM (2019).

Large Language Models

- Encoder-Only
- Decoder-Only
- Encoder-Decoder

GPT

- OpenAI GPT uses a left-to-right Transformer

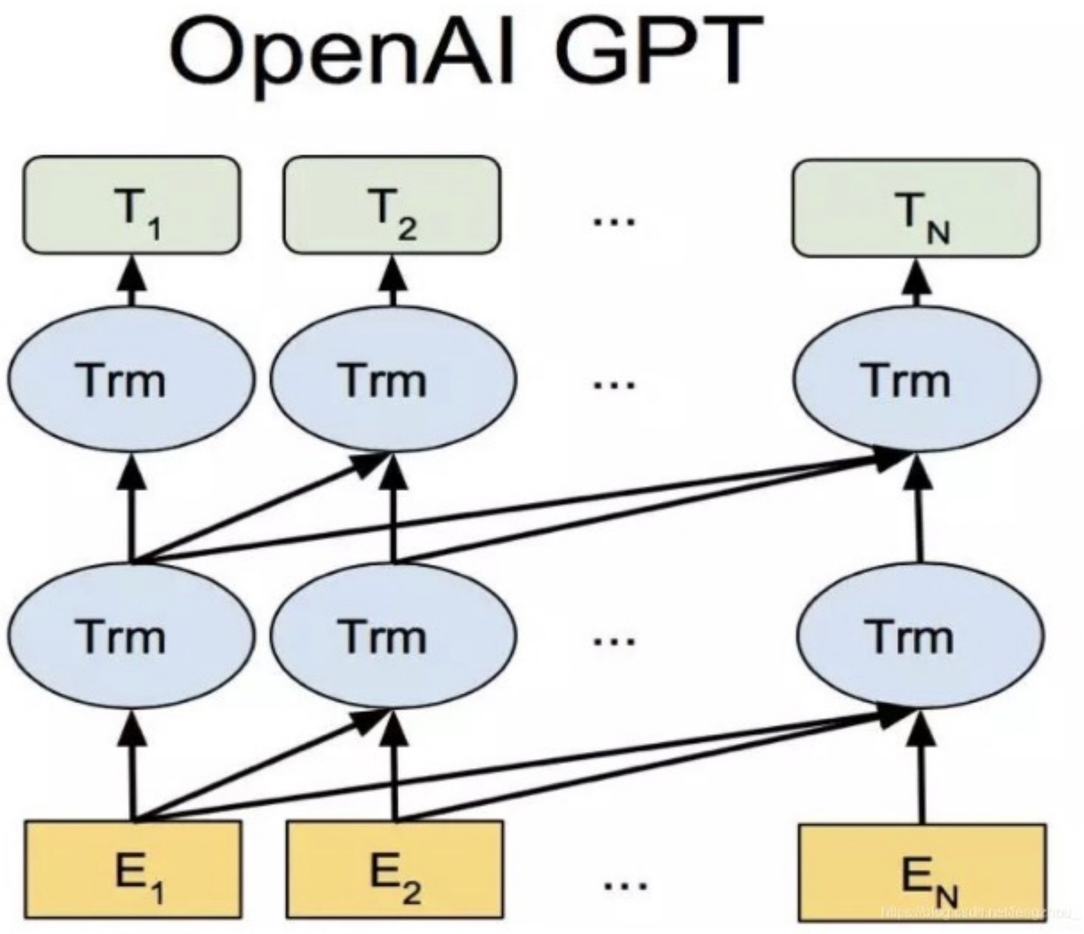

"Improving language understanding by generative pre-training." (2018).

PETER

- ❑ Utilizing the IDs to predict the words in the target explanation

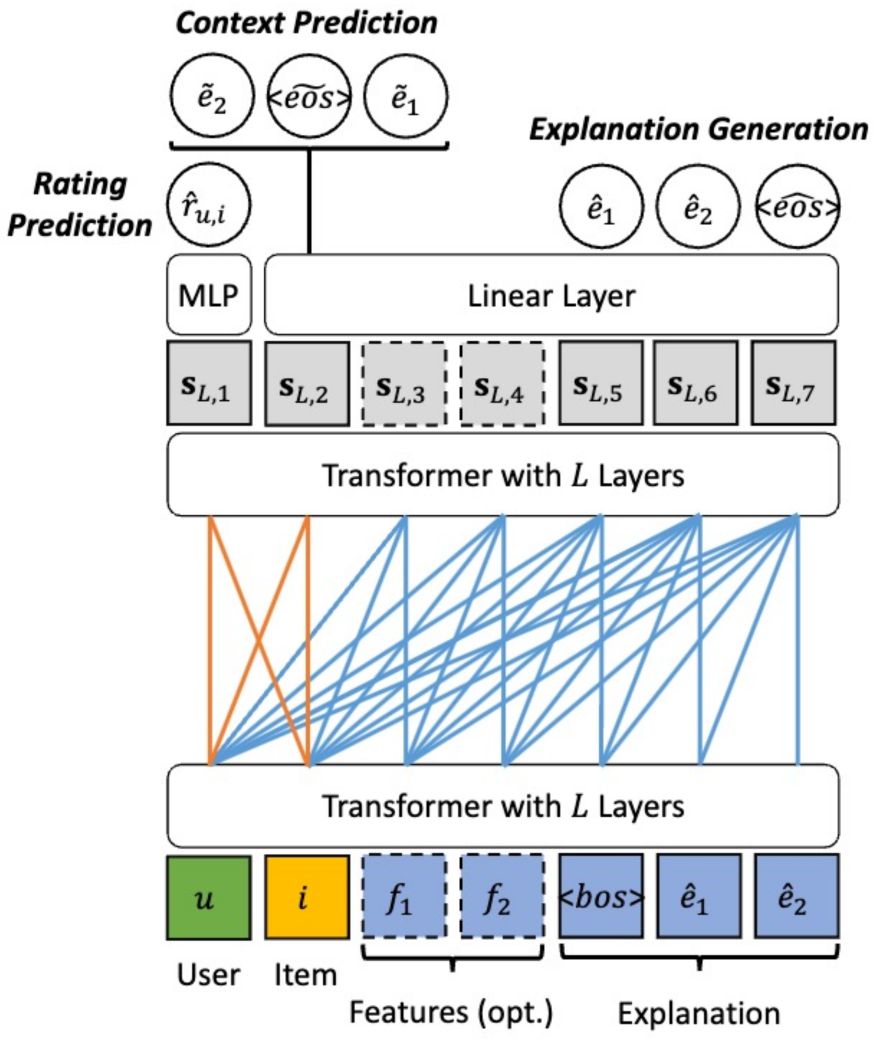

"Personalized transformer for explainable recommendation." ACL-IJCNLP (2021).

Large Language Models

Encoder-Only

Decoder-Only

Encoder-Decoder

T5

- T5 handles any text-to-text task by converting every natural language processing problem into a text generation problem.

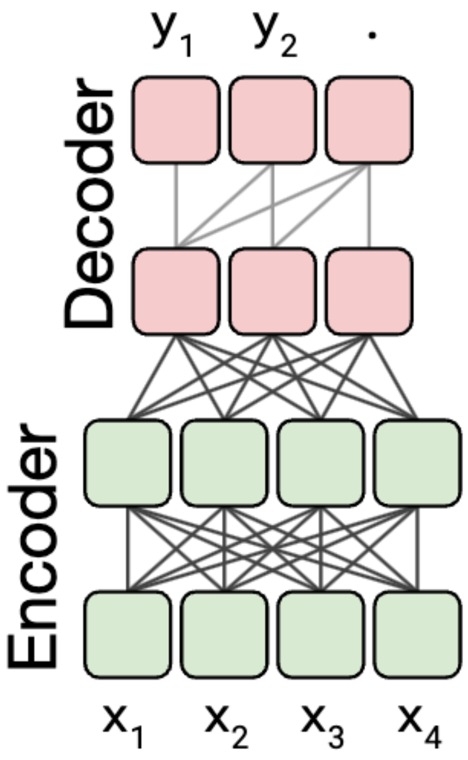

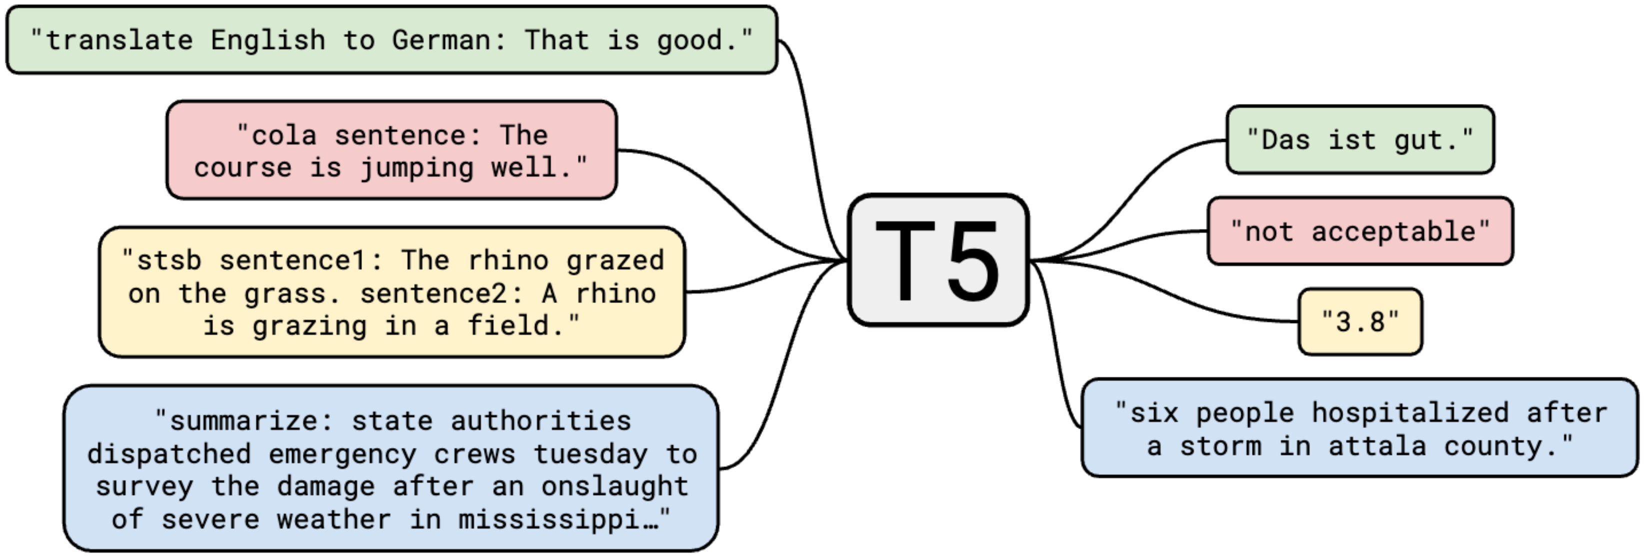

"Exploring the limits of transfer learning with a unified text-to-text transformer." The Journal of Machine Learning Research (2020).

P5

- Text-to-text paradigm - “Pretrain, Personalized Prompt, and Predict Paradigm” (P5) for recommendation

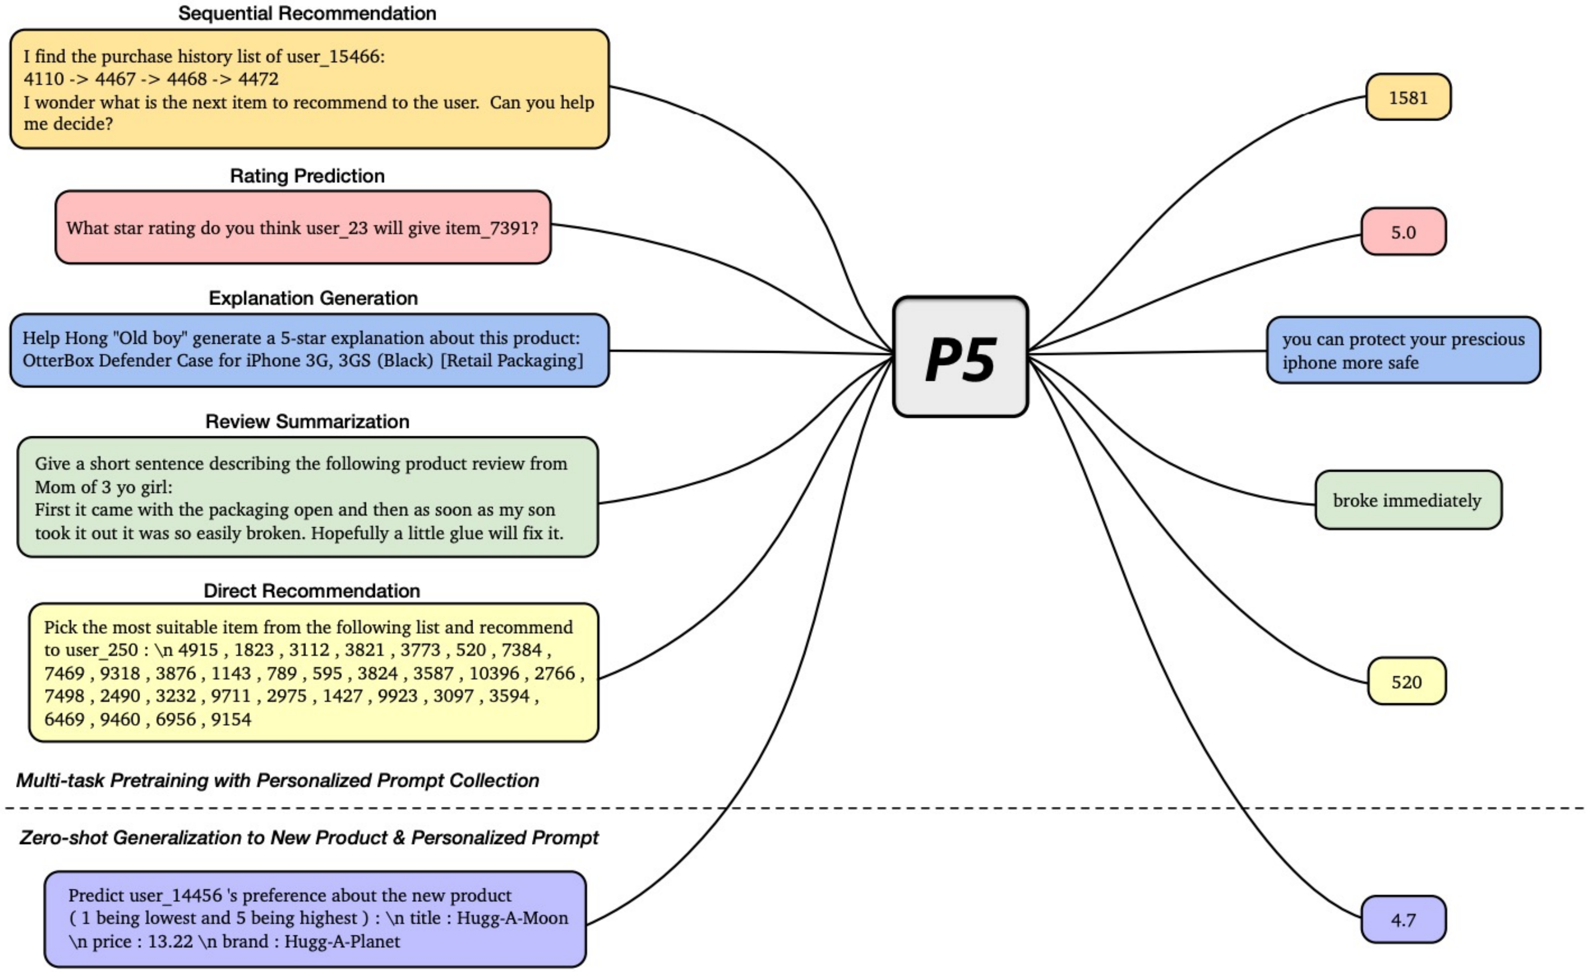

"Recommendation as language processing (rlp): A unified pretrain, personalized prompt & predict paradigm (p5)." RecSys (2022).

Overview

Presenter:
Jiatong Li

User & Item Representation

- ID-based LLM RecSys
- Text-based LLM RecSys

Pre-training

- Pre-training in NLP
- Pre-training LLM-based RecSys

Zoom ID:
91649466943
Password:
202312

Website QR Code

User & Item Representation

- Users and Items can be represented in various ways

○ **User ID**

U8189cf6745fc0d808977bdb0b9f22995

Username: Jack0513

Poster	Movie Name	Numeric ID
	In Broad Daylight	1697292155
	The Marvels	1699436461
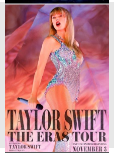	TAYLOR SWIFT | THE ERAS TOUR	1695730583
	The Dark Knight Rises	1699611567
	Oppenheimer	1687513232

User & Item Representation in LLMs

ID-based LLM RecSys

Text-based LLM RecSys

ID-based LLM RecSys

- ❑ Various ways of assigning IDs

	Randomly	Based on Popularity	Based on Time
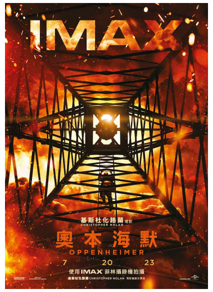	AXGGWD027	01	1687513232
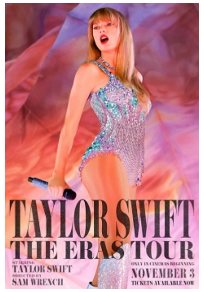	XJSGDG0881	02	1695730583
	BXGW2UD803	03	1699436461

ID-based LLM RecSys

- IDs are originally for unique identification
- However, the embedding of LLMs can not hold millions of items and users

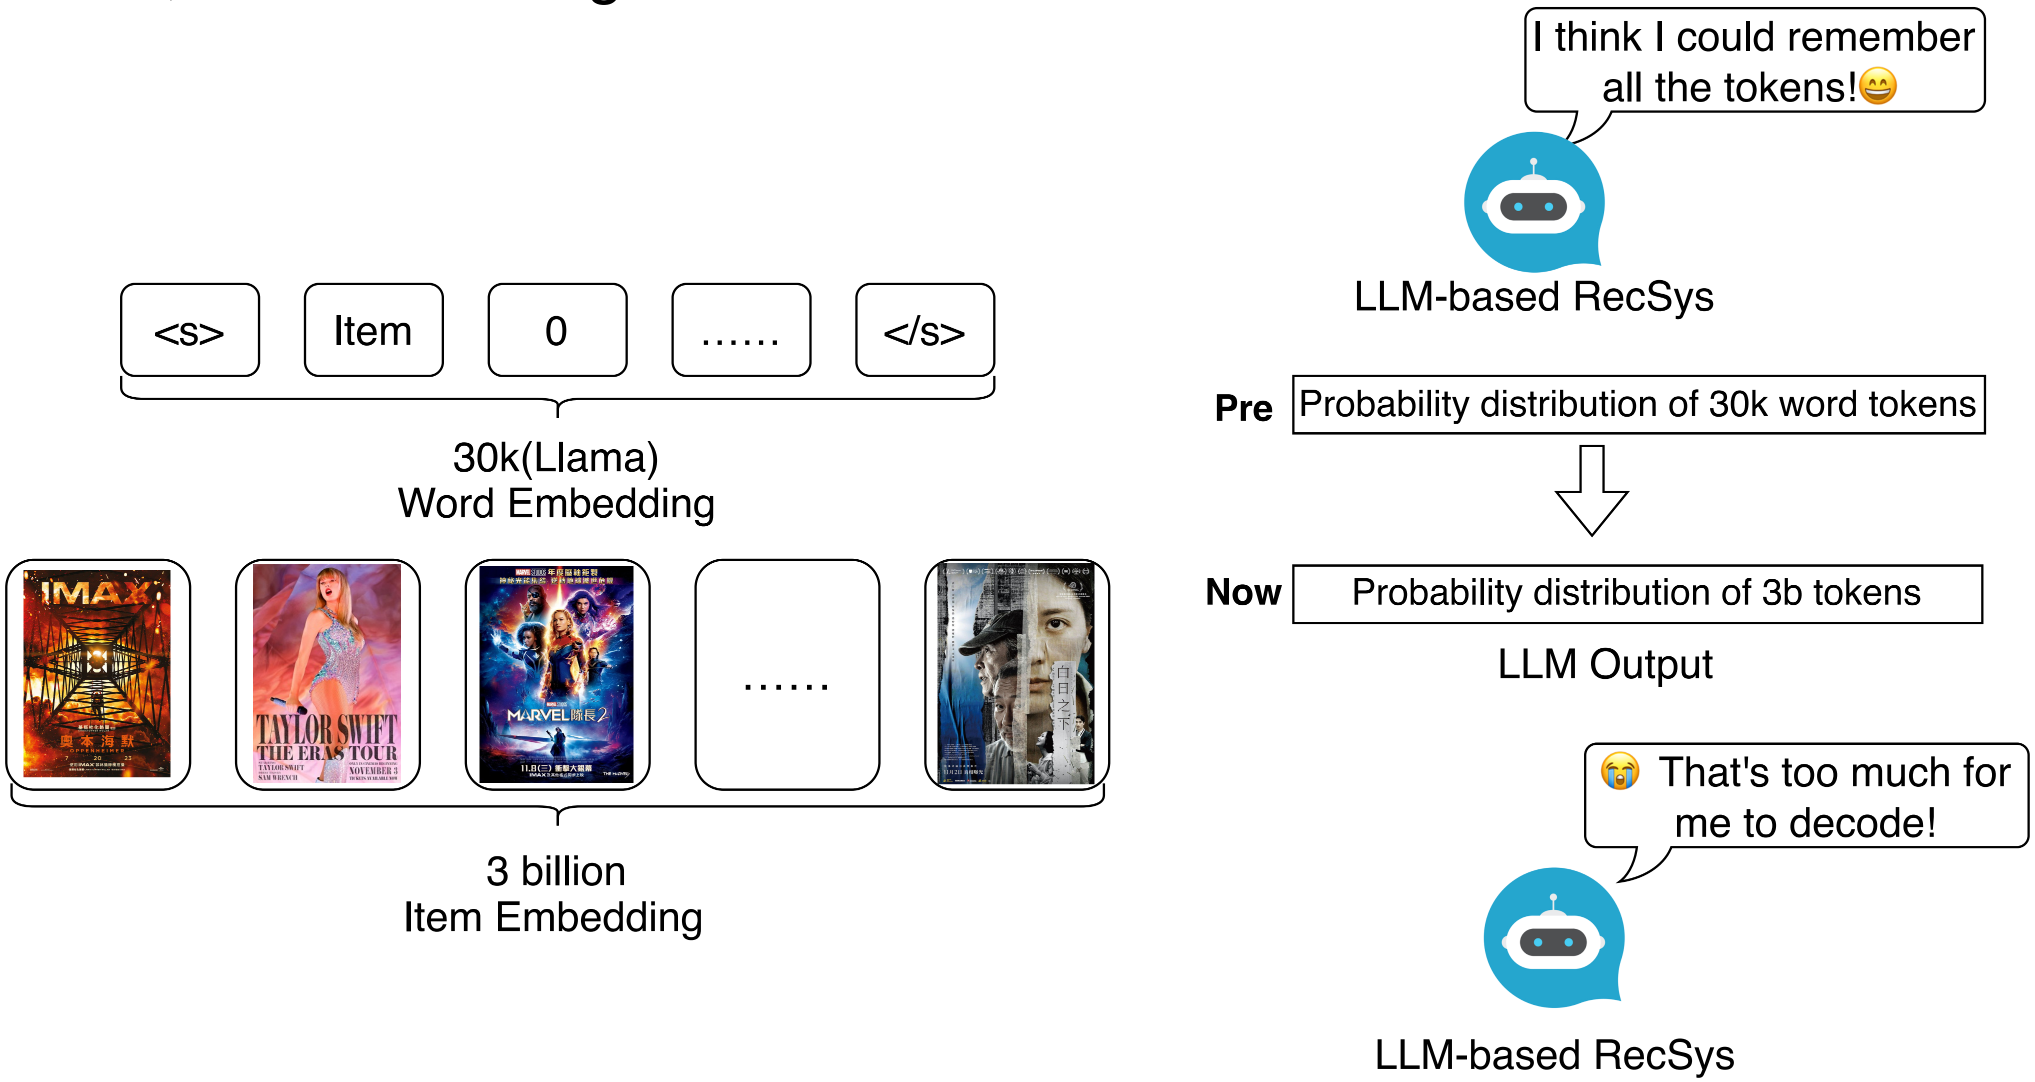

ID-based LLM RecSys

- Normally, we can represent users and items with a span of tokens.
- The format is like "[Prefix]_[ID]". Examples:
  - User_0123 : ["User", "_", "0", "1", "2", "3"]
  - Item_5471 : ["Item", "_", "5", "4", "7", "1"]
  - However, for Item_1003, it could be ["Item", "_", "100", "3"], which might be confusing for LLMs!

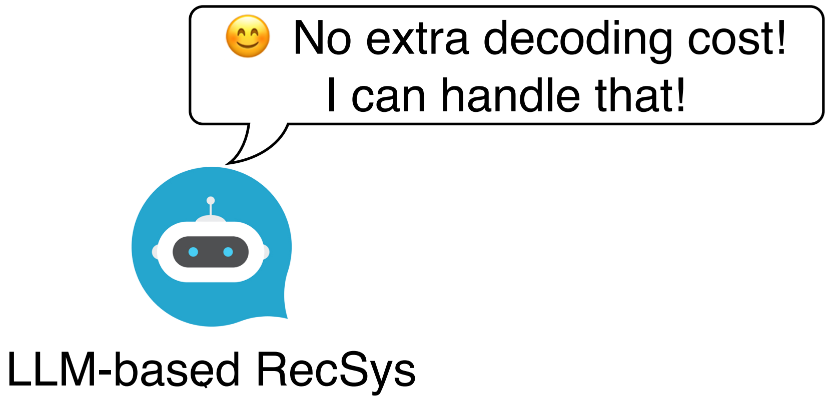

ID-based LLM RecSys

- Indexing methods might affect the performance of RecSys

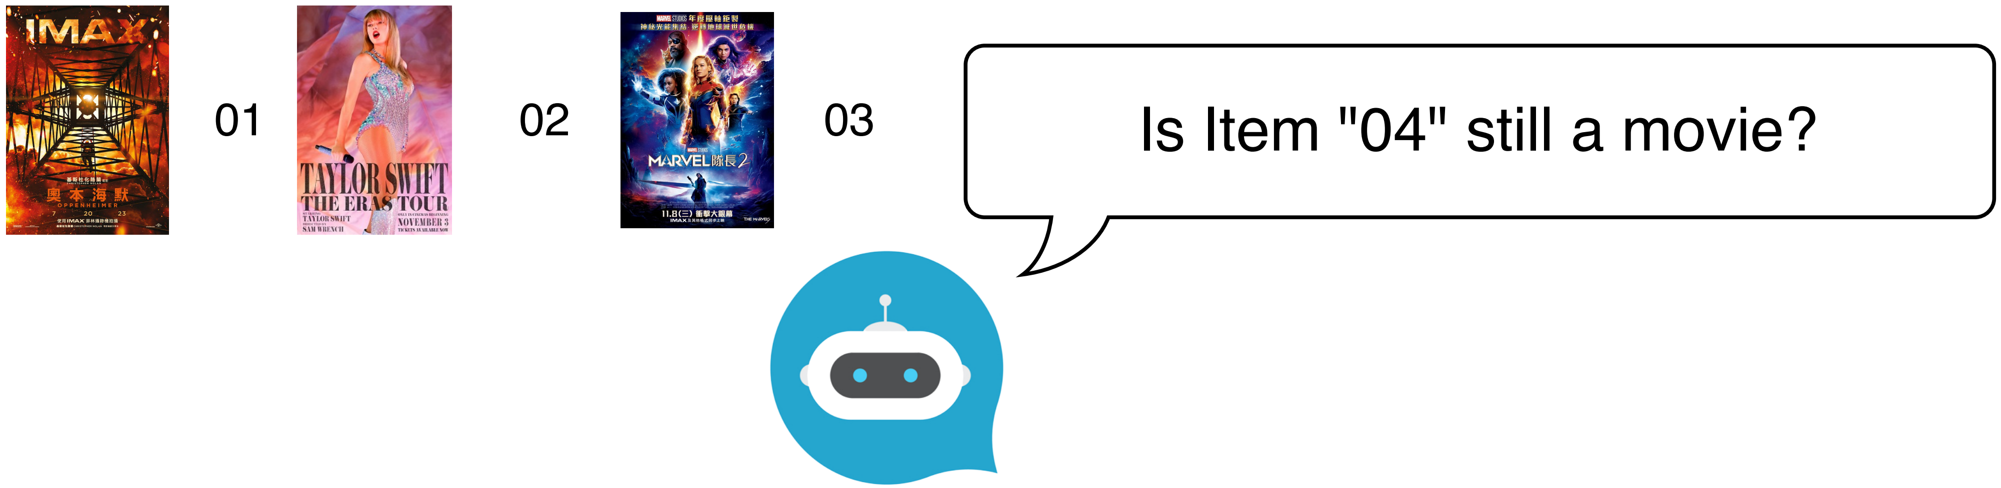

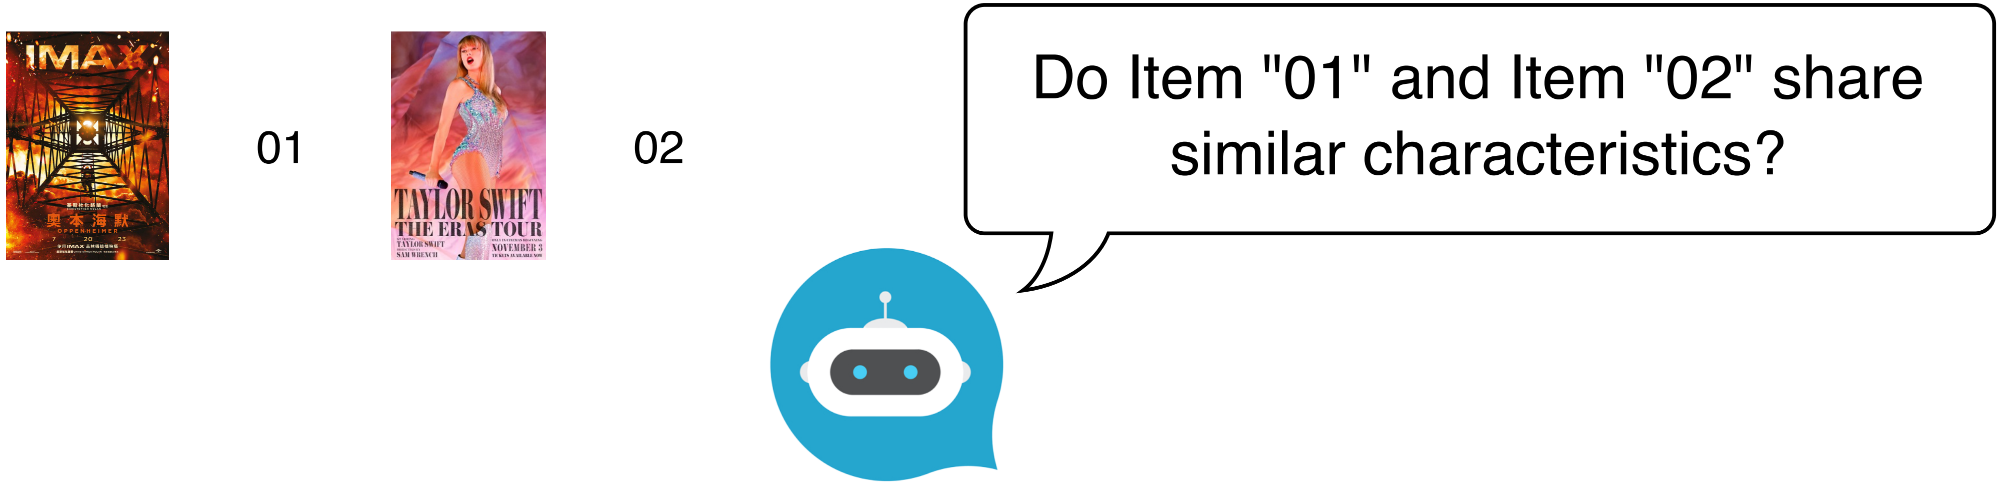

ID-based LLM RecSys

- Introducing more Information to the ID representation
  - Collaborative Indexing

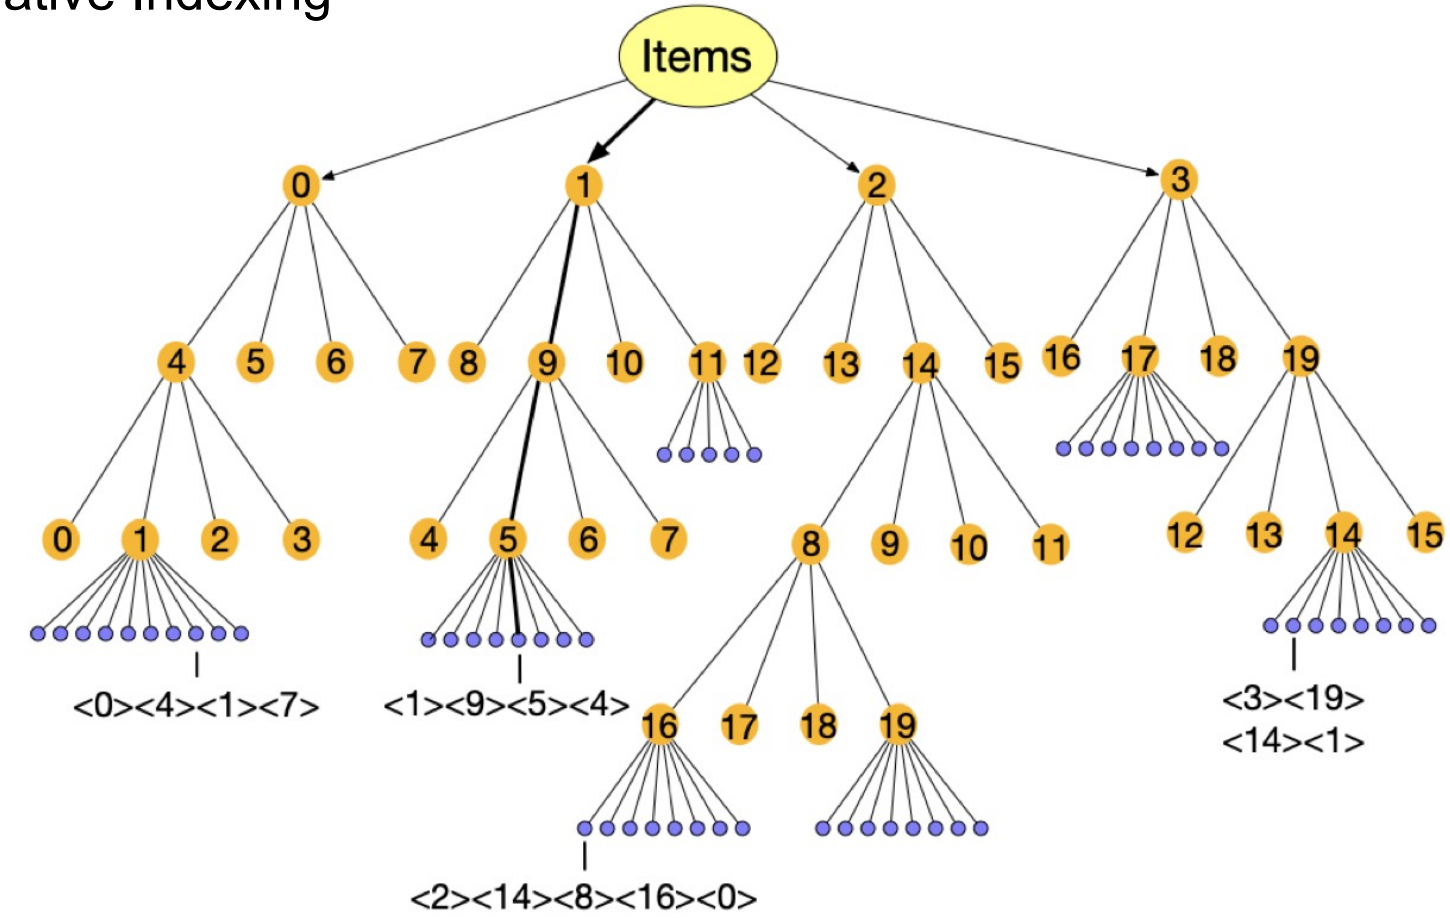

"How to Index Item IDs for Recommendation Foundation Models." arXiv preprint (2023).

ID-based LLM RecSys

- Introducing more Information to the ID representation
  - Semantic Indexing

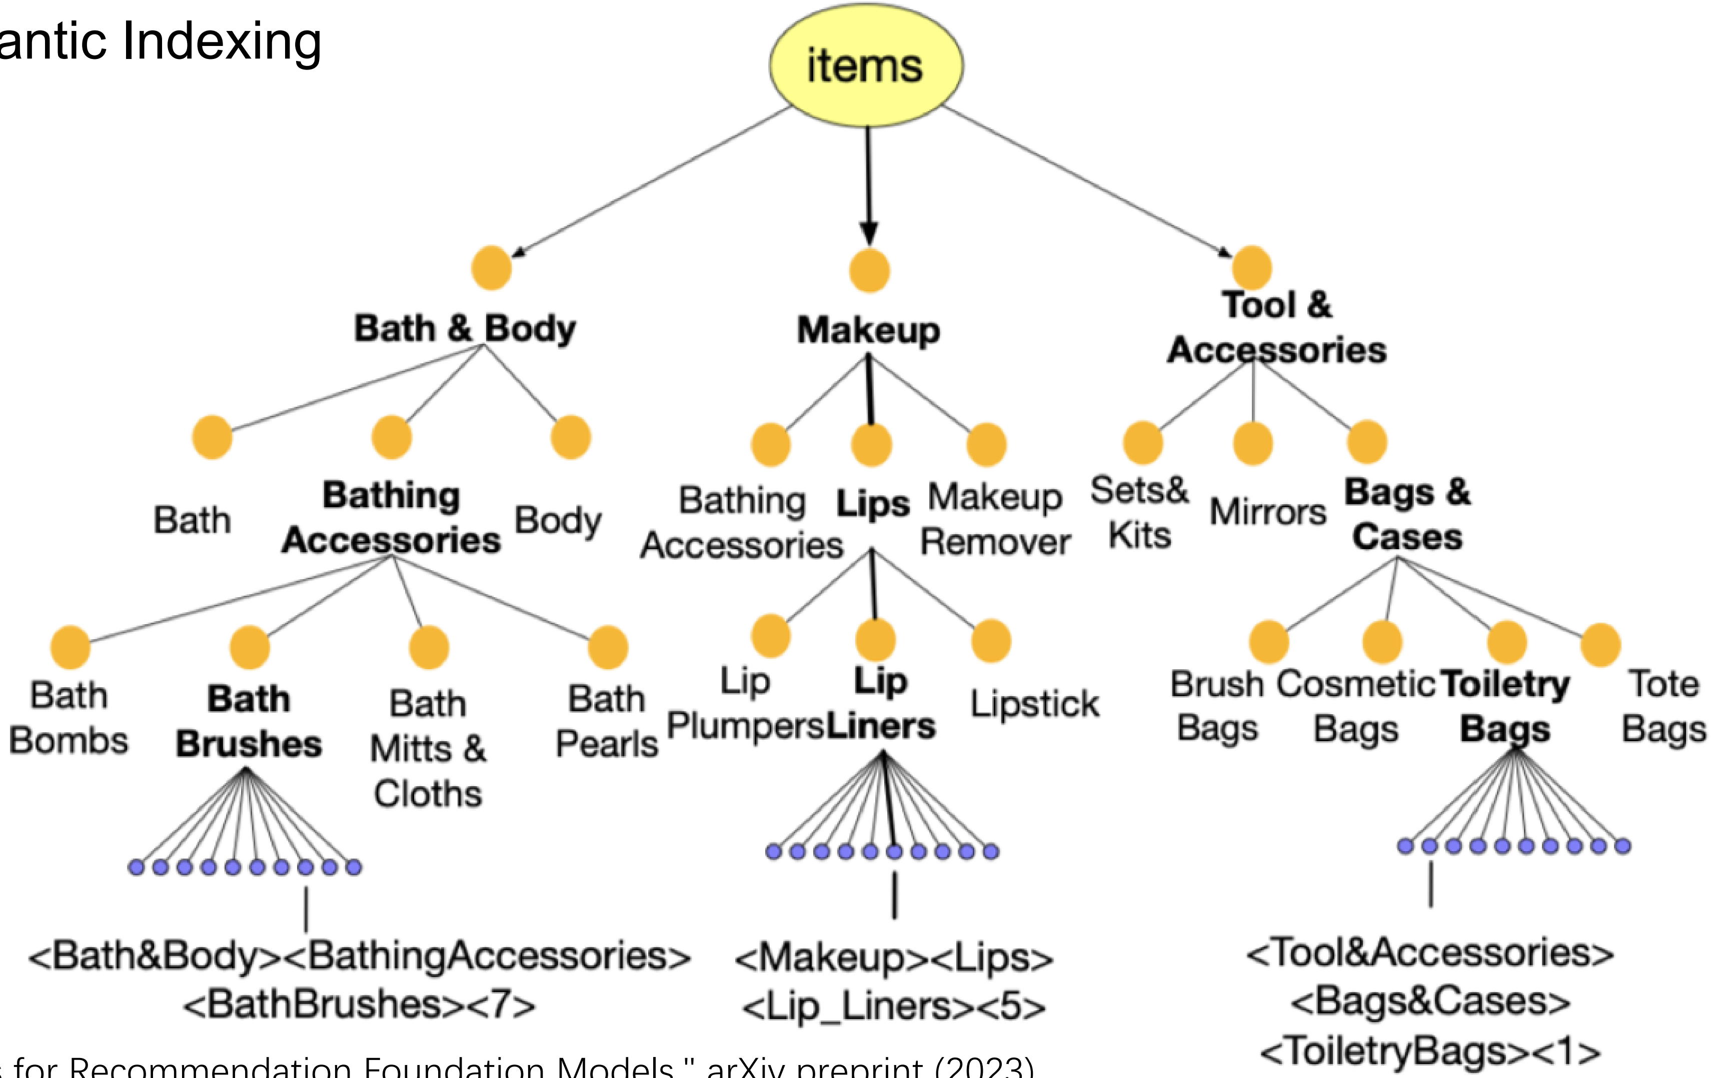

"How to Index Item IDs for Recommendation Foundation Models." arXiv preprint (2023).

ID-based LLM RecSys

- Modelling user interaction history with Markov chain

User_4782 has bought
3472, 7653, 0192, 4271.
What will he buy next?

Question

User_4782 will buy 7251.

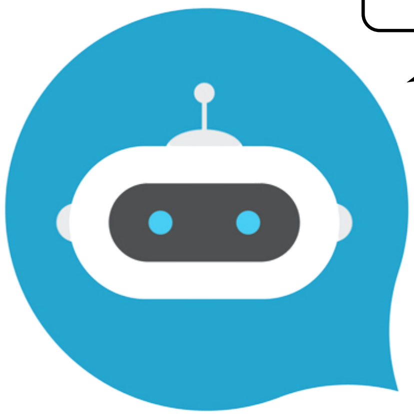

RecSys

ID-based LLM RecSys

- Modelling user interaction history with Markov chain

User_0513 bought 3472, 1784, 2563, 8892, 4271, 0988

User_4781 bought 7562, 8892, 4057, 0192, 7251

……

User_2788 bought 4271, 7251, 9082, 0192, 5672

Pre-training & Fine-tuning

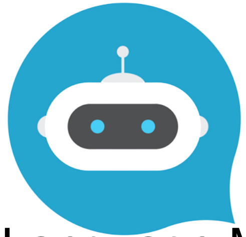

Large Language Models

$$v_{i+1} = \arg\max_{v} P(v_{i+1}|v_1, v_2, \ldots, v_i)$$

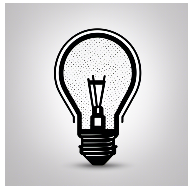

User_4782 will buy 7251.

Modelling the probability of the next item

ID-based LLM RecSys

- The N-gram probability in NLP
  - Unigram

$$P(\text{“3472”}) = \frac{1}{16} \qquad P(\text{“1784”}) = \frac{1}{16}$$

$$P(\text{“2563”}) = \frac{1}{16} \qquad \boxed{P(\text{“8892”}) = \frac{2}{16}}$$

$$\boxed{P(\text{“4271”}) = \frac{2}{16}} \qquad P(\text{“0988”}) = \frac{1}{16}$$

$$P(\text{“7562”}) = \frac{1}{16} \qquad P(\text{“4057”}) = \frac{1}{16}$$

$$\boxed{P(\text{“0192”}) = \frac{2}{16}} \qquad \boxed{P(\text{“7251”}) = \frac{2}{16}}$$

$$P(\text{“9082”}) = \frac{1}{16} \qquad P(\text{“5672”}) = \frac{1}{16}$$

User_0513 bought 3472, 1784, 2563, 8892, 4271, 0988

User_4781 bought 7562, 8892, 4057, 0192, 7251

……

User_2788 bought 4271, 7251, 9082, 0192, 5672

User_4782 has bought 3472, 7653, 0192, 4271. What will he buy next?

Question

ID-based LLM RecSys

- The N-gram probability in NLP
  - Bigram

$$P(\text{“0988”} \mid \text{“4271”}) = \frac{1}{2}$$

$$P(\text{“7251”} \mid \text{“4271”}) = \frac{1}{2}$$

  - Which one to choose?

$$P(\text{“0988”}) = \frac{1}{16}$$

$$P(\text{“7251”}) = \frac{2}{16}$$

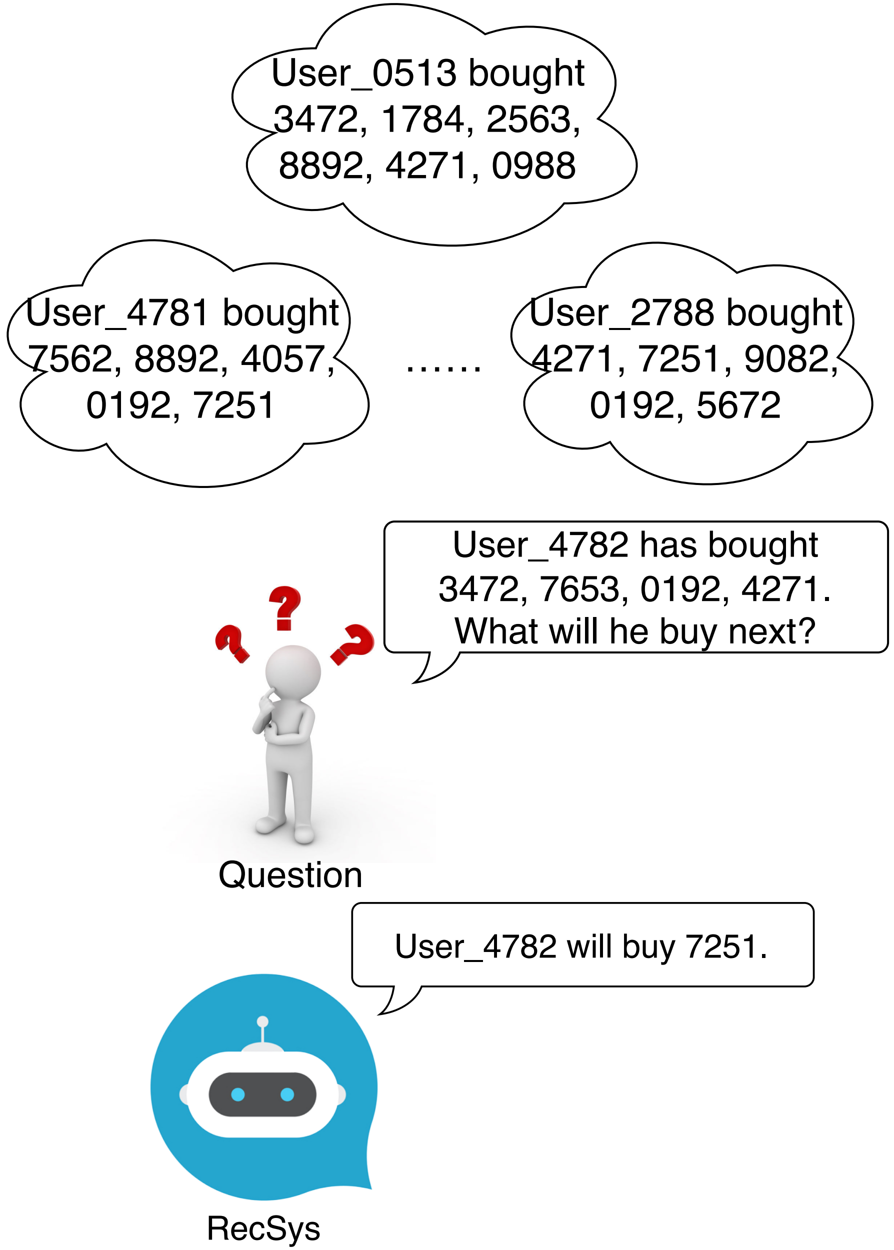

ID-based LLM RecSys

- The N-gram probability in NLP
  - The co-occurrence of item IDs
  - User_0513 bought 3472, …, 4271, 0988
  - User_4782 bought 3472, …, 4271, ?

  - Is “0988” a better answer than “7251”?

User_0513 bought 3472, 1784, 2563, 8892, 4271, 0988

User_4781 bought 7562, 8892, 4057, 0192, 7251

……

User_2788 bought 4271, 7251, 9082, 0192, 5672

User_4782 has bought 3472, 7653, 0192, 4271. What will he buy next?

Question

ID-based LLM RecSys

- Contextual representations of words in LLMs
  - User_0513 bought 3472, 1784, 2563, 8892, 4271, 0988
  - User_4782 bought 3472, 7653, 0192, 4271, ?
  - The item representations can vary for different contexts

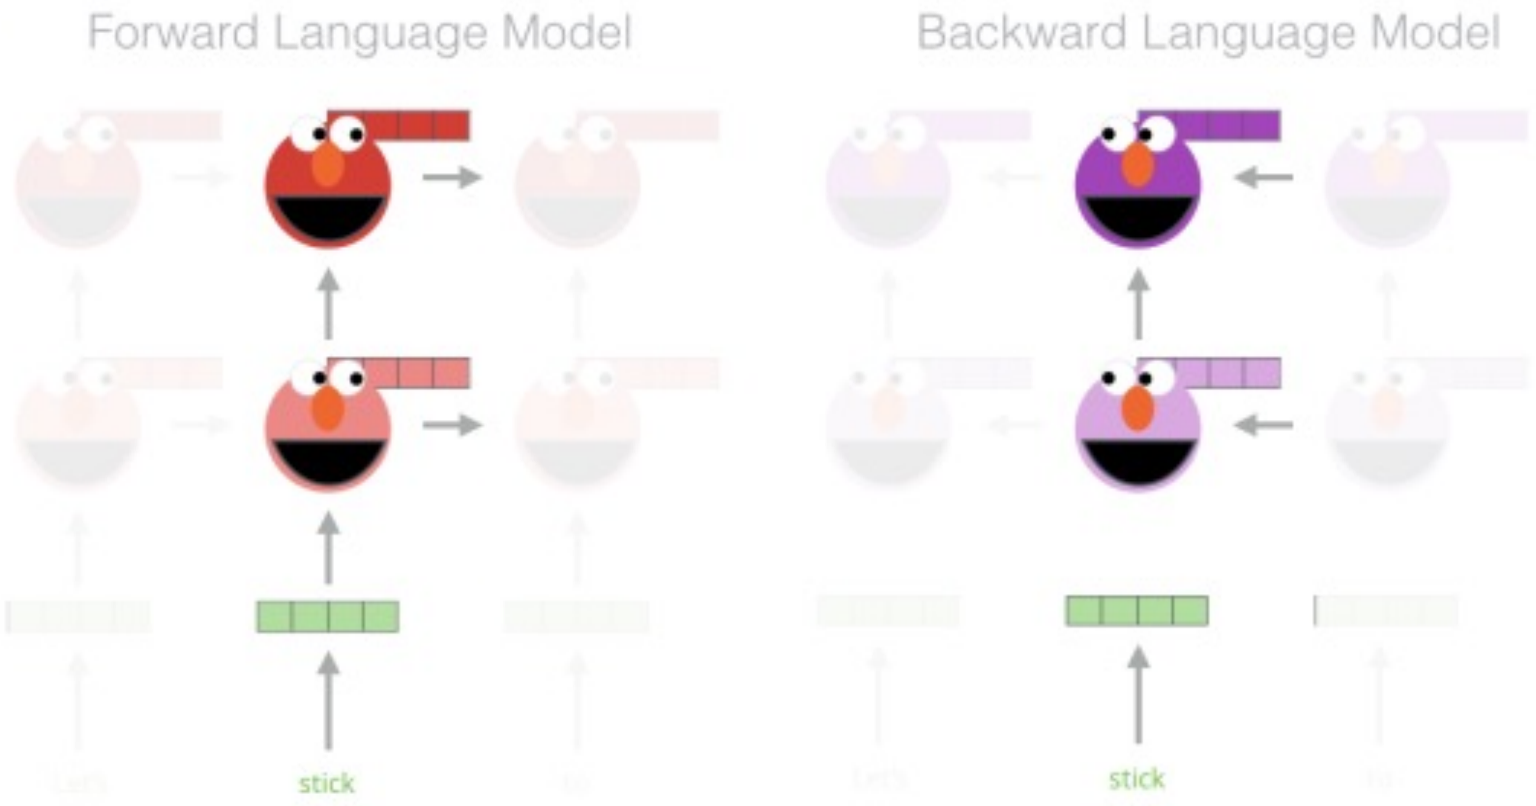

User & Item Representation in LLMs

ID-based LLM RecSys

Text-based LLM RecSys

Text-based LLM RecSys

❑ GPT4Rec

- ❖ Item title contains rich semantic information
- ❖ It's a natural way to use text to describe items

Previously, the customer has bought

Ben Nye Banana Luxury Face Powder 3.0 oz Makeup Kim Kardashian NEW!!!. Rosallini Women Stainless Steel Extension Eyelash Applicator Tool Fish Tail Clip. Beauty Flawless Makeup Blender Sponge Puff (size 1). Fruit Of The Earth 100% Aloe Vera 24oz Gel Pump.

In the future, the customer wants to buy

Fine-tuned GPT-2

Ben Nye Luxury Powders - Banana 1.5oz.
Beautyblender Solid Blendercleanser 1 oz.
Professional 15 Color Concealer Camouflage Makeup Palette.
Pro Beauty Makeup Sponge Blender Flawless Smooth Shaped Water Droplets Puff (Random Color).
L'Oreal Paris True Match Super Blendable Makeup, Natural Buff, 1.0 Ounces.

"GPT4Rec: A generative framework for personalized recommendation and user interests interpretation." arXiv preprint (2023).

Text-based LLM RecSys

GPT4Rec

- In the era of LLMs, **Retrieval-Augmented Generation** (**RAG**) could be a way to improve the capability of LLMs
- RAG also enhances the explainability of LLM-based RecSys

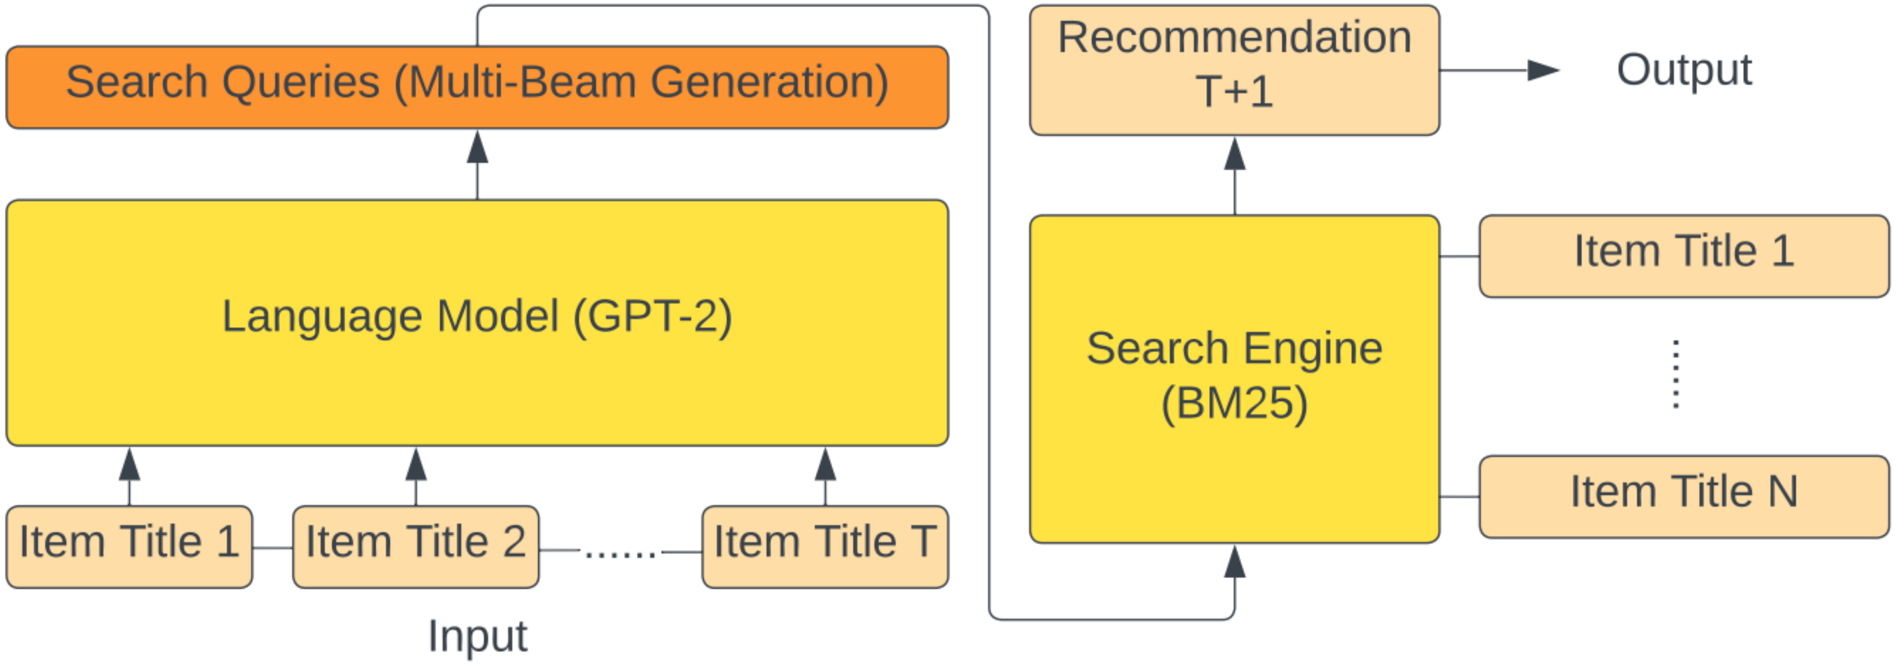

"GPT4Rec: A generative framework for personalized recommendation and user interests interpretation." arXiv preprint (2023).

Text-based LLM RecSys

TF-DCon

- Content-level condensation for recommendation
- Condense Item title and description to refine item representation

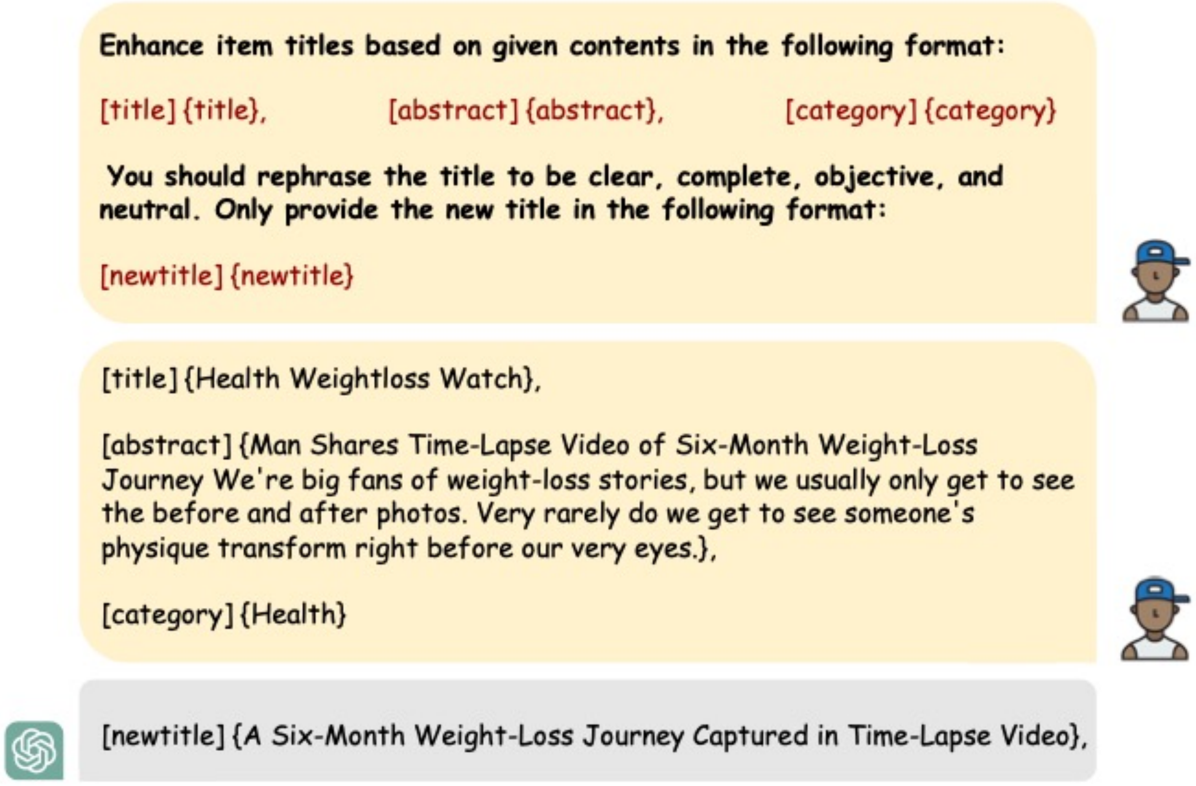

Content-Level Condensation

"Leveraging Large Language Models (LLMs) to Empower Training-Free Dataset Condensation for Content-Based Recommendation." arXiv preprint (2023).

Overview

Presenter:
Jiatong Li

User & Item Representation

- ID-based LLM RecSys
- Text-based LLM RecSys

Pre-training

- Pre-training in NLP
- Pre-training LLM-based RecSys

Zoom ID:
91649466943
Password:
202312

Website QR Code

Pre-training

Pre-training in NLP

Pre-training LLM-based RecSys

Pre-training in NLP

- What is pre-training?
  - Core Idea: knowledge transfer
  - Technically:

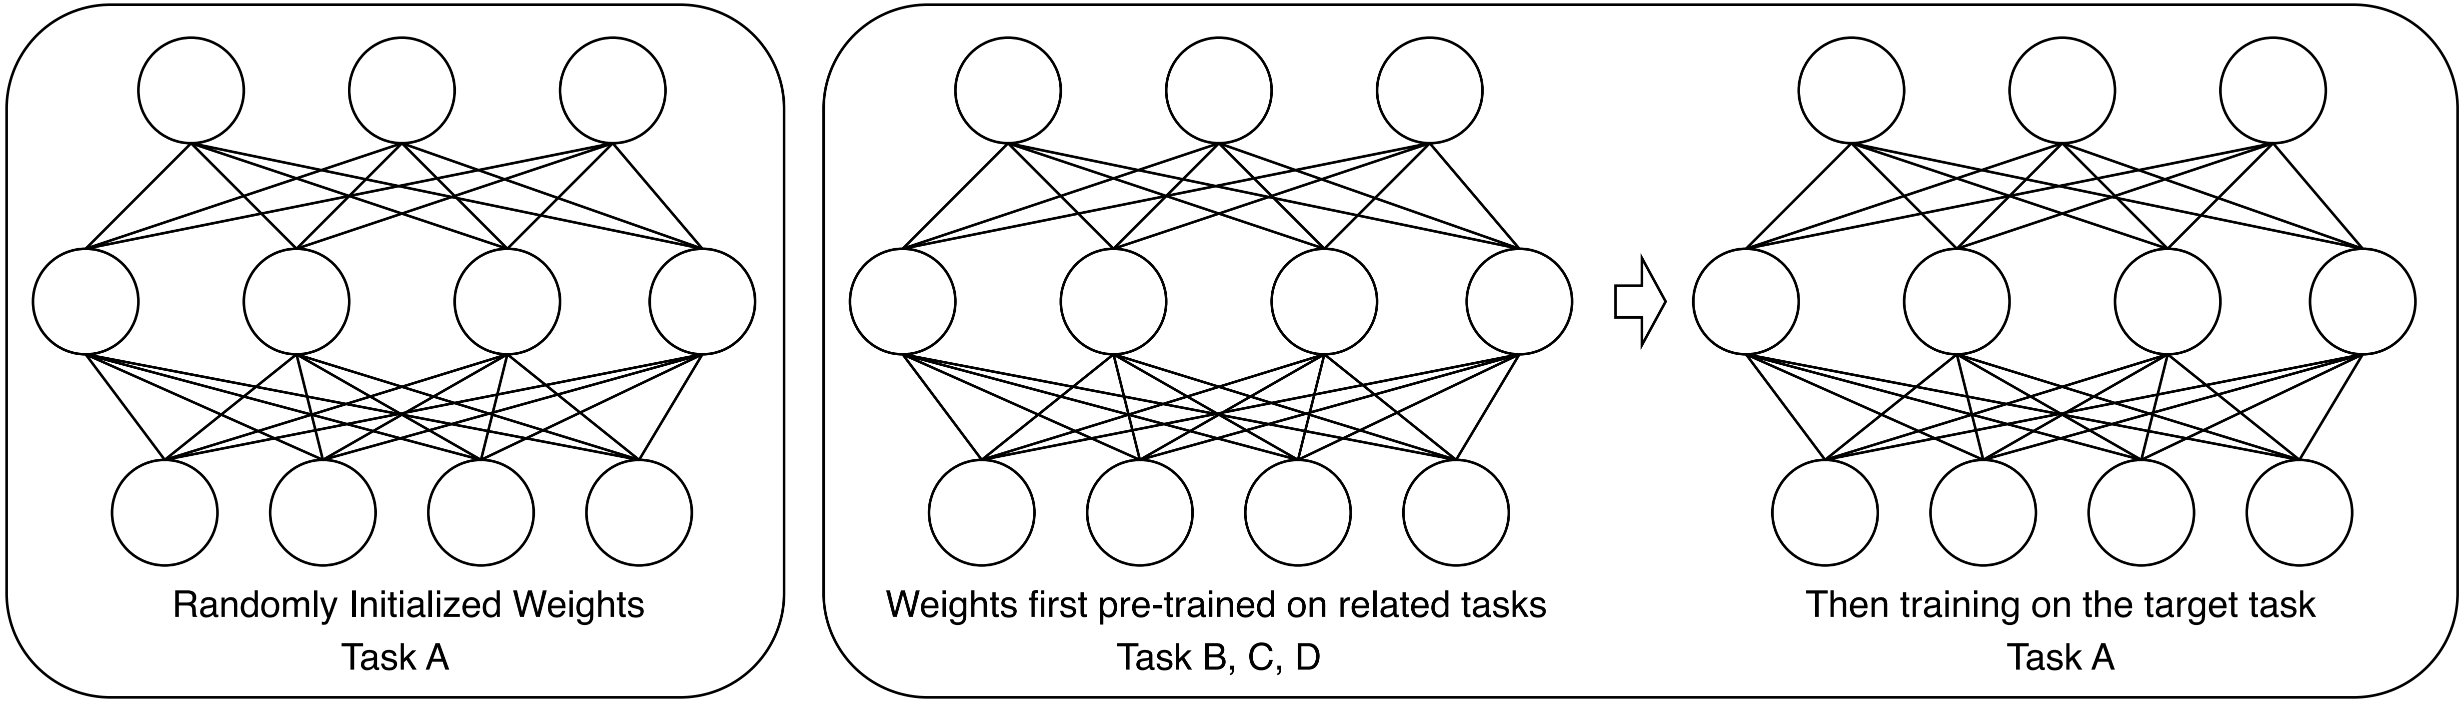

Pre-training in NLP

- Why pre-training?
  - Recall: Tokenization

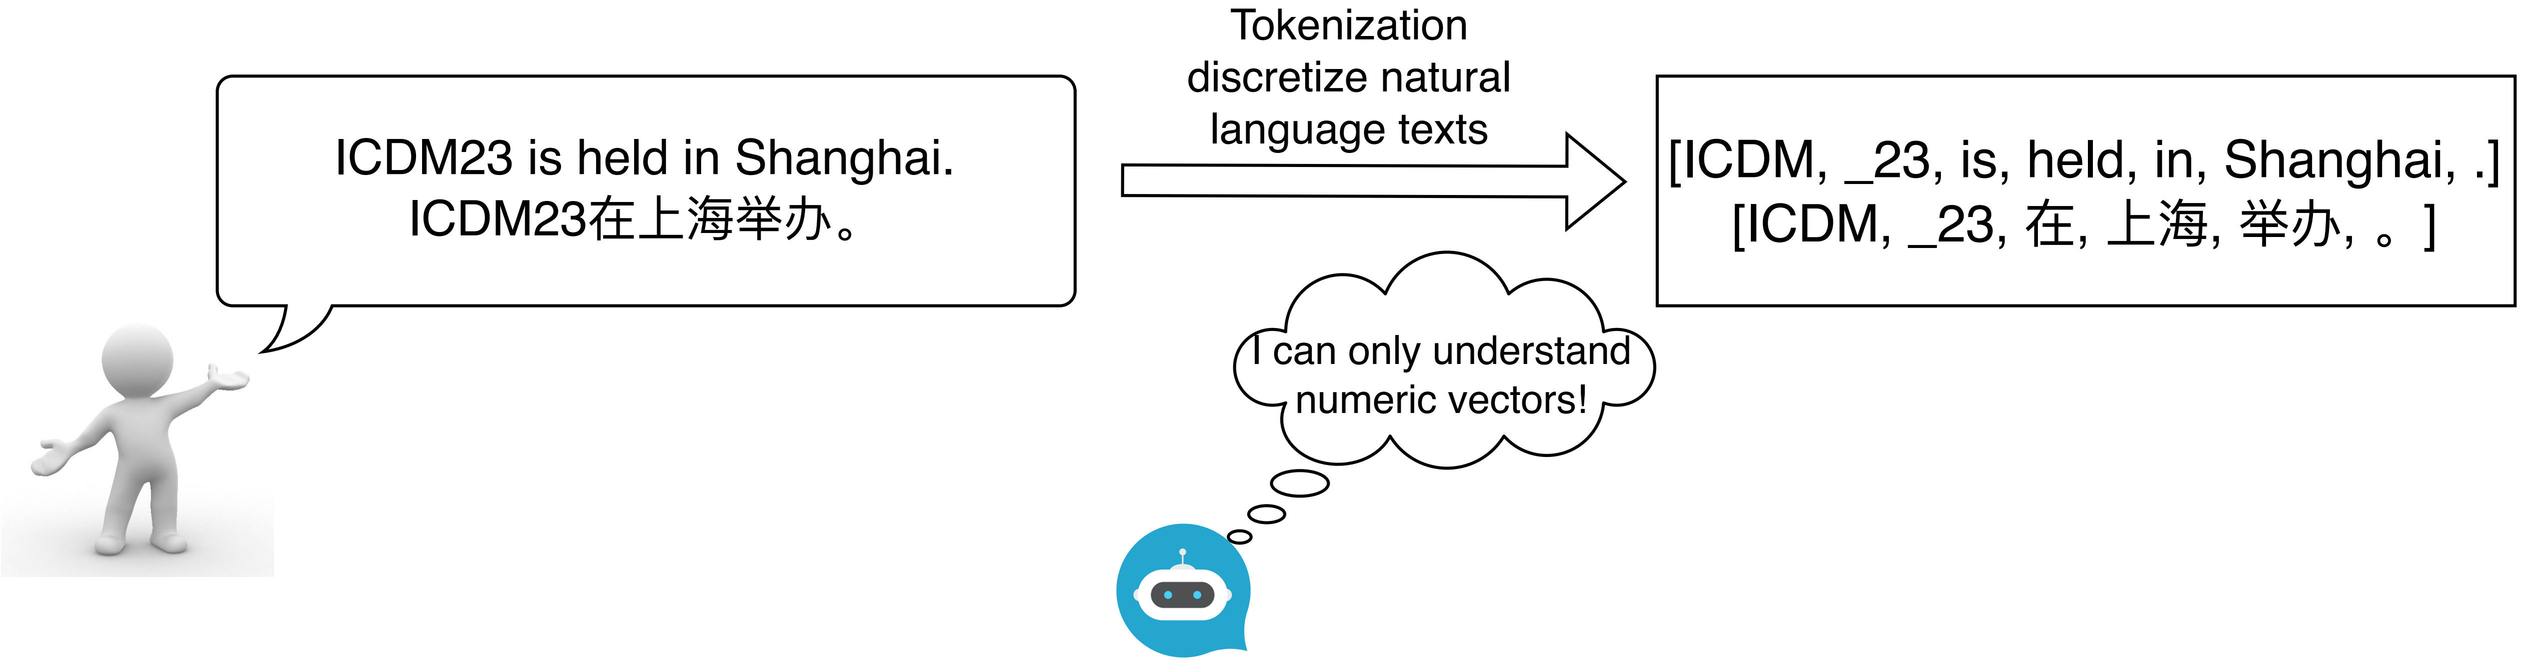

Pre-training in NLP

- Why pre-training?
  - Recall: Tokenization

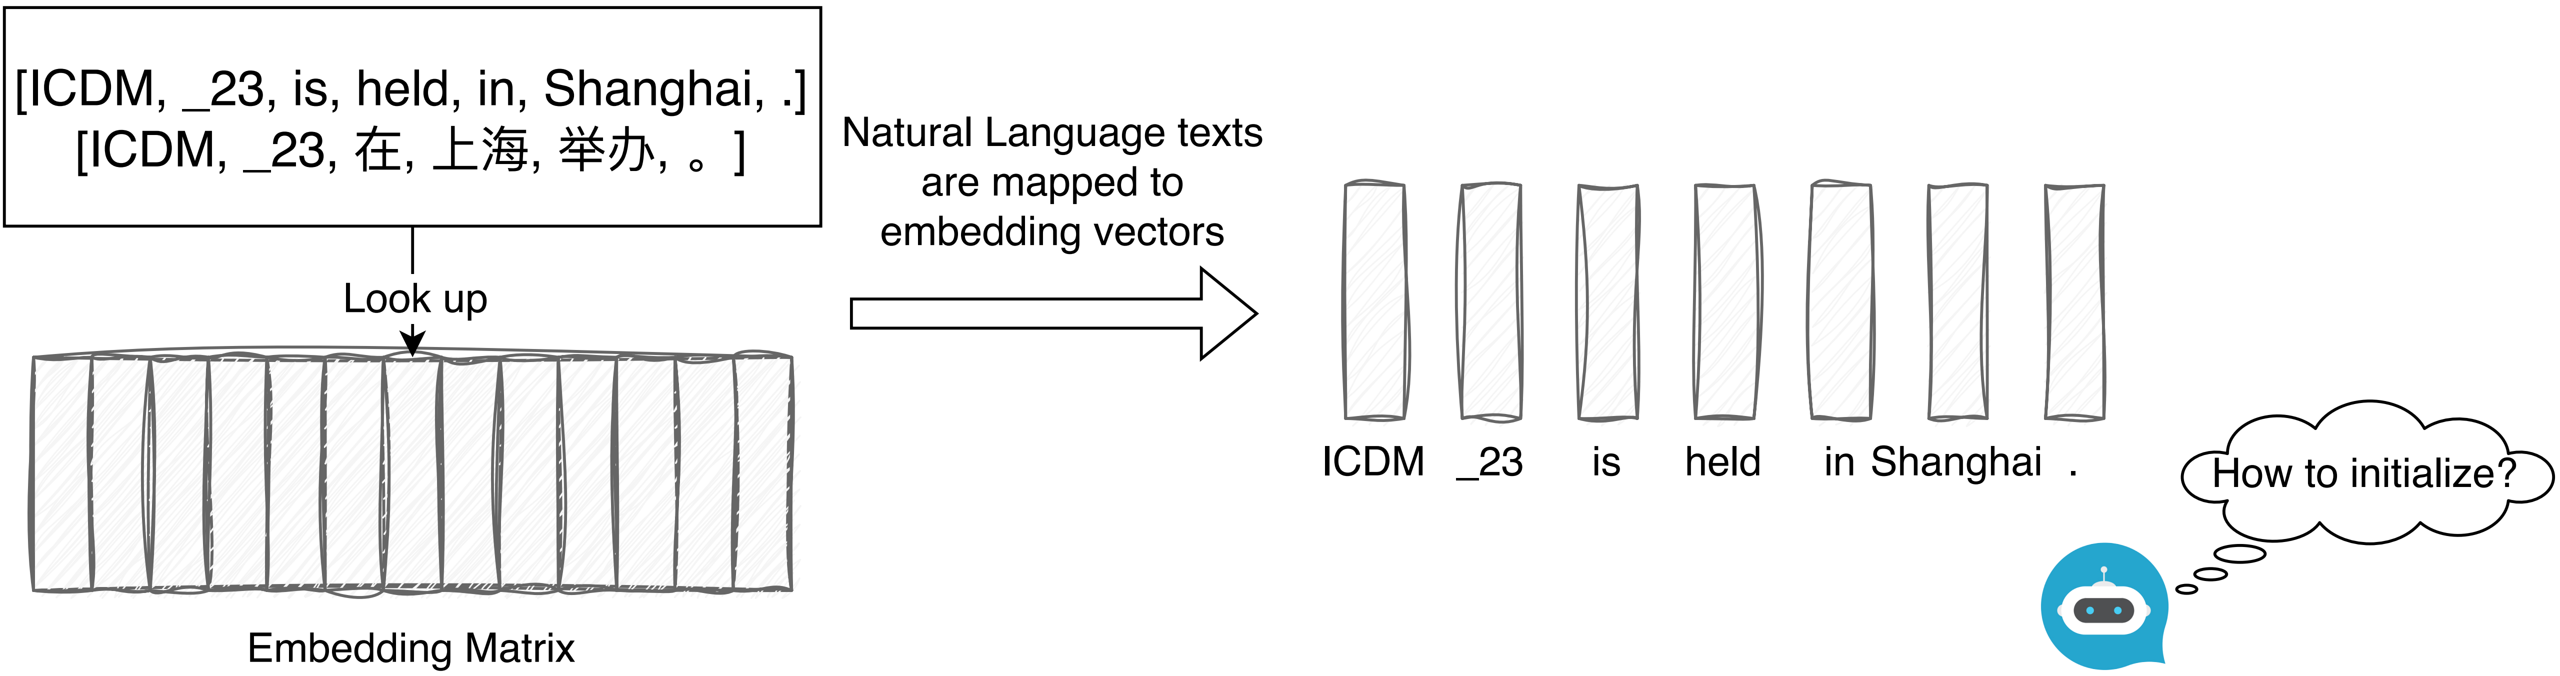

Pre-training in NLP

- Word embeddings?

  - king: [-0.5, -0.9, 1.4, …]
  - queen: [-0.6, -0.8, -0.2, …]
  - woman: [-0.1, -0.1, -1.6, …]

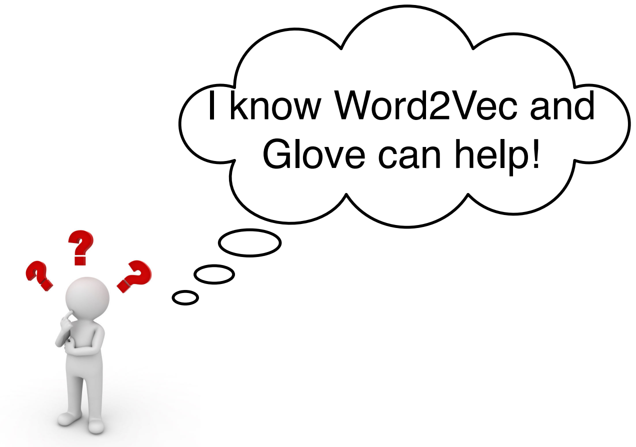

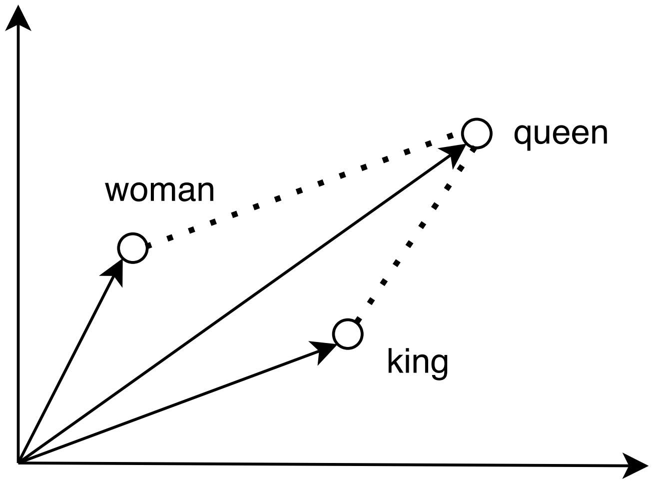

- Static word embeddings (word2vec, Glove) are pre-trained on text corpus from co-occurrence statistics

  - He is the king of the country

Inner Product

  - She is the queen of the country

Pre-training in NLP

- Problem of static word embedding – Context-Free

He is in charge of the lab.
He and I charge the phone.

Different meanings

- How to solve it? – Contextual representations
  - He is in charge of the lab
    - charge: [0.2, 0.8, 1.4, …]
  - He and I charge the phone
    - charge: [-0.3, -0.4, 0.7, …]

Pre-training in NLP

- Semi-Supervised Sequence Learning

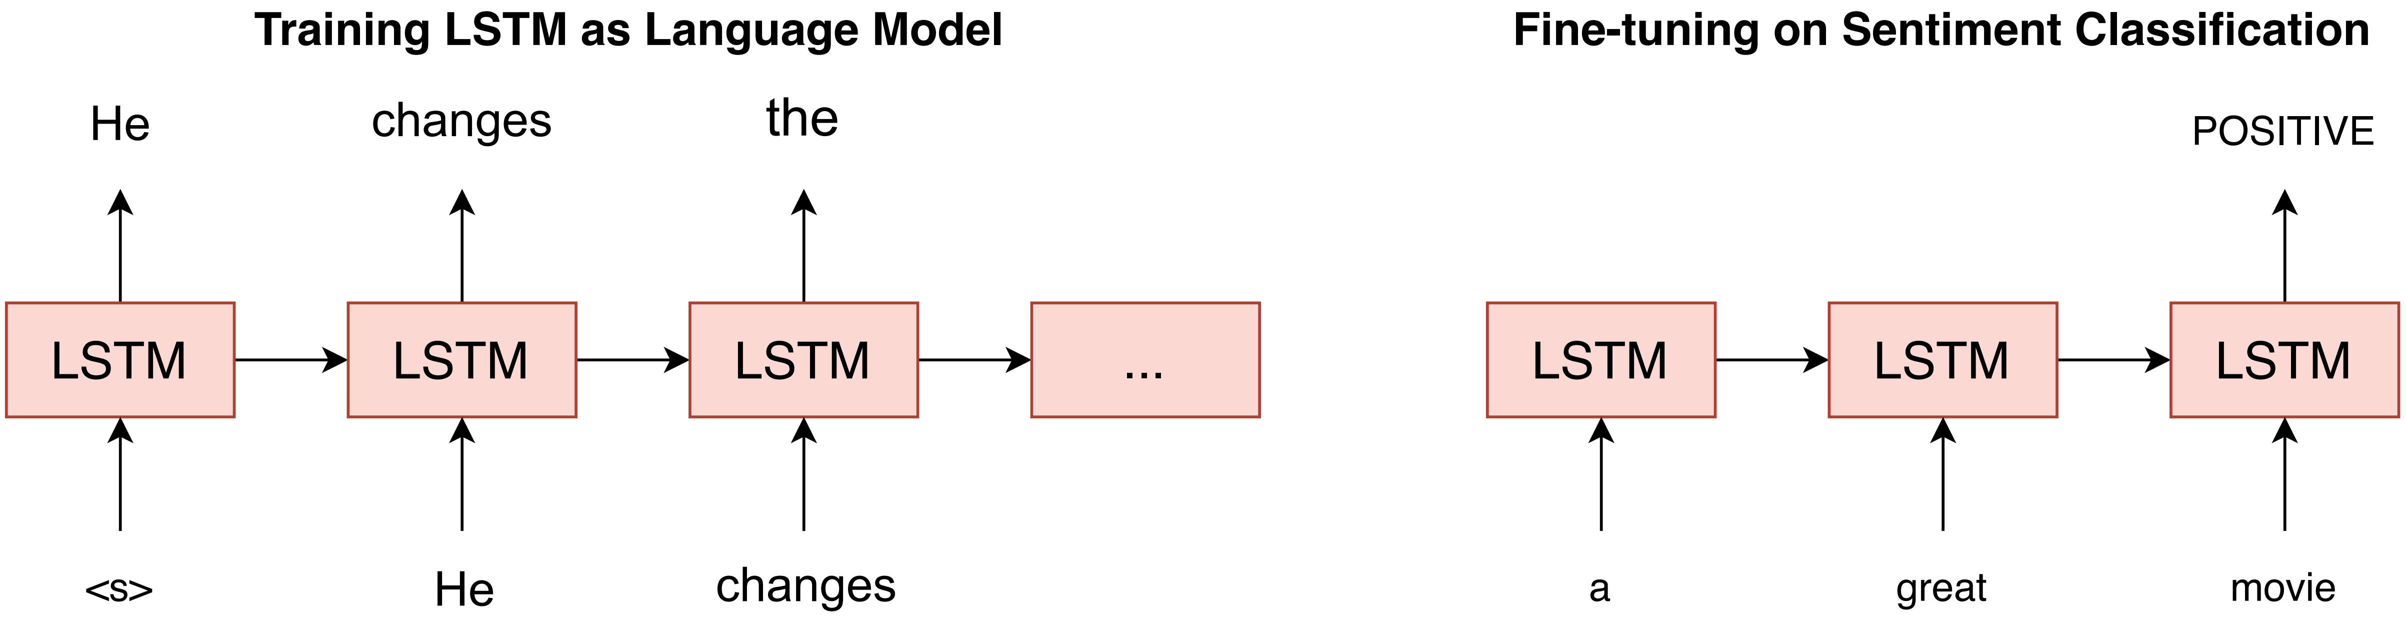

"Semi-supervised sequence learning." Advances in neural information processing systems 28 (2015).

Pre-training in NLP

- ELMo: Deep Contextual Word Embeddings

Training Separate Left-to-Right and Right-to-Left Language Models

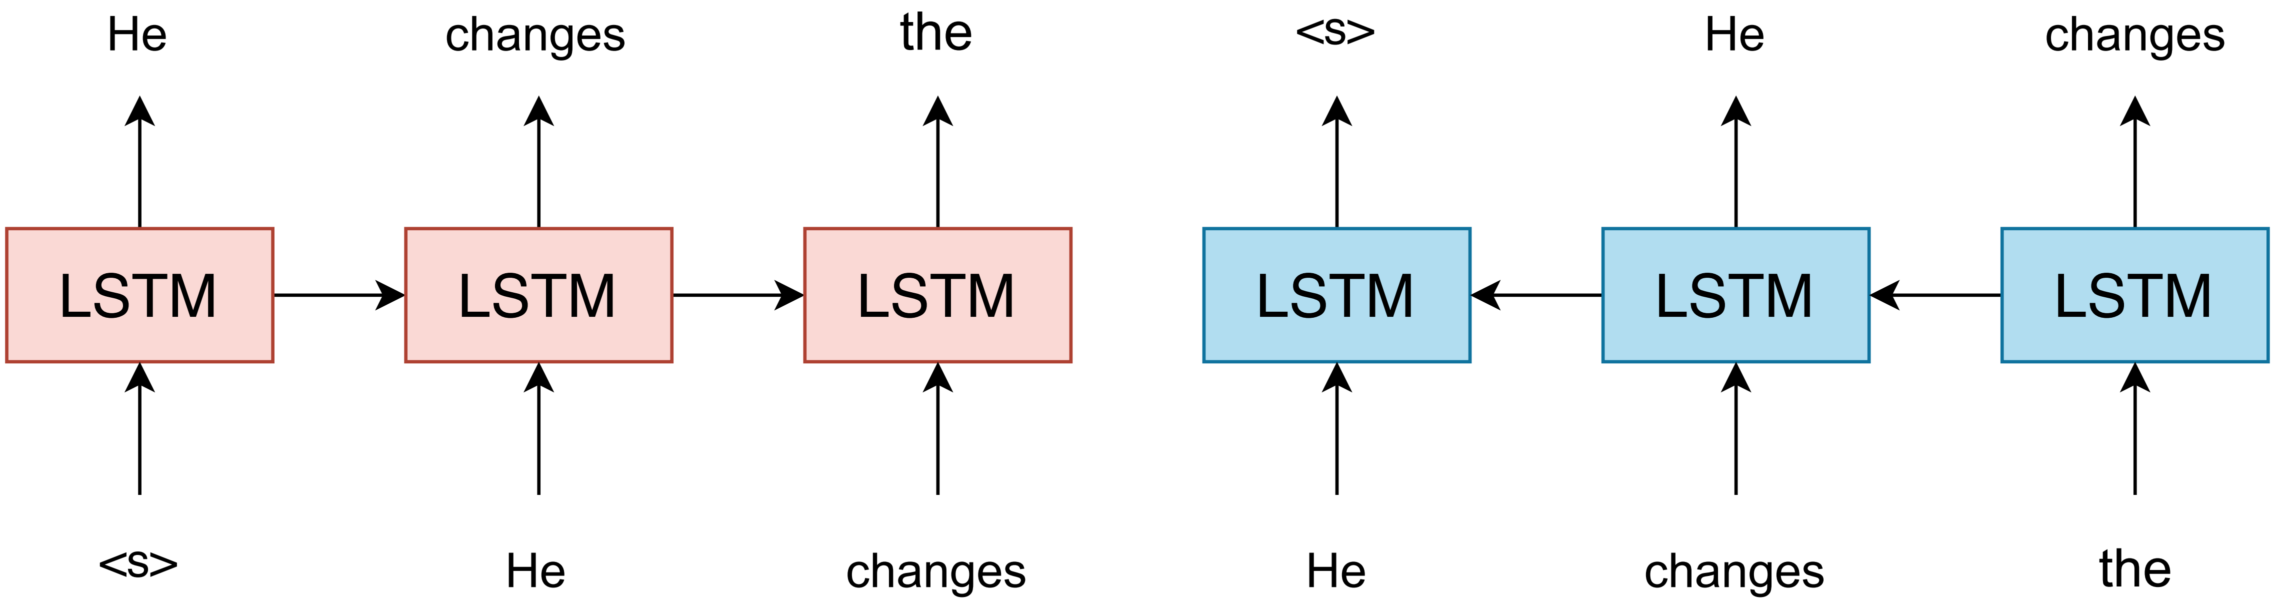

"Deep Contextualized Word Representations." NAACL (2018).

Pre-training in NLP

- Most Favored Pre-training Tasks in NLP
  - Design specific pre-training tasks that could introduce knowledge
    - Masked Language Modelling (For Encoder-Decoder and Encoder-only Structures)
    - Next Token Prediction (For Decoder-only Structures)

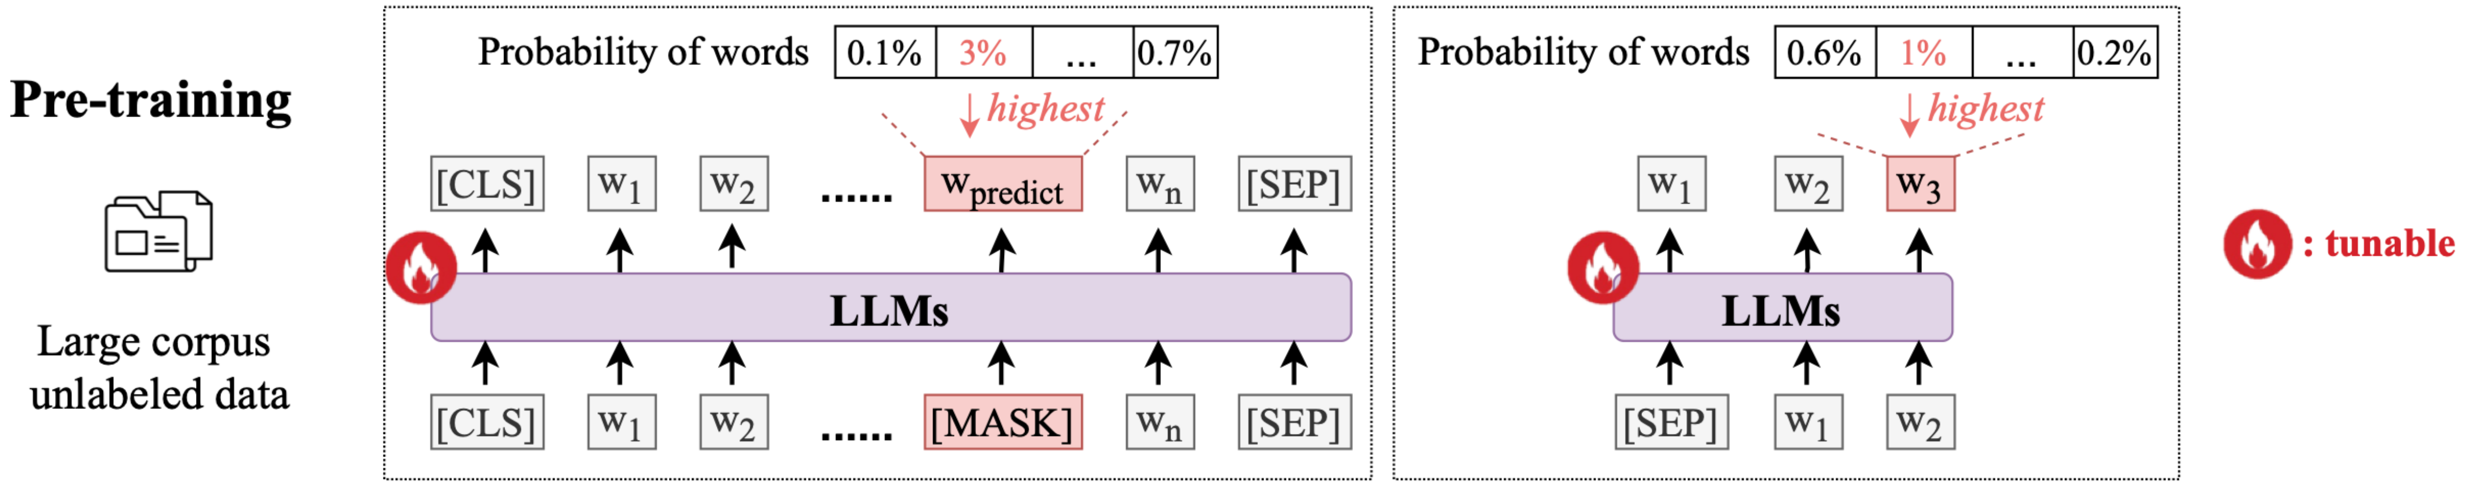

(Masked Language Modeling)

(Next Token Prediction)

Pre-training in NLP

- BERT: Bidirectional Encoder Representations from Transformers

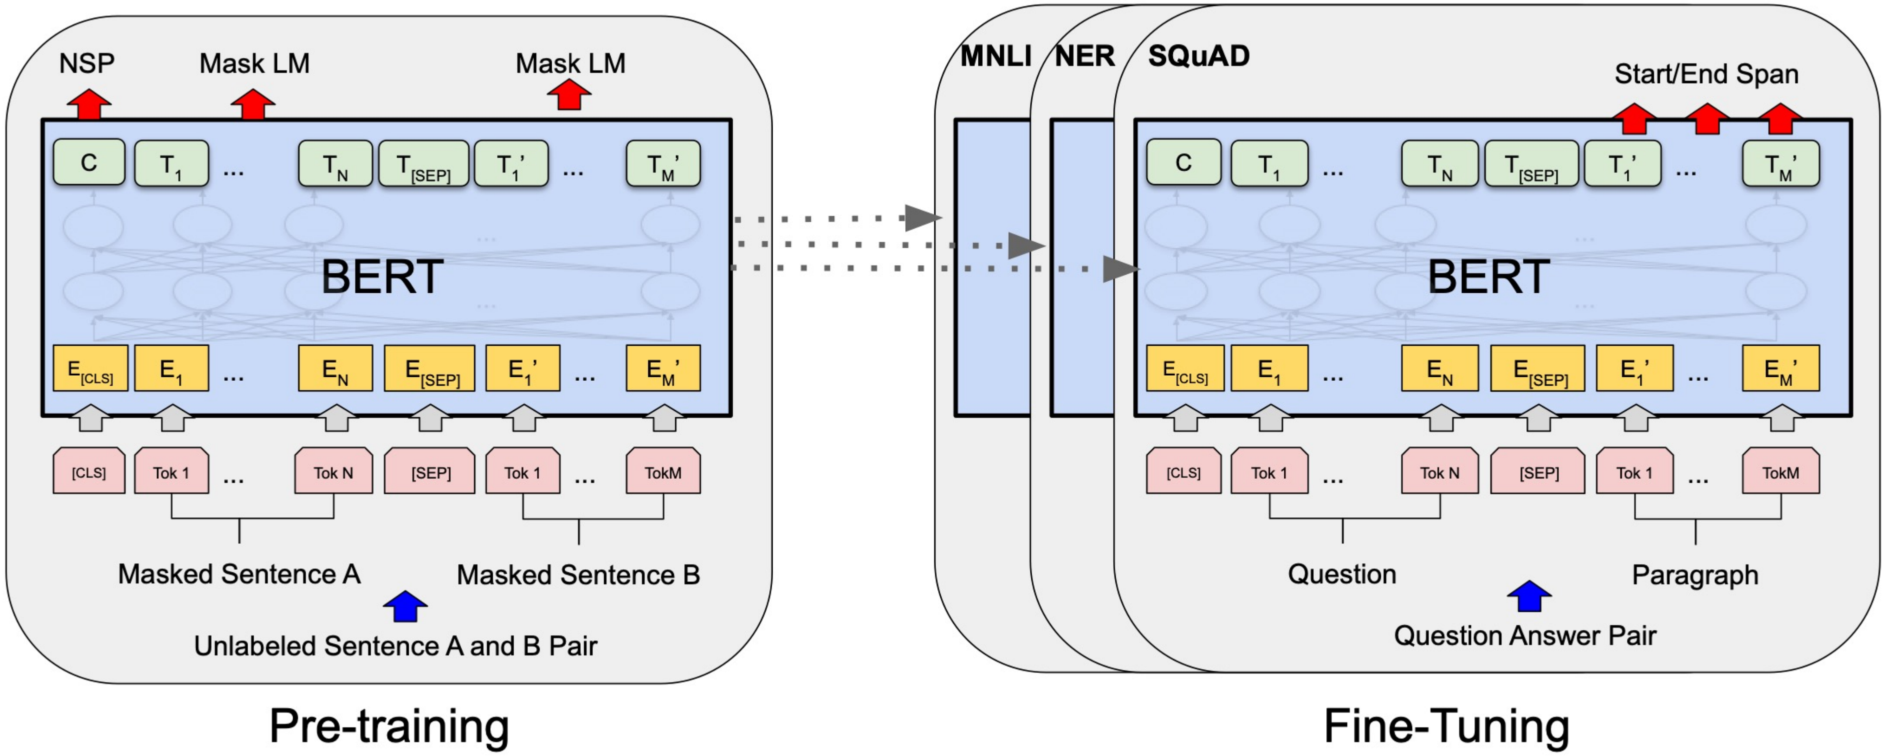

"Bert: Pre-training of deep bidirectional transformers for language understanding." NAACL (2019).

Pre-training in NLP

- GPT

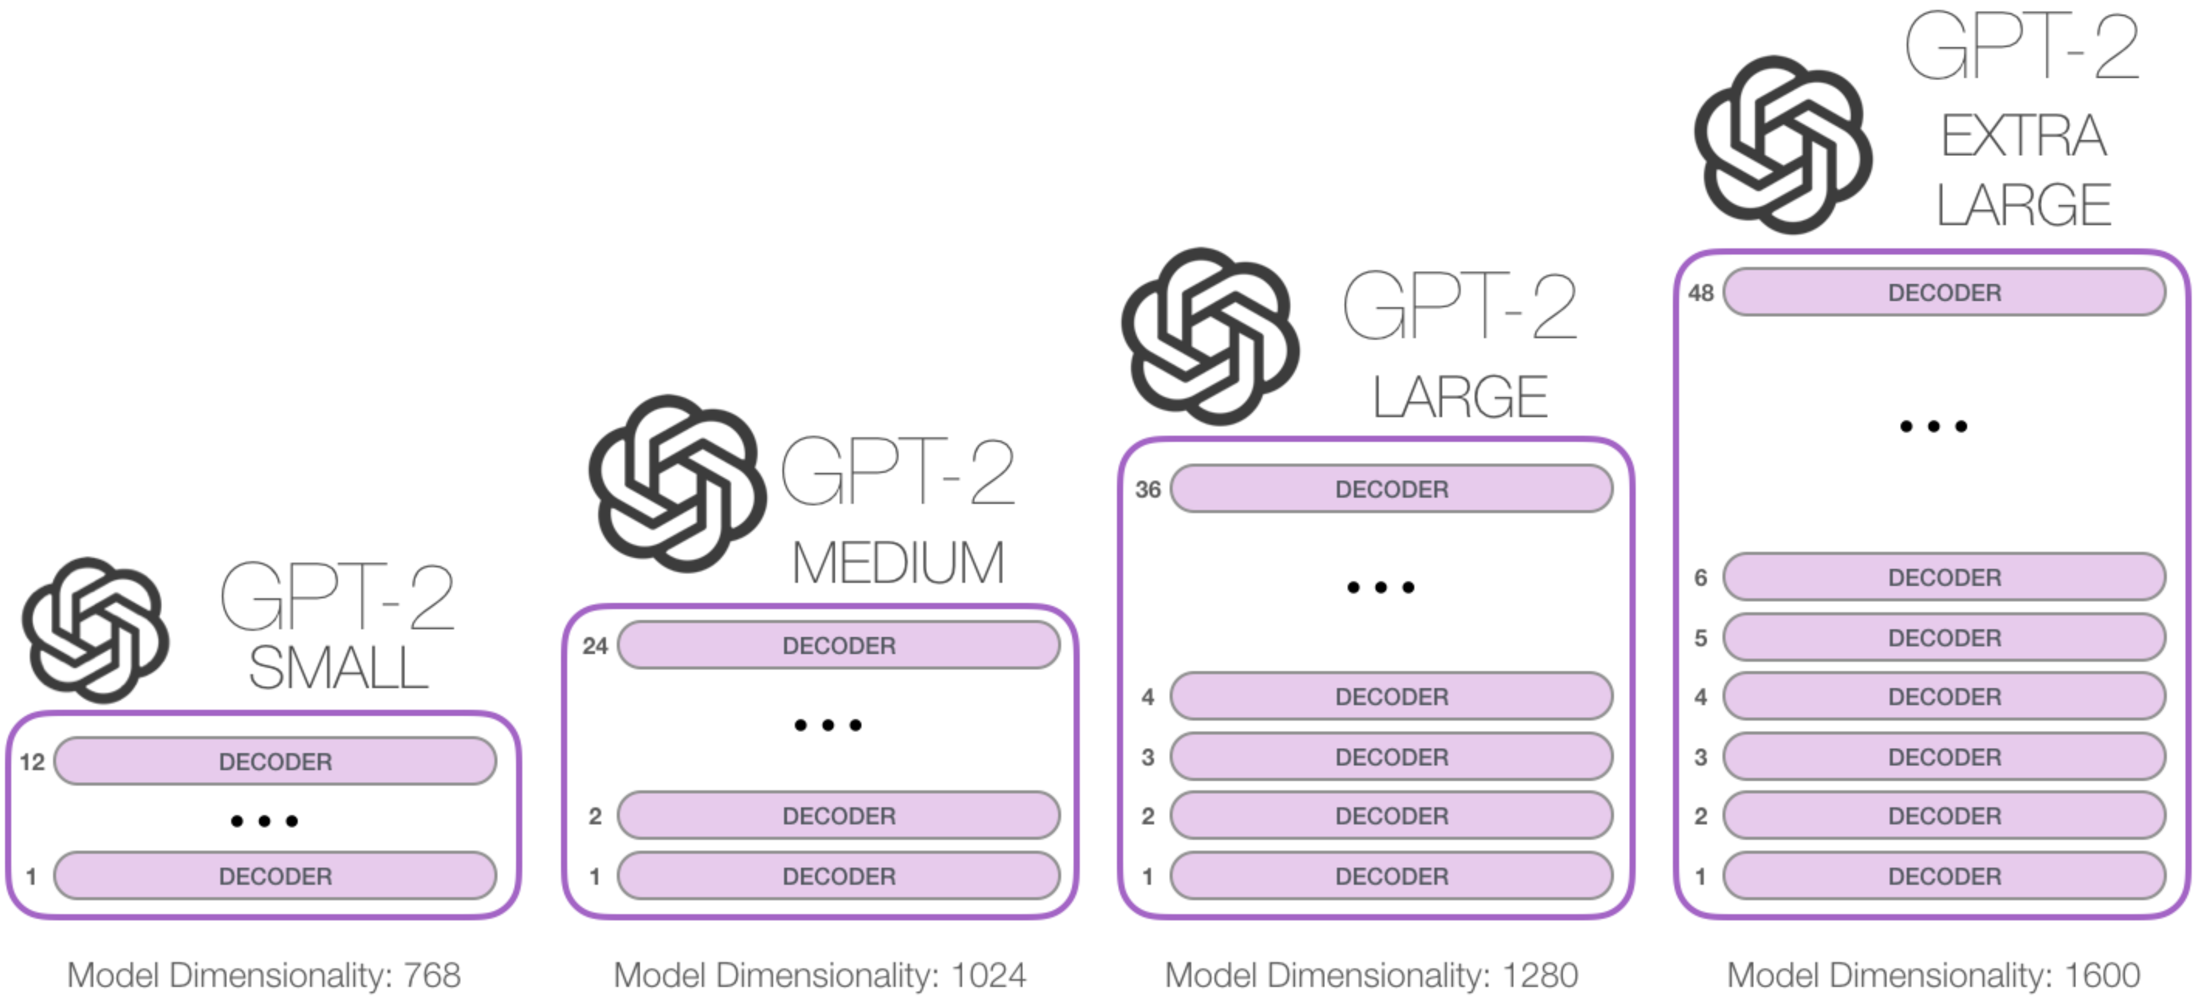

"Language models are unsupervised multitask learners." OpenAI blog 1.8 (2019).

Pre-training

Pre-training in NLP

Pre-training LLM-based RecSys

Pre-training LLM-based RecSys

❑ What is Pre-training in LLM-based RecSys and Why is it Necessary?

- General pre-training vs. domain-specific pre-training
- Domain knowledge is essential for relieving the knowledge gap

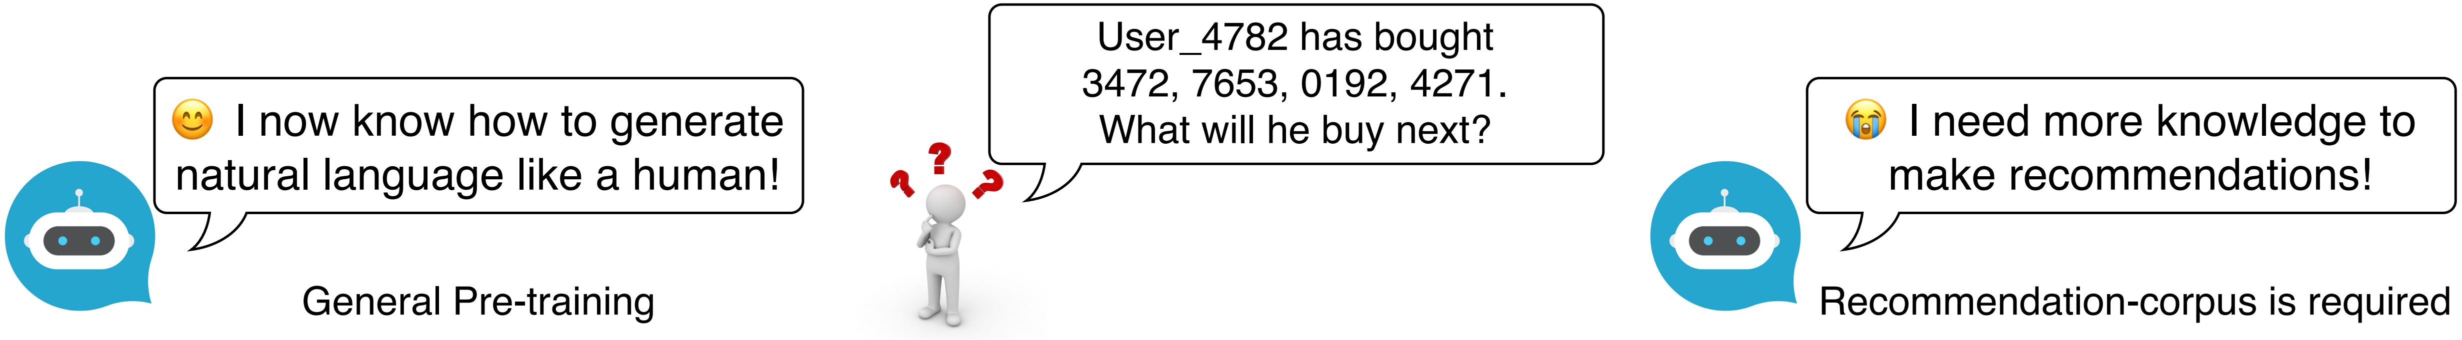

PTUM

- ❑ Masked Behavior Prediction (MBP)
- ❑ Next K Behaviors Prediction (NBP)

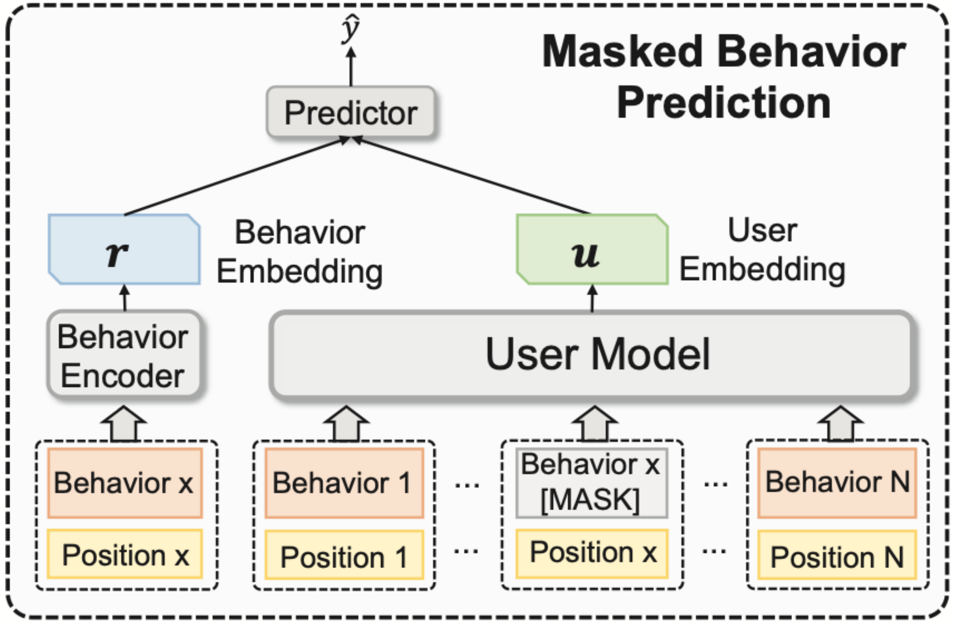

(a) Masked Behavior Prediction (MBP) task.

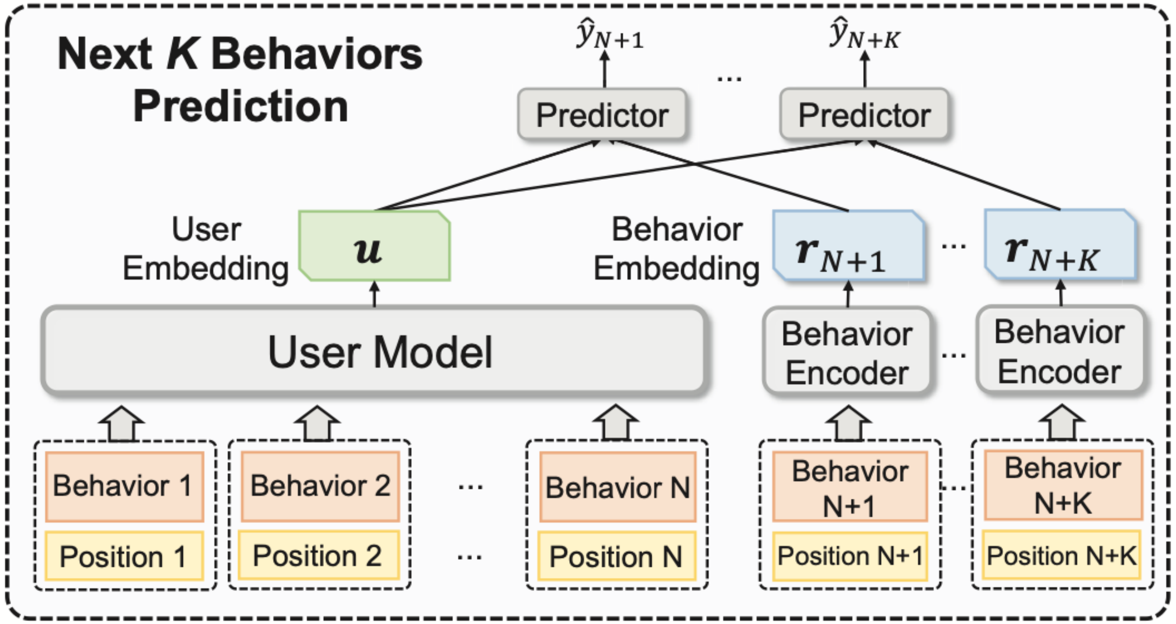

(b) Next K Behaviors Prediction (NBP) task.

$$\mathcal{L}_{MBP} = -\sum_{y \in \mathcal{S}_1} \sum_{i=1}^{P+1} y_i \log(\hat{y}_i).$$

$$\mathcal{L}_{NBP} = -\frac{1}{K} \sum_{y \in \mathcal{S}_2} \sum_{k=1}^{K} \sum_{i=1}^{P+1} y_{i,k} \log(\hat{y}_{i,k}).$$

$$\mathcal{L} = \mathcal{L}_{MBP} + \lambda \mathcal{L}_{NBP}.$$

"PTUM: Pre-training User Model from Unlabeled User Behaviors via Self-supervision." EMNLP (2020).

M6-Rec

- Text-infilling

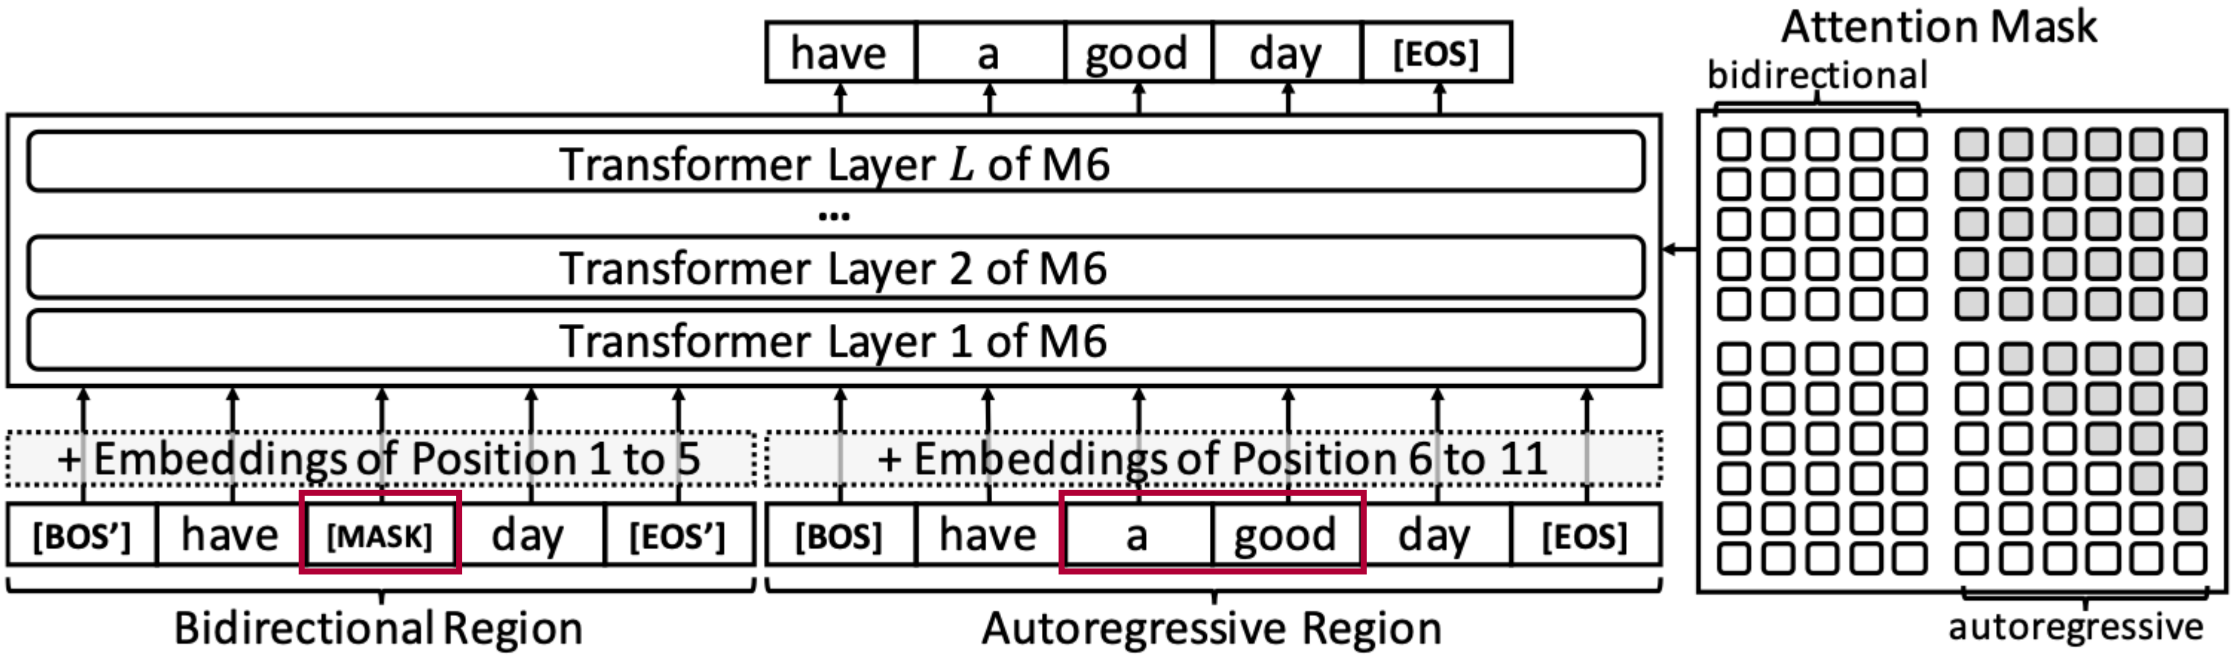

"M6-rec: Generative pretrained language models are open-ended recommender systems." arXiv preprint (2022).

P5

❑ Multi-task Pretraining with Personalized Prompt Collection

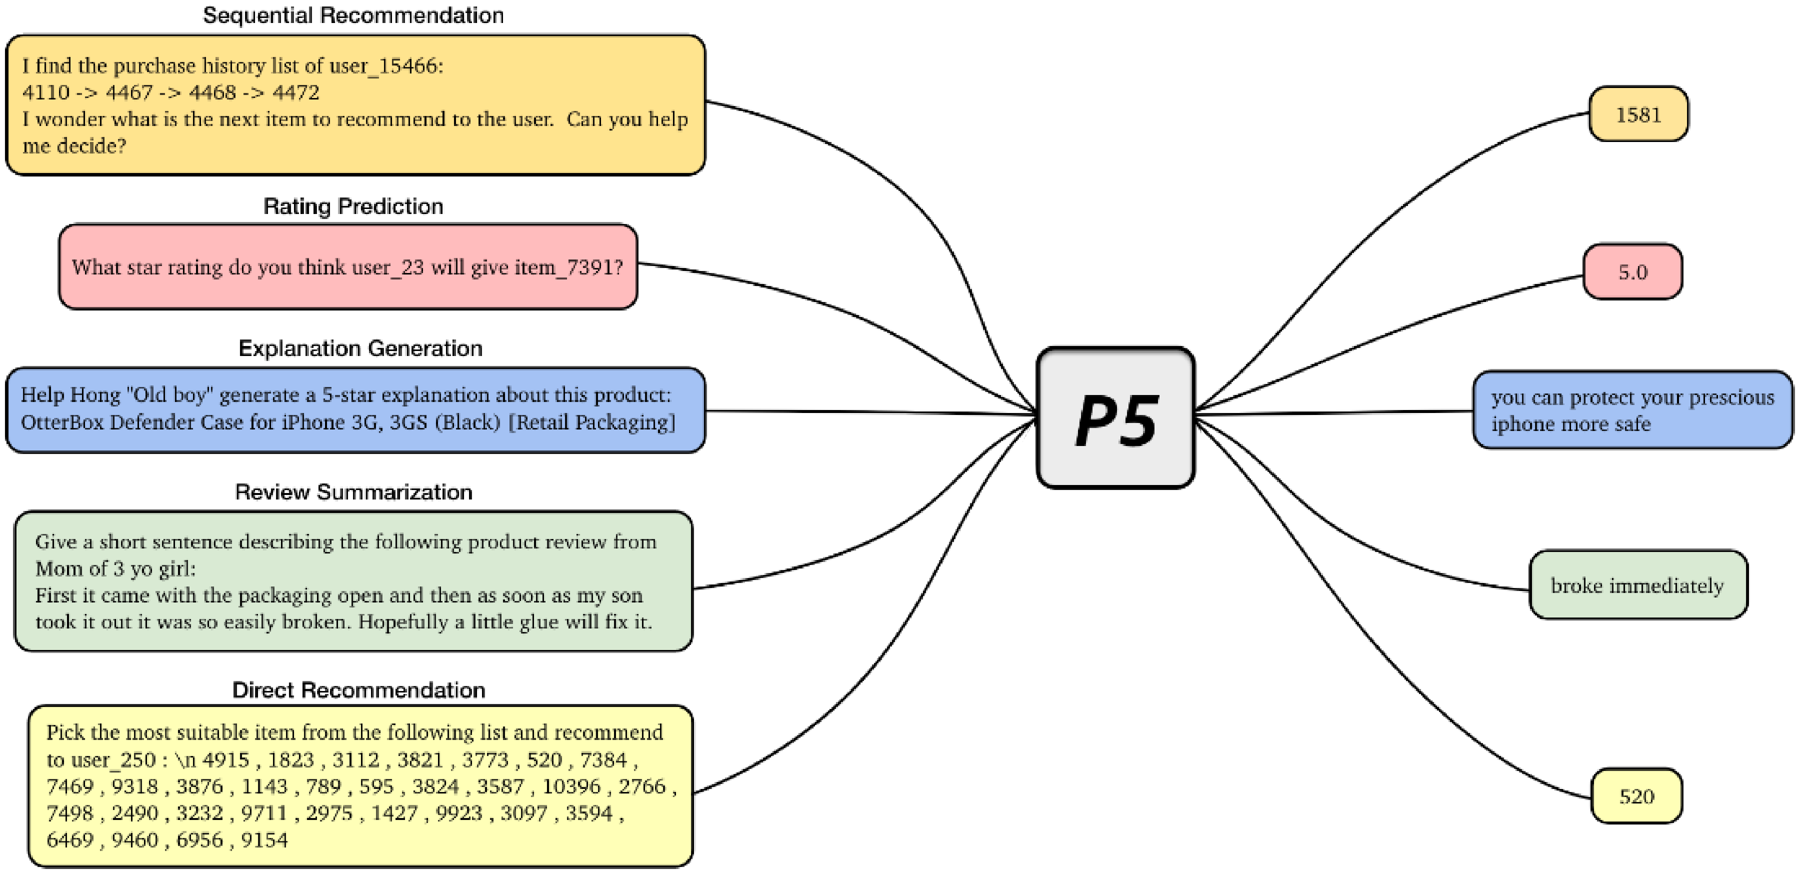

Multi-task Pretraining with Personalized Prompt Collection

"Recommendation as language processing (rlp): A unified pretrain, personalized prompt & predict paradigm (p5)." RecSys (2022).

Overview

Presenter:
Jiatong Li

Fine-tuning

- ⦿ Fine-tuning in NLP
- ○ Fine-tuning LLM-based RecSys

Parameter Efficient Fine-tuning

Zoom ID:
91649466943
Password:
202312

Website QR Code

Fine-tuning in NLP

- What is Fine-tuning and Why Fine-tuning?
  - Gaps between the pre-training tasks and downstream tasks still exist
    - Masked Language Modelling v.s. Sentiment Classification
  - Fine-tuning means training pre-trained LLMs on downstream tasks to fit the requirements
  - Supervised Fine-tuning (SFT) and Fine-tuning with Reinforcement Learning

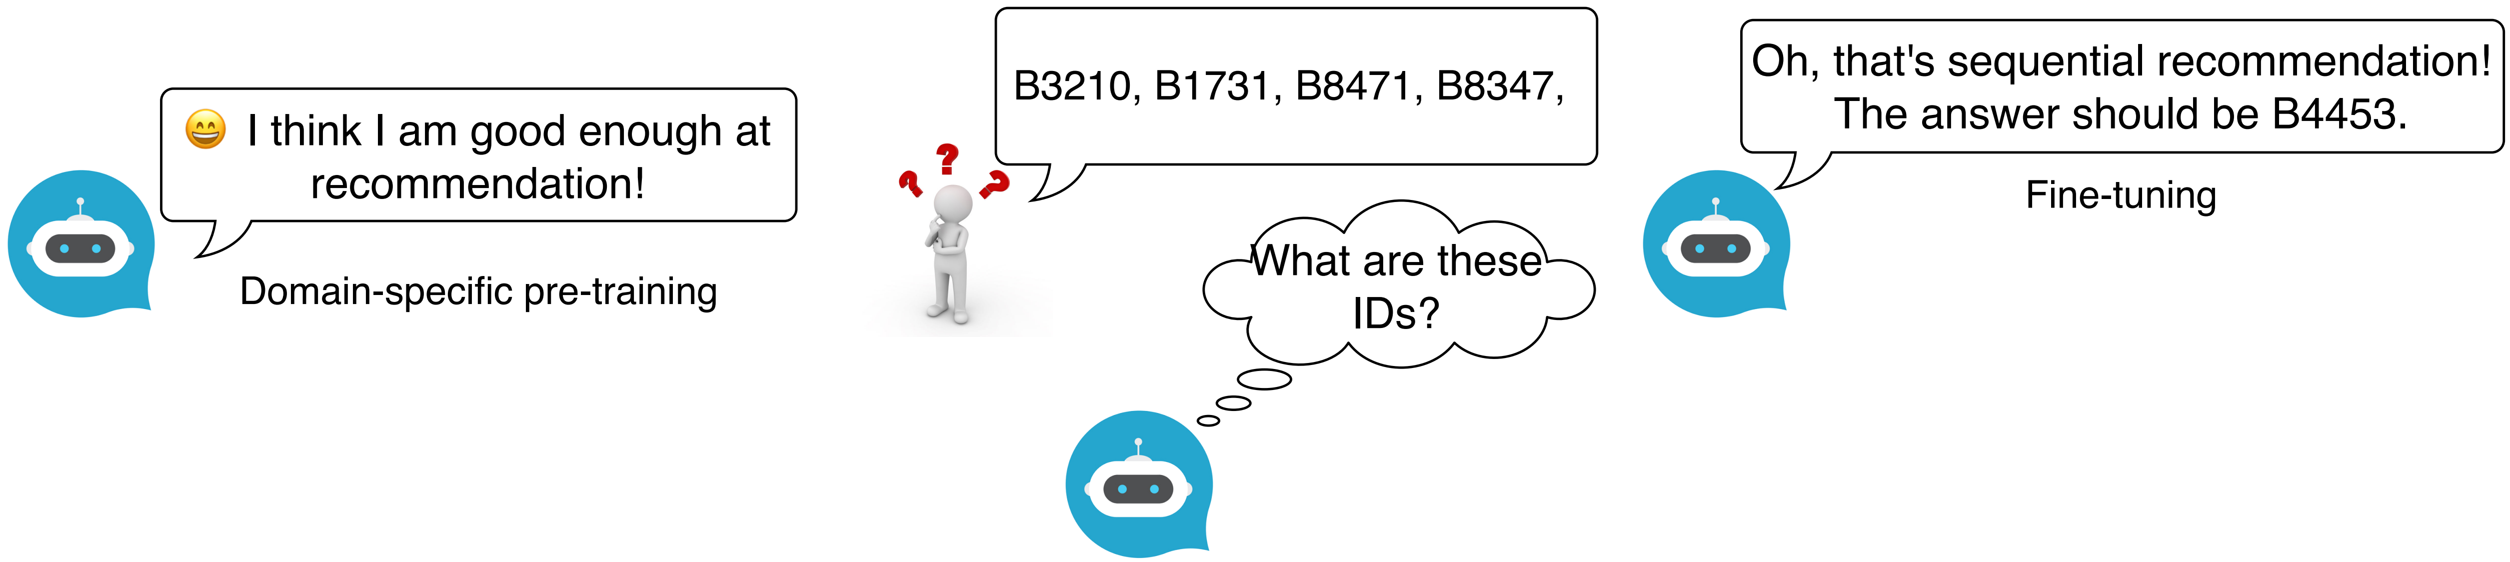

Fine-tuning with Reinforcement Learning

- Reinforcement Learning based on Human Feedbacks (RLHF)
- Proximal Policy Optimization (PPO)
- Direct Preference Optimization (DPO)

Fine-tuning with Reinforcement Learning

- RLHF

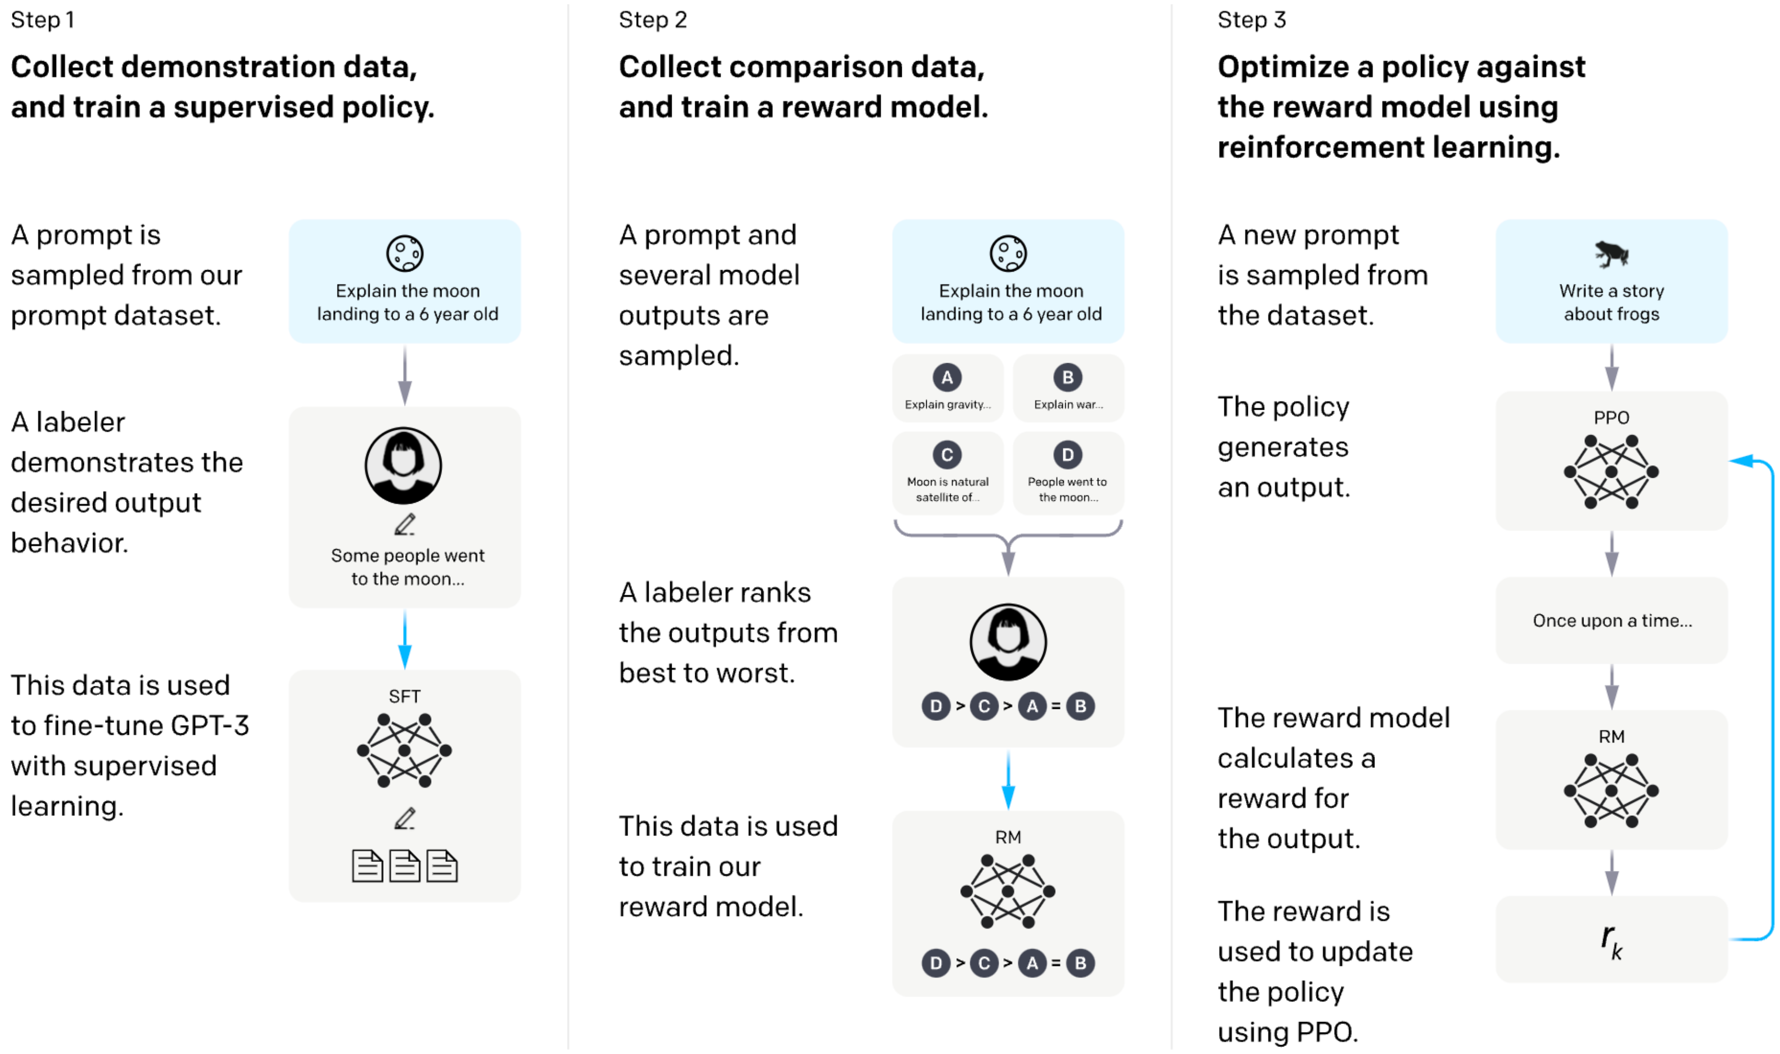

"Deep reinforcement learning from human preferences." Advances in neural information processing systems (2017).

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

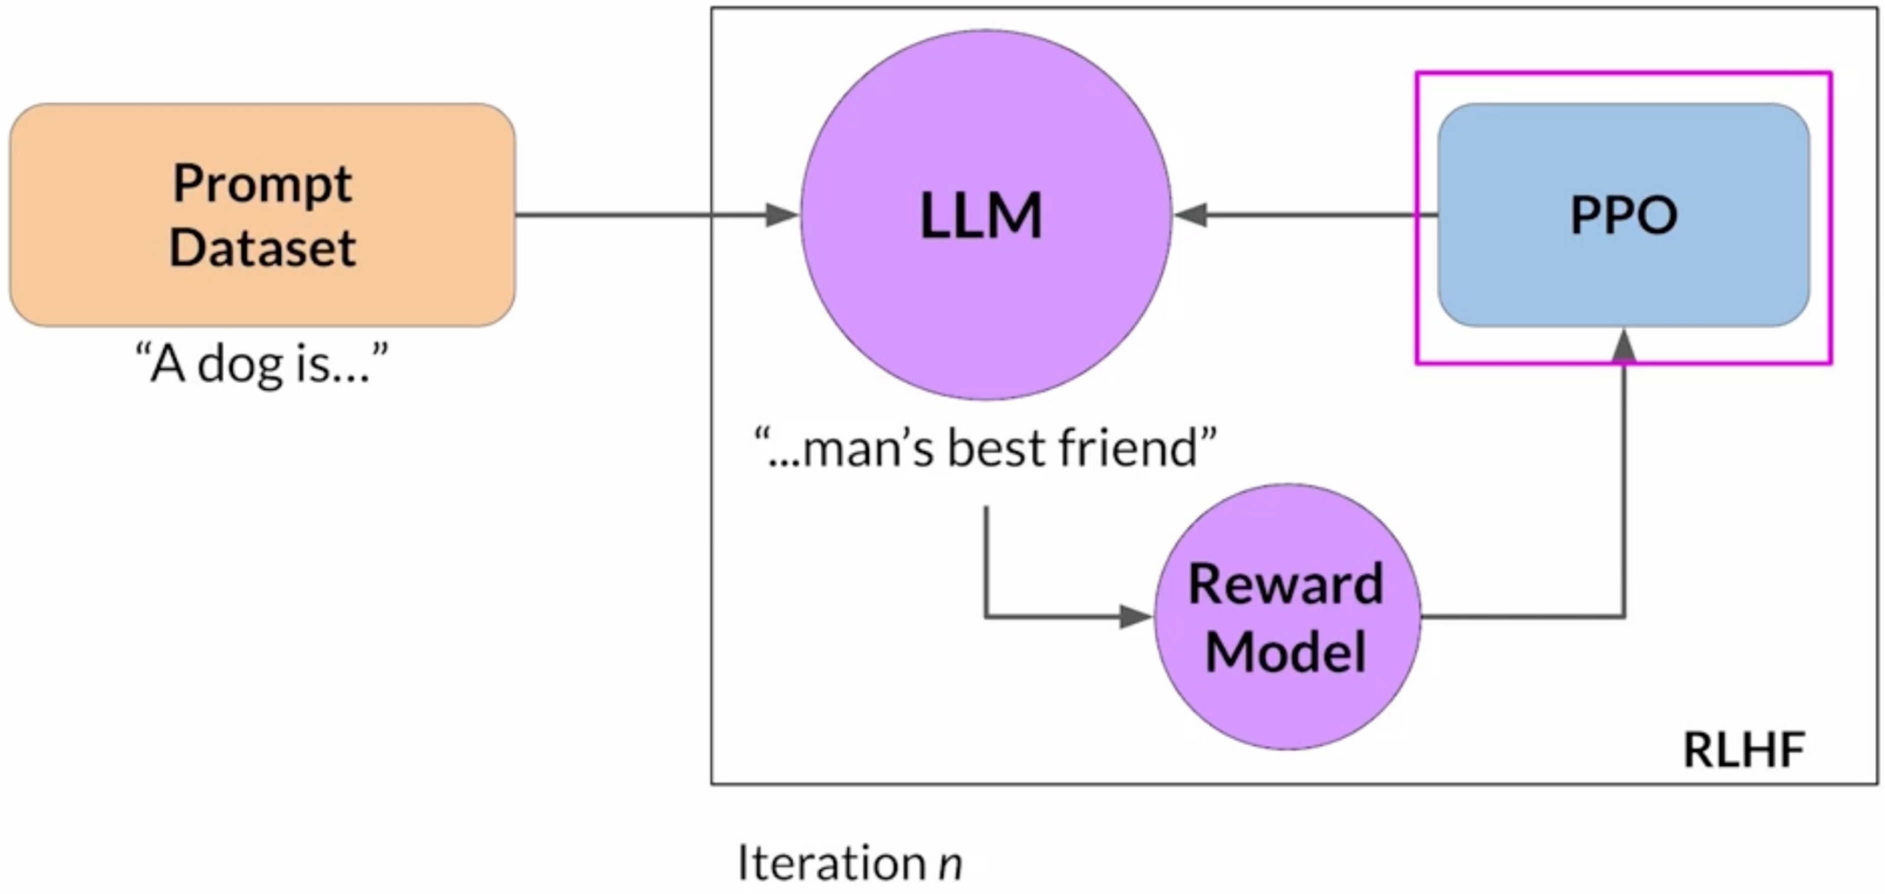

https://www.coursera.org/lecture/generative-ai-with-llms/optional-video-proximal-policy-optimization-1iZJO

Fine-tuning with Reinforcement Learning

- PPO

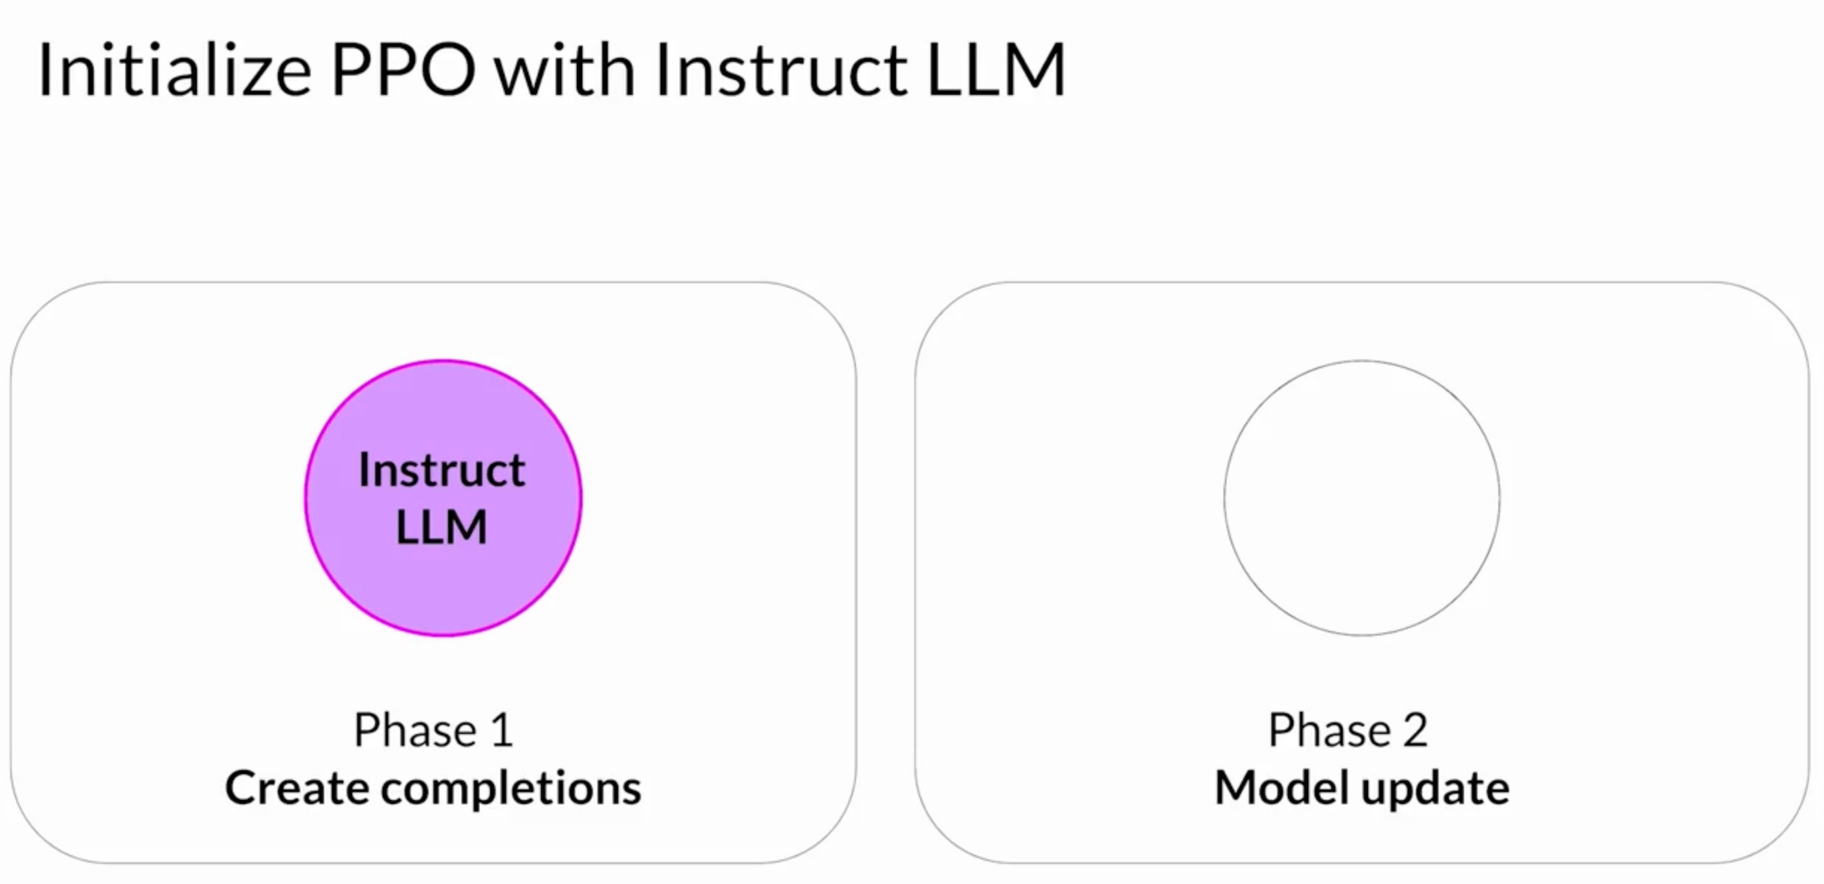

https://www.coursera.org/lecture/generative-ai-with-llms/optional-video-proximal-policy-optimization-1iZJO

Fine-tuning with Reinforcement Learning

- PPO

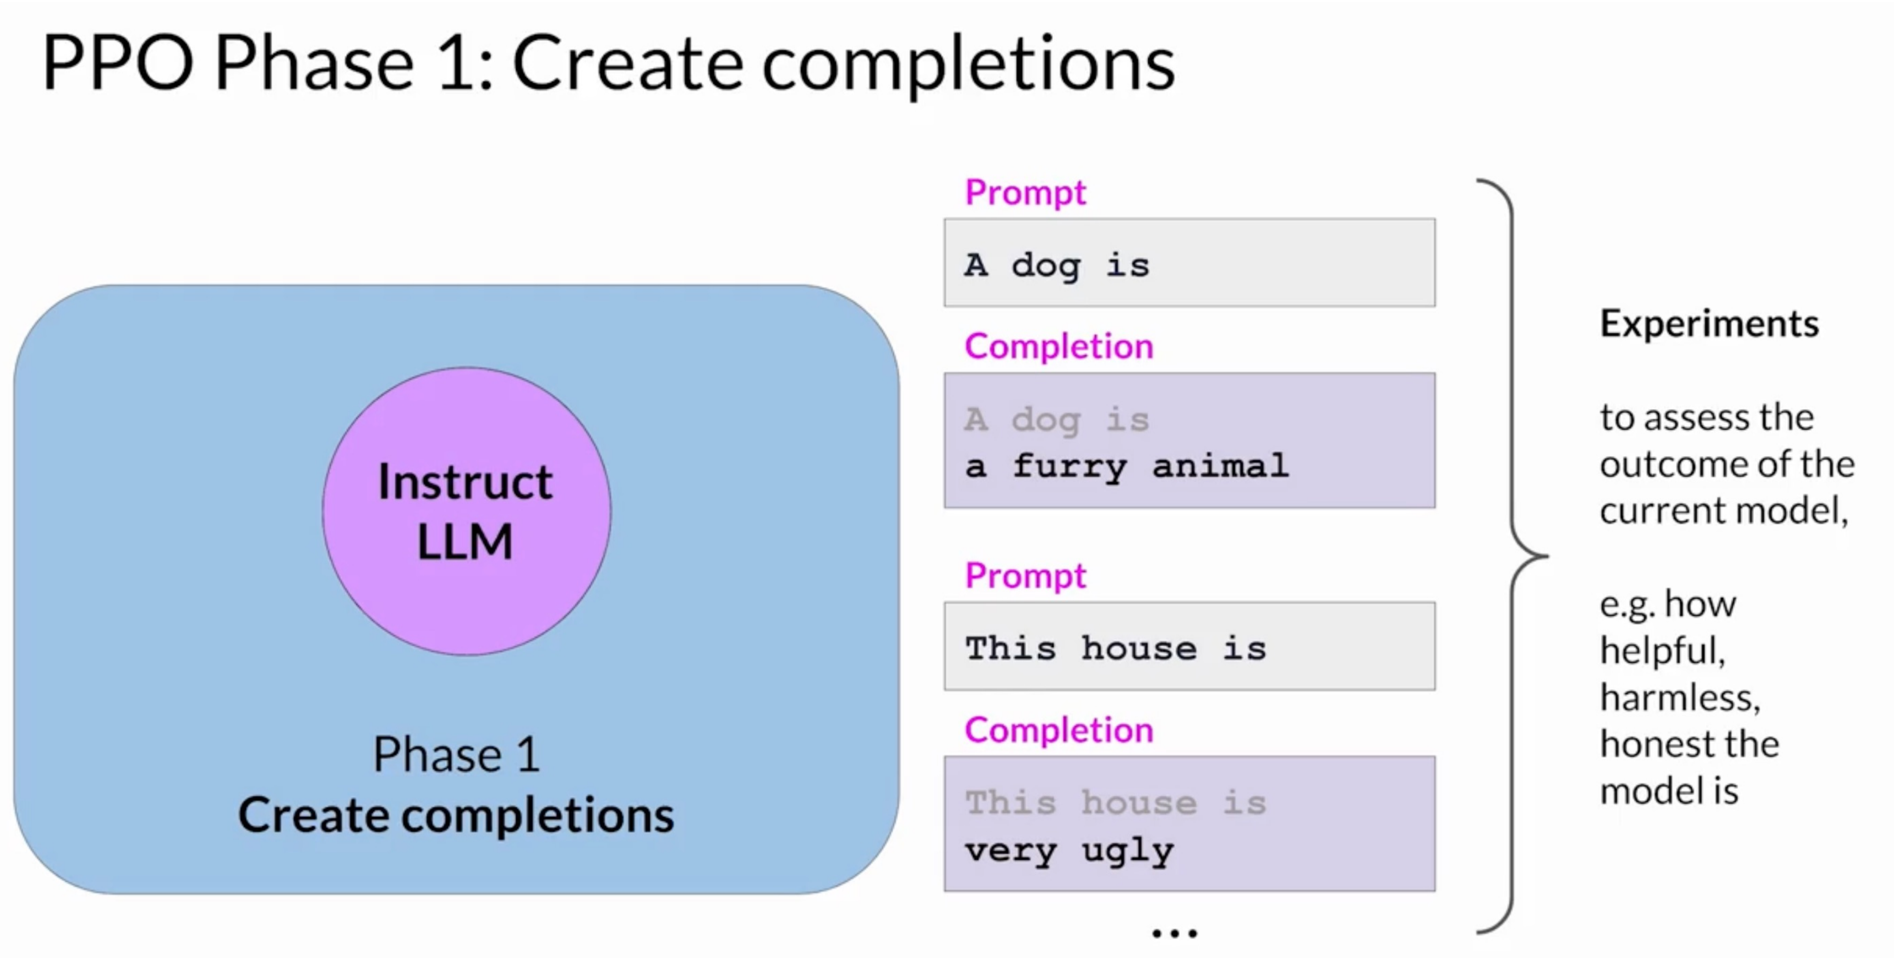

https://www.coursera.org/lecture/generative-ai-with-llms/optional-video-proximal-policy-optimization-1iZJO

Fine-tuning with Reinforcement Learning

- PPO

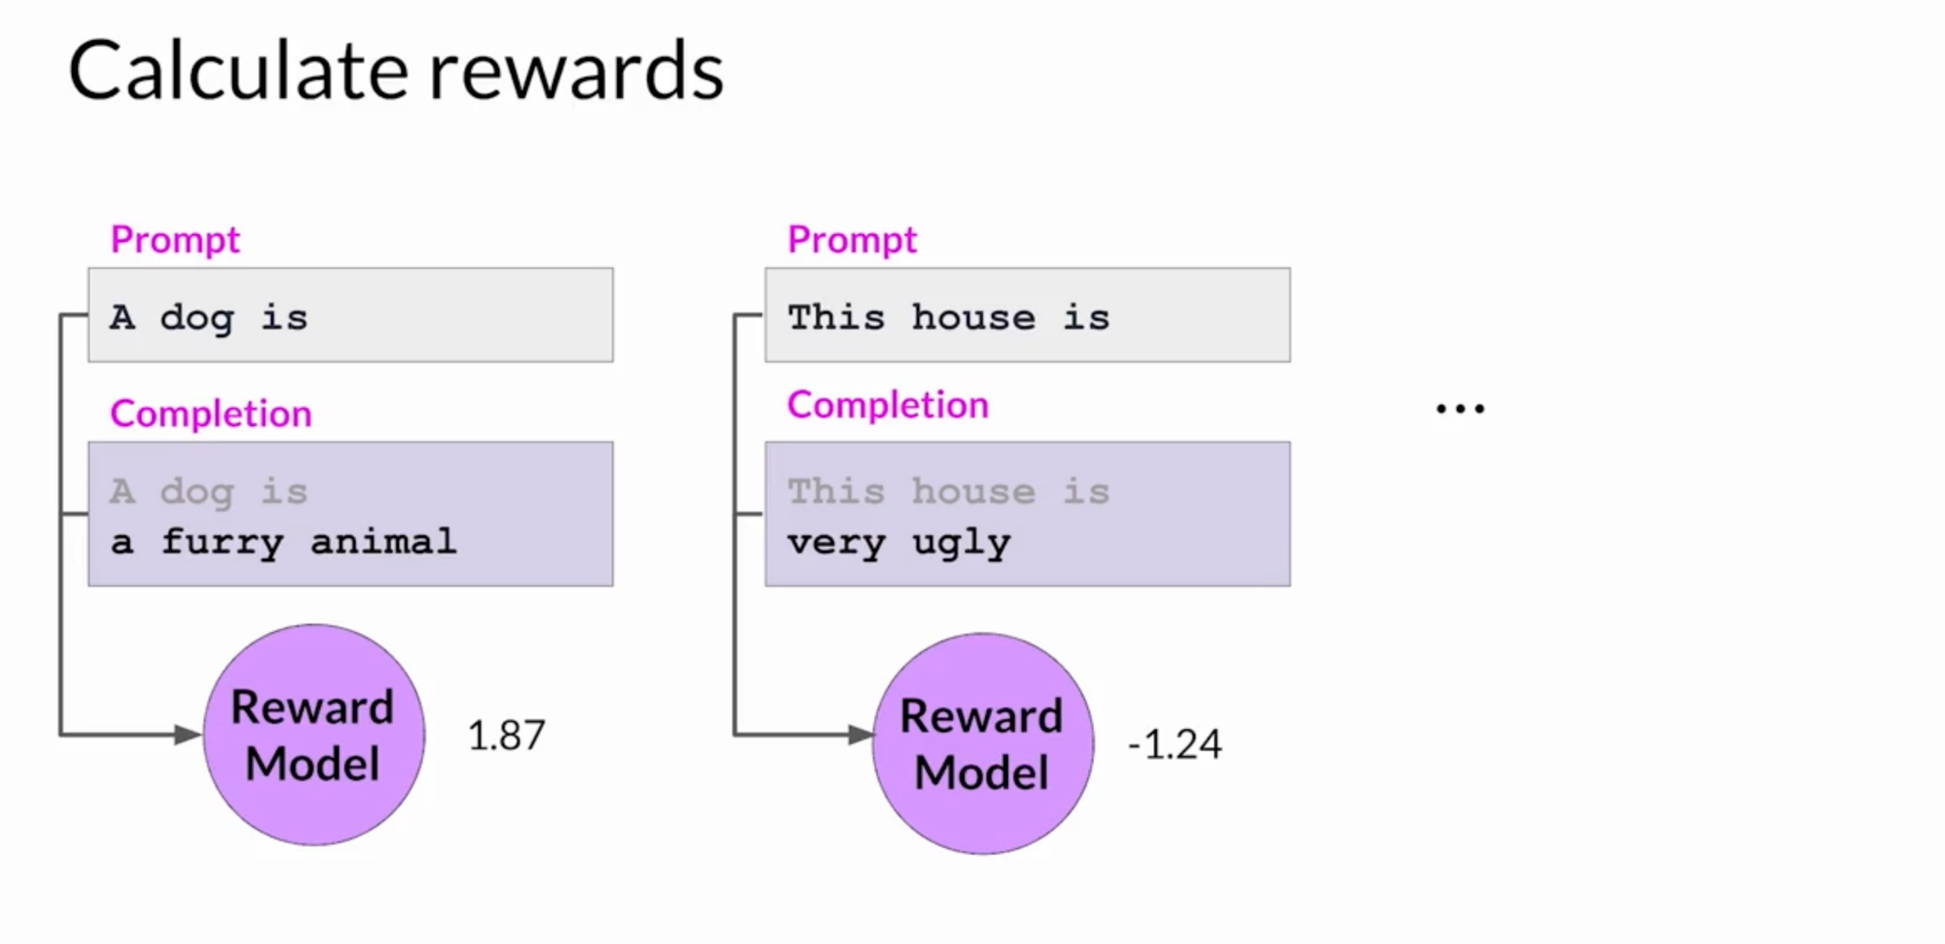

https://www.coursera.org/lecture/generative-ai-with-llms/optional-video-proximal-policy-optimization-1iZJO

Fine-tuning with Reinforcement Learning

- PPO

Calculate value loss

Prompt

```
A dog is
```

Completion

```
A dog is
a furry...
```

Value loss

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

Estimated future total reward: 1.23

Known future total reward: 1.87

https://www.coursera.org/lecture/generative-ai-with-llms/optional-video-proximal-policy-optimization-1iZJO

Fine-tuning with Reinforcement Learning

- PPO

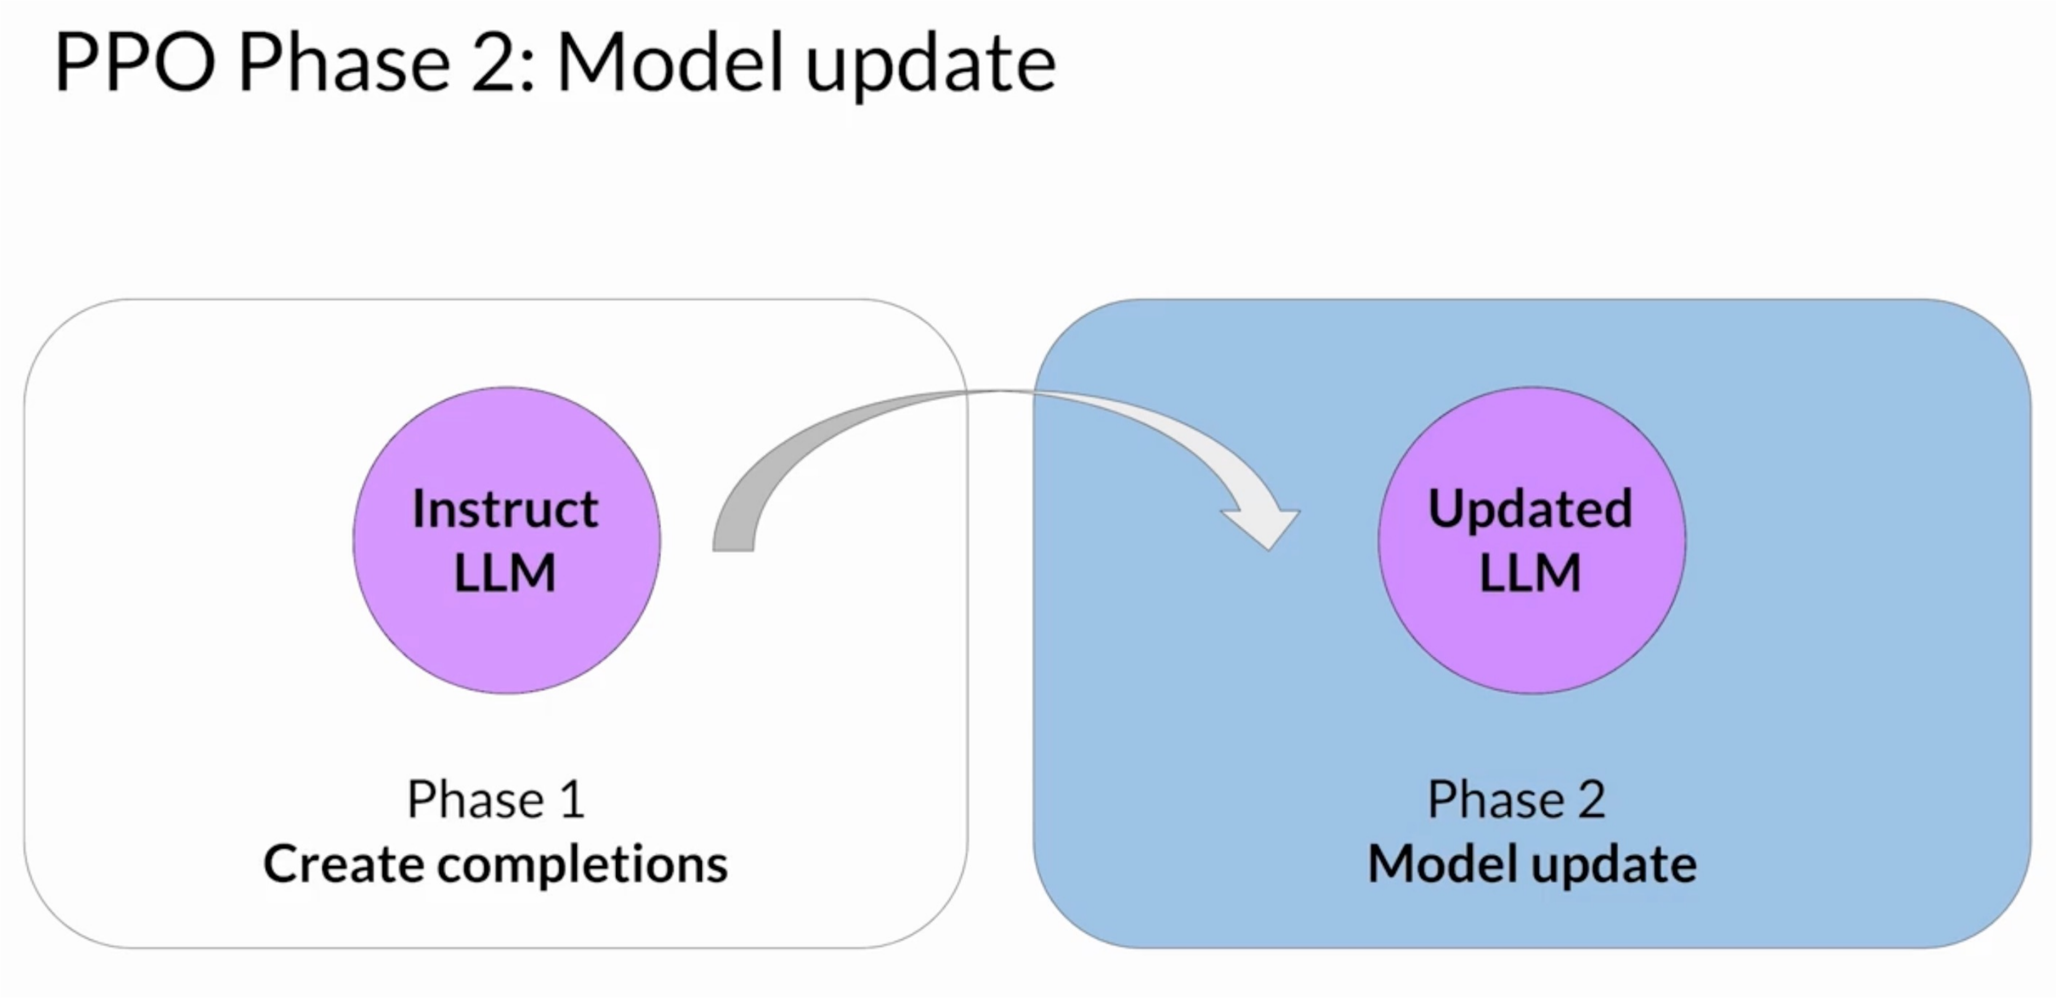

https://www.coursera.org/lecture/generative-ai-with-llms/optional-video-proximal-policy-optimization-1iZJO

Fine-tuning with Reinforcement Learning

- PPO

PPO Phase 2: Calculate policy loss

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

https://www.coursera.org/lecture/generative-ai-with-llms/optional-video-proximal-policy-optimization-1iZJO

Fine-tuning with Reinforcement Learning

- PPO

PPO Phase 2: Calculate entropy loss

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

Low entropy:		High entropy:
Prompt	Prompt	Prompt
`A dog is`	`A dog is`	`A dog is`
Completion	Completion	Completion
`A dog is a domesticated carnivorous mammal`	`A dog is a small carnivorous mammal`	`A dog is is one of the most popular pets around the world`

https://www.coursera.org/lecture/generative-ai-with-llms/optional-video-proximal-policy-optimization-1iZJO

Fine-tuning with Reinforcement Learning

- PPO

PPO Phase 2: Objective function

Hyperparameters

$$L^{PPO} = \underbrace{L^{POLICY}}_{\text{Policy loss}} + \underbrace{c_1 L^{VF}}_{\text{Value loss}} + \underbrace{c_2 L^{ENT}}_{\text{Entropy loss}}$$

https://www.coursera.org/lecture/generative-ai-with-llms/optional-video-proximal-policy-optimization-1iZJO

Fine-tuning with Reinforcement Learning

- PPO

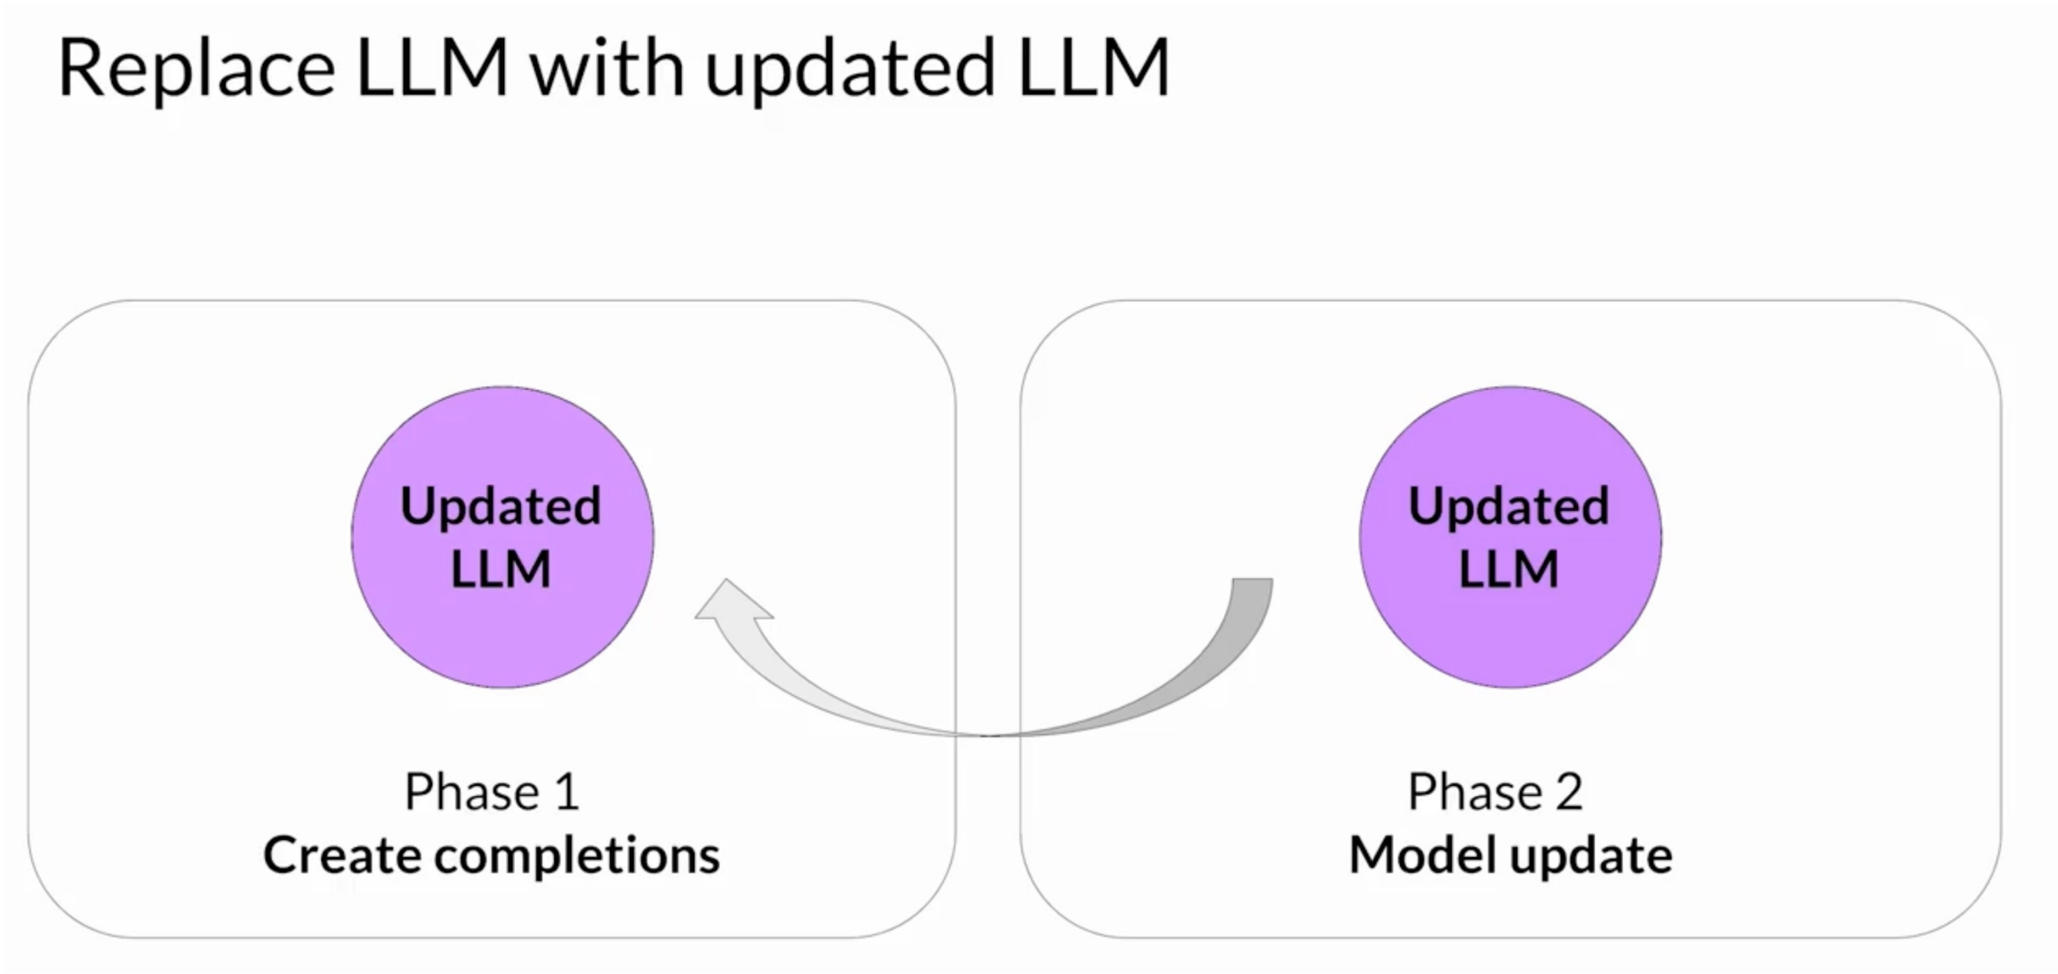

https://www.coursera.org/lecture/generative-ai-with-llms/optional-video-proximal-policy-optimization-1iZJO

Fine-tuning with Reinforcement Learning

❑ PPO

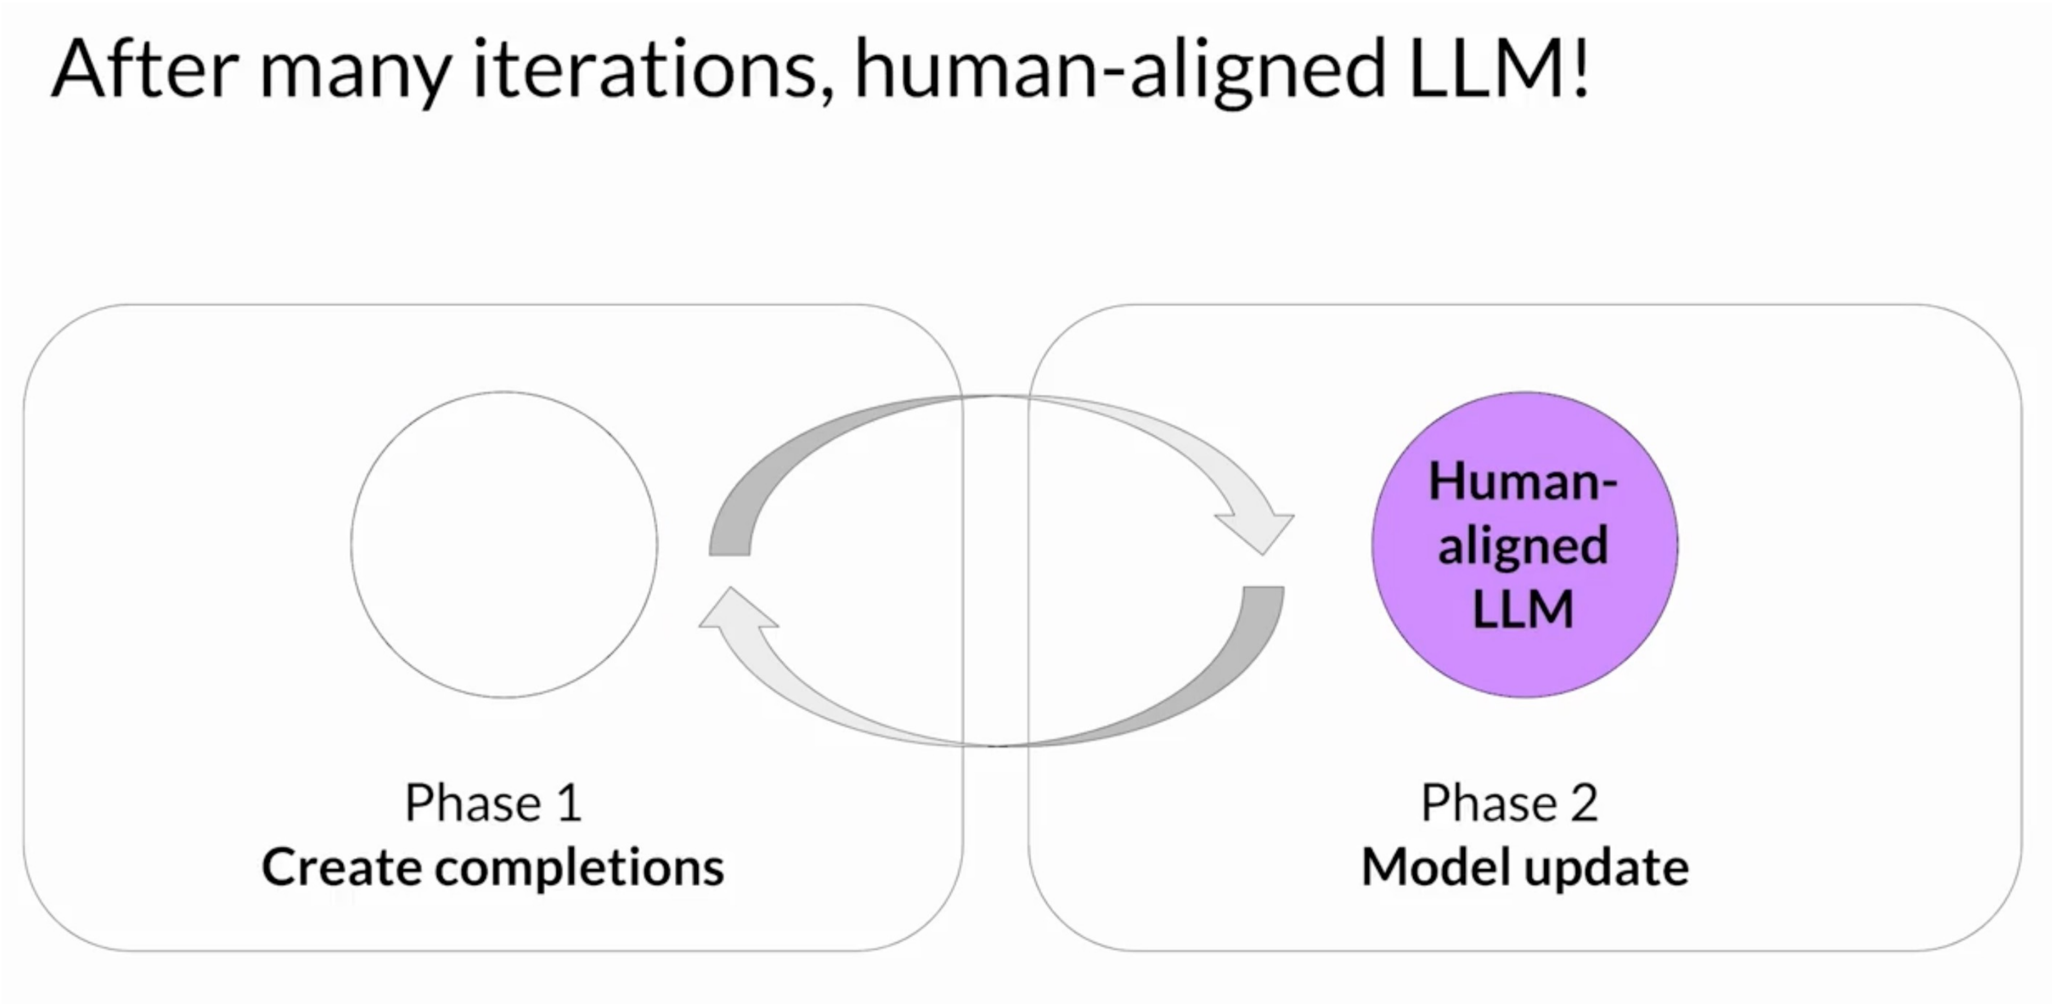

https://www.coursera.org/lecture/generative-ai-with-llms/optional-video-proximal-policy-optimization-1iZJO

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

$$\mathcal{L}_{\text{DPO}}(\underbrace{\pi_\theta}_{\text{Policy to optimize}}; \underbrace{\pi_{\text{ref}}}_{\text{Reference policy (used to control behavior of LLMs)}}) = -\underbrace{\mathbb{E}_{(x,y_w,y_l)\sim\mathcal{D}}}_{\text{Aggregation over preference data}}\left[\log \underbrace{\sigma}_{\text{Logistic function}}\left(\underbrace{\beta \log \frac{\pi_\theta(y_w \mid x)}{\pi_{\text{ref}}(y_w \mid x)}}_{\text{Shift in preferred completion}} - \underbrace{\beta \log \frac{\pi_\theta(y_l \mid x)}{\pi_{\text{ref}}(y_l \mid x)}}_{\text{Shift in dispreferred completion}}\right)\right]$$

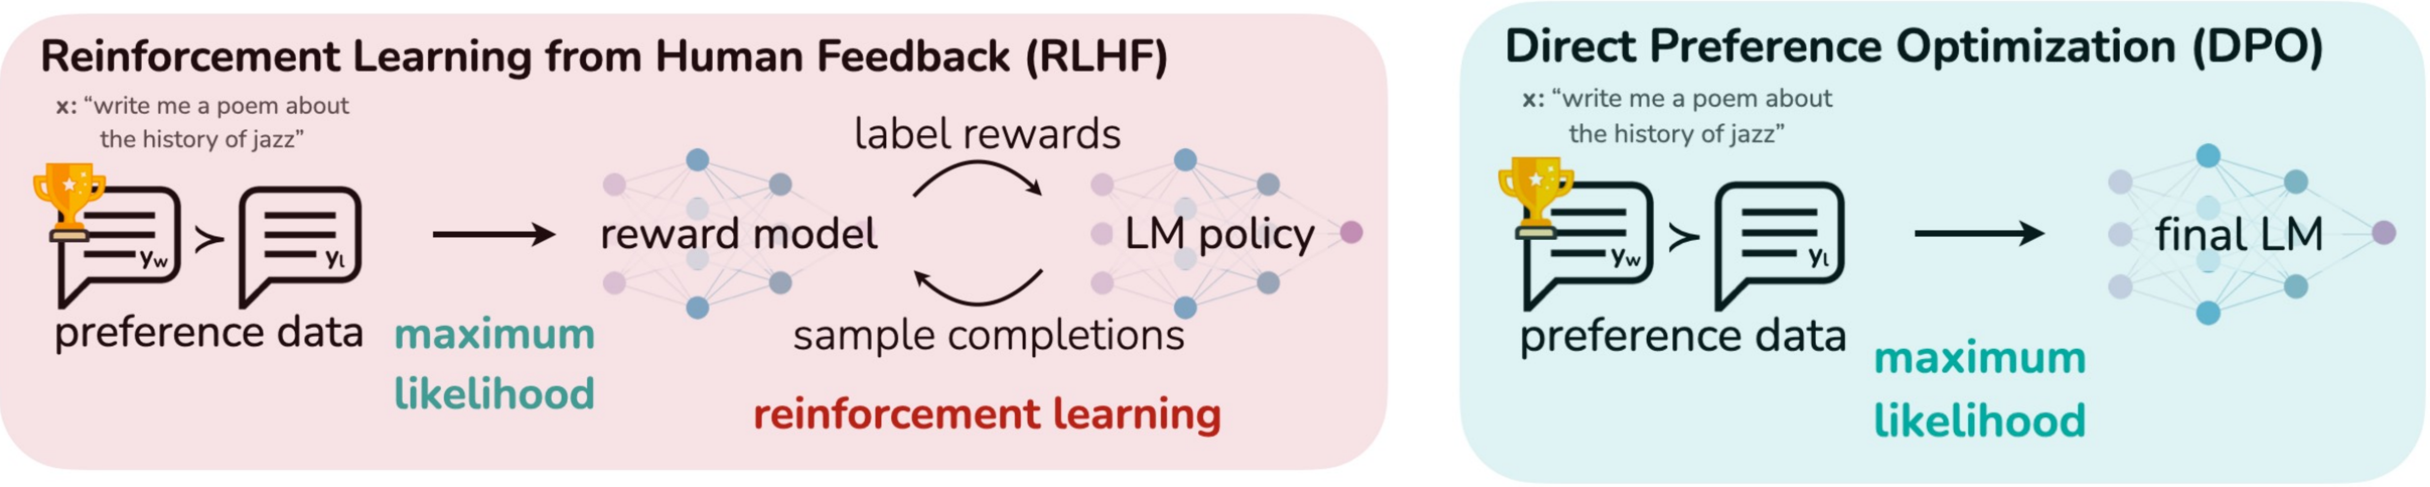

"Direct preference optimization: Your language model is secretly a reward model." arXiv preprint (2023).

Preliminaries | Pre-training | **Fine-tuning** | Prompting | Future Directions

Overview

Presenter:
Jiatong Li

Fine-tuning

- ○ Fine-tuning in NLP
- ⦿ Fine-tuning LLM-based RecSys

Parameter Efficient Fine-tuning

Zoom ID:
91649466943
Password:
202312

Website QR Code

GIRL

❑ Multi-steps of Fine-tuning with SFT and RL

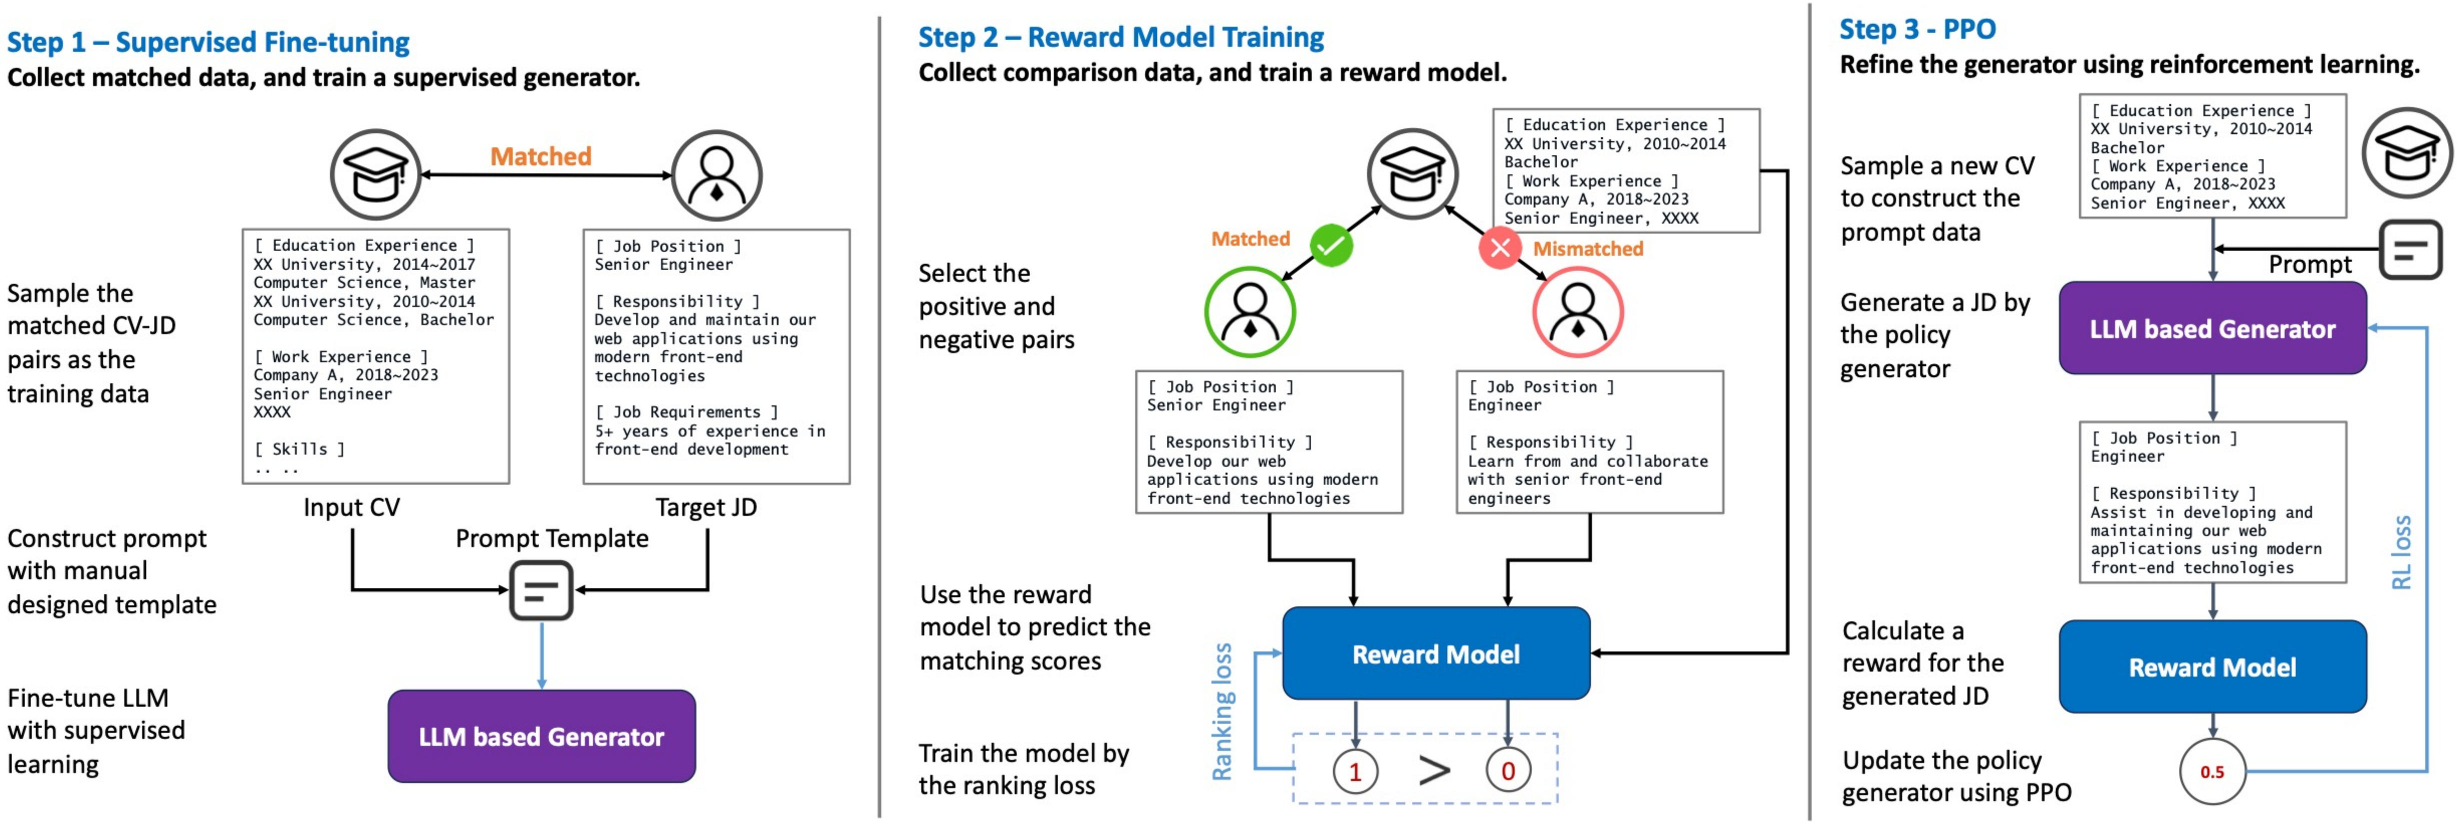

$$\mathcal{L}_{sft} = -\log \Pr(C|J,T,\mathcal{G}) = -\sum_{i=1}^{|l_j|} \log \Pr(v_i|v_{<i},C,T,\mathcal{G})$$

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

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

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

"Generative job recommendations with large language model." arXiv preprint (2023).

TransRec

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

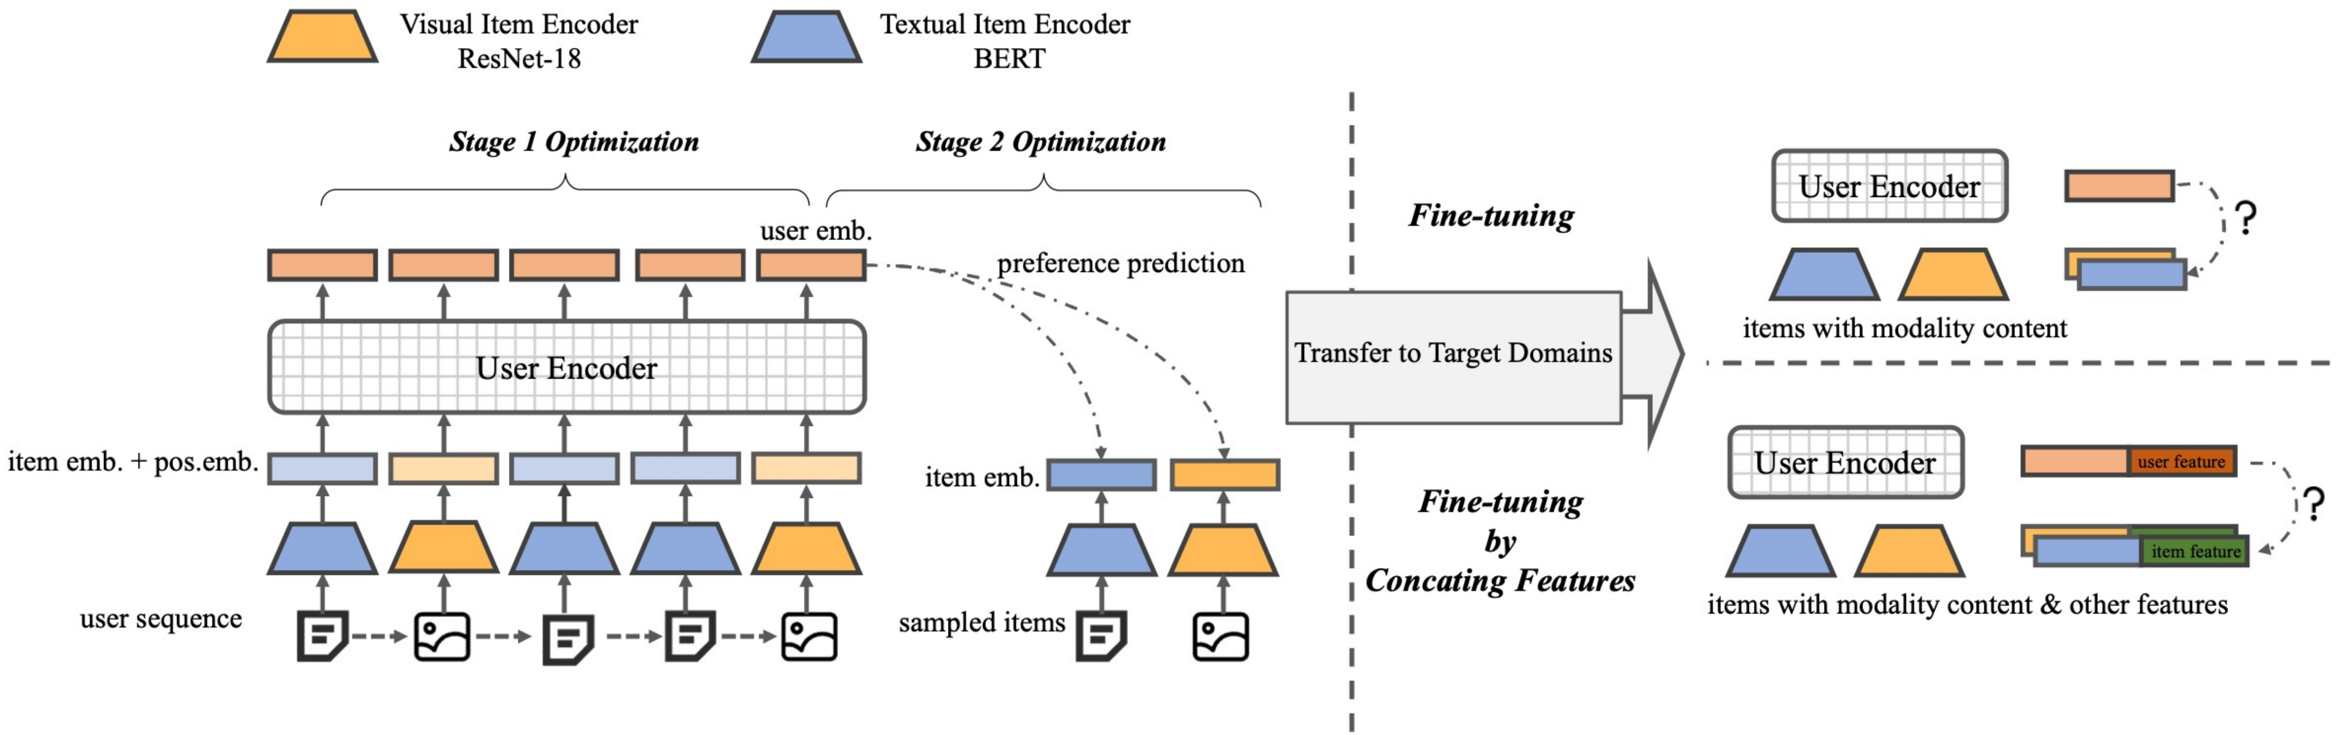

(a) Pre-training on Source Domain

(b) Fine-tuning on Target Domains

"Transrec: Learning transferable recommendation from mixture-of-modality feedback." arXiv preprint (2022).

SBERT

❑ Fine-tuning LLM-based RecSys with Contrastive Learning

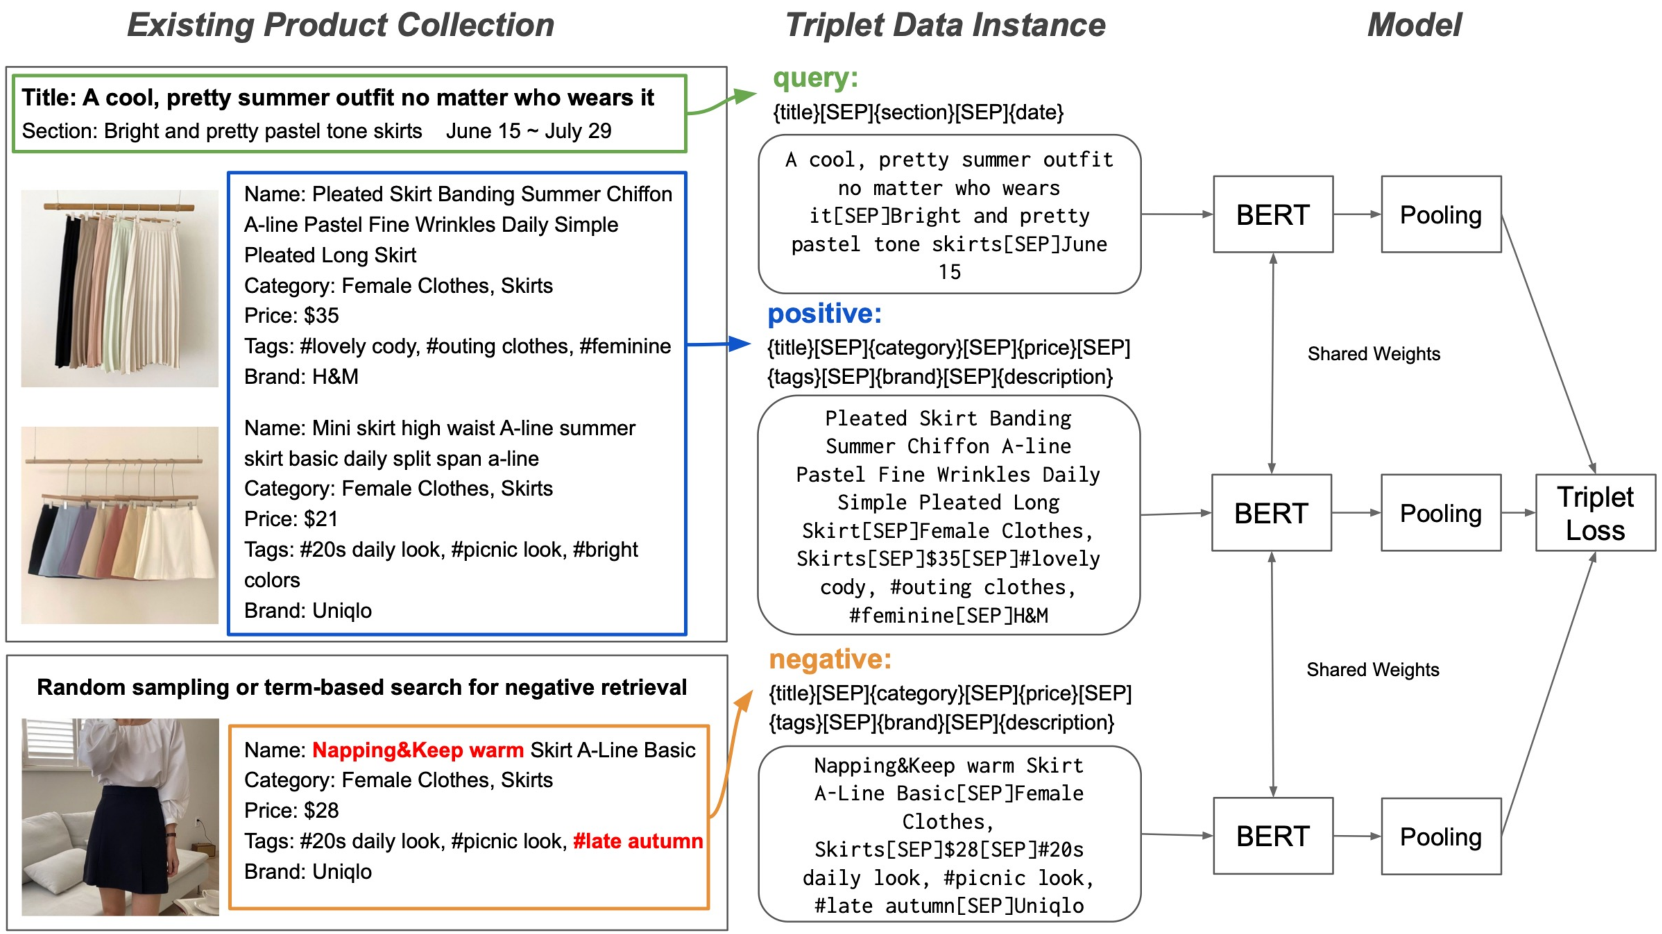

"Intent-based Product Collections for E-commerce using Pretrained Language Models." ICDMW (2021).

Preliminaries › Pre-training › **Fine-tuning** › Prompting › Future Directions

Overview

Presenter:
Jiatong Li

Fine-tuning

- Fine-tuning in NLP
- Fine-tuning LLM-based RecSys

Parameter Efficient Fine-tuning

Zoom ID:
91649466943
Password:
202312

Website QR Code

Parameter Efficient Fine-tuning

❑ 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

M6-Rec

- Option Adapter Fine-tunes LLMs

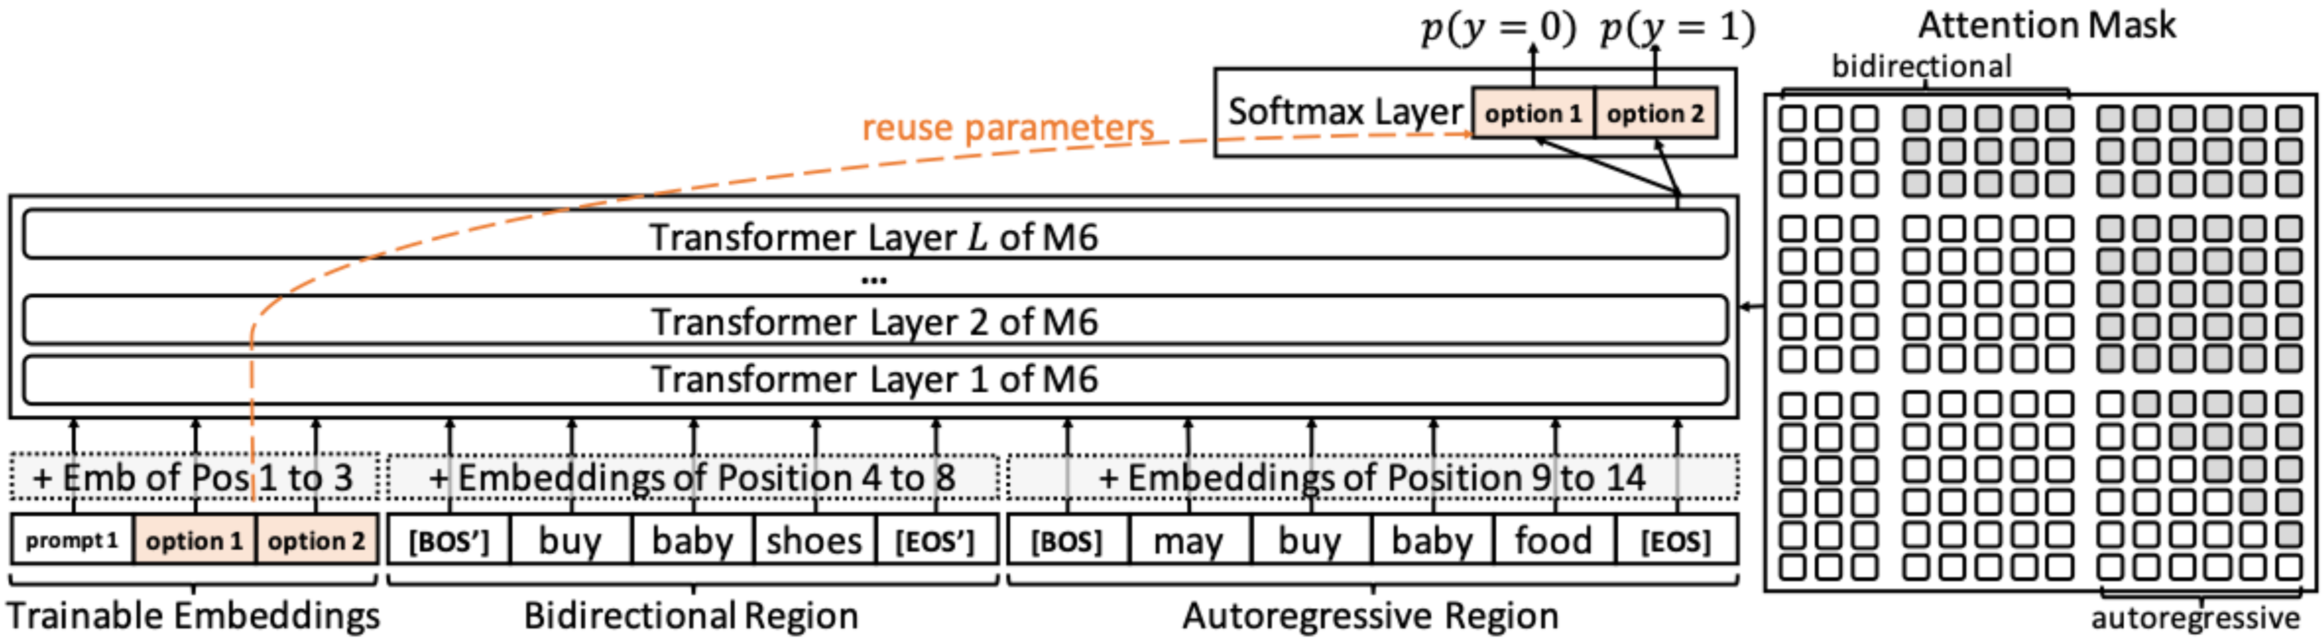

"M6-rec: Generative pretrained language models are open-ended recommender systems." arXiv preprint (2022).

LoRA Adaptation

- Low-Rank Adaptation of LLMs (LoRA) fine-tuning
  - Fine-tuning a 7B model needs 7,000,000,000∗8/1024^3 ≅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

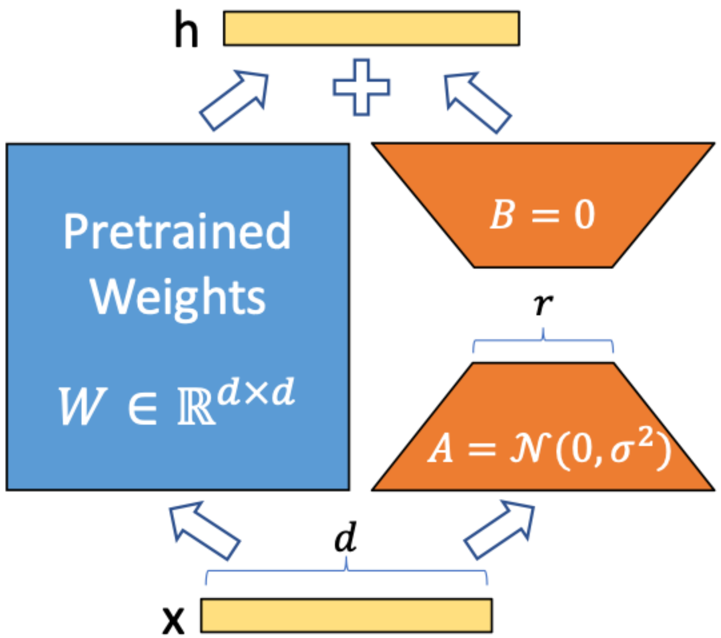

"Lora: Low-rank adaptation of large language models." arXiv preprint (2021).

TALLRec

❑ LoRA Fine-tune LLMs

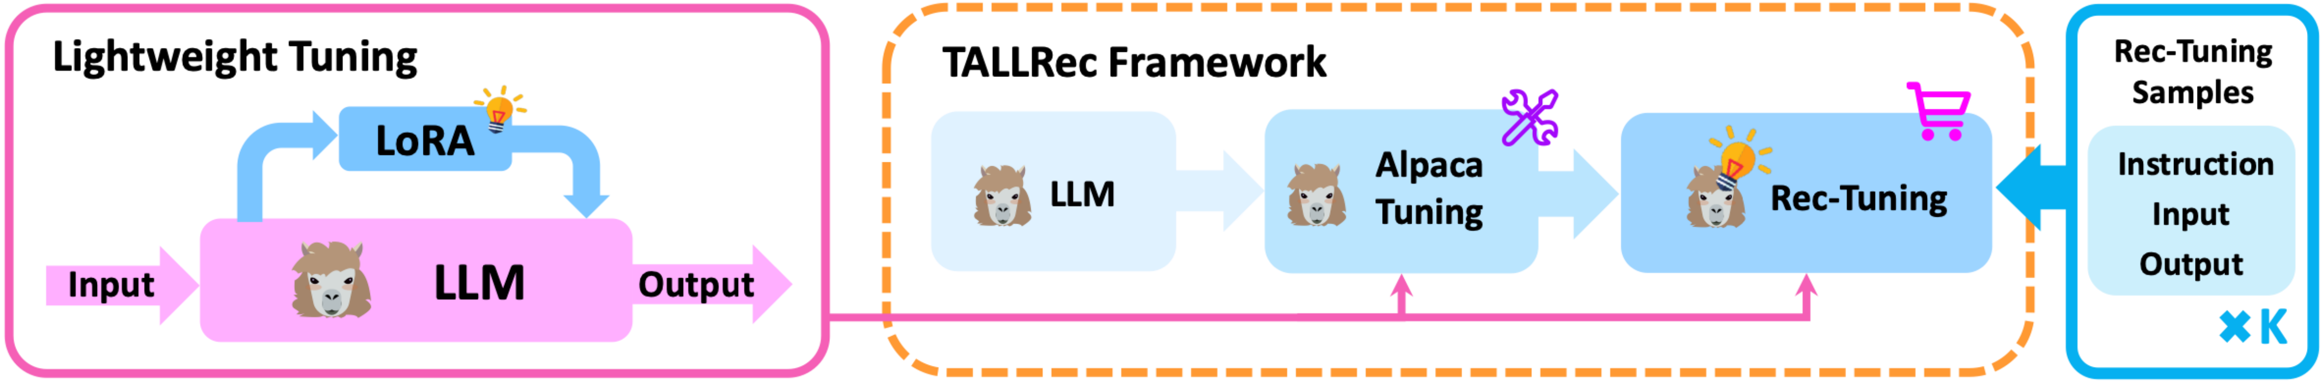

"Tallrec: An effective and efficient tuning framework to align large language model with recommendation." RecSys (2023).

Overview

Presenter:
Zihuai Zhao

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

Zoom ID:
91649466943
Password:
202312

Website QR Code

Brief Ideas of Prompt

- An **intuitive prompt design** for ChatGPT

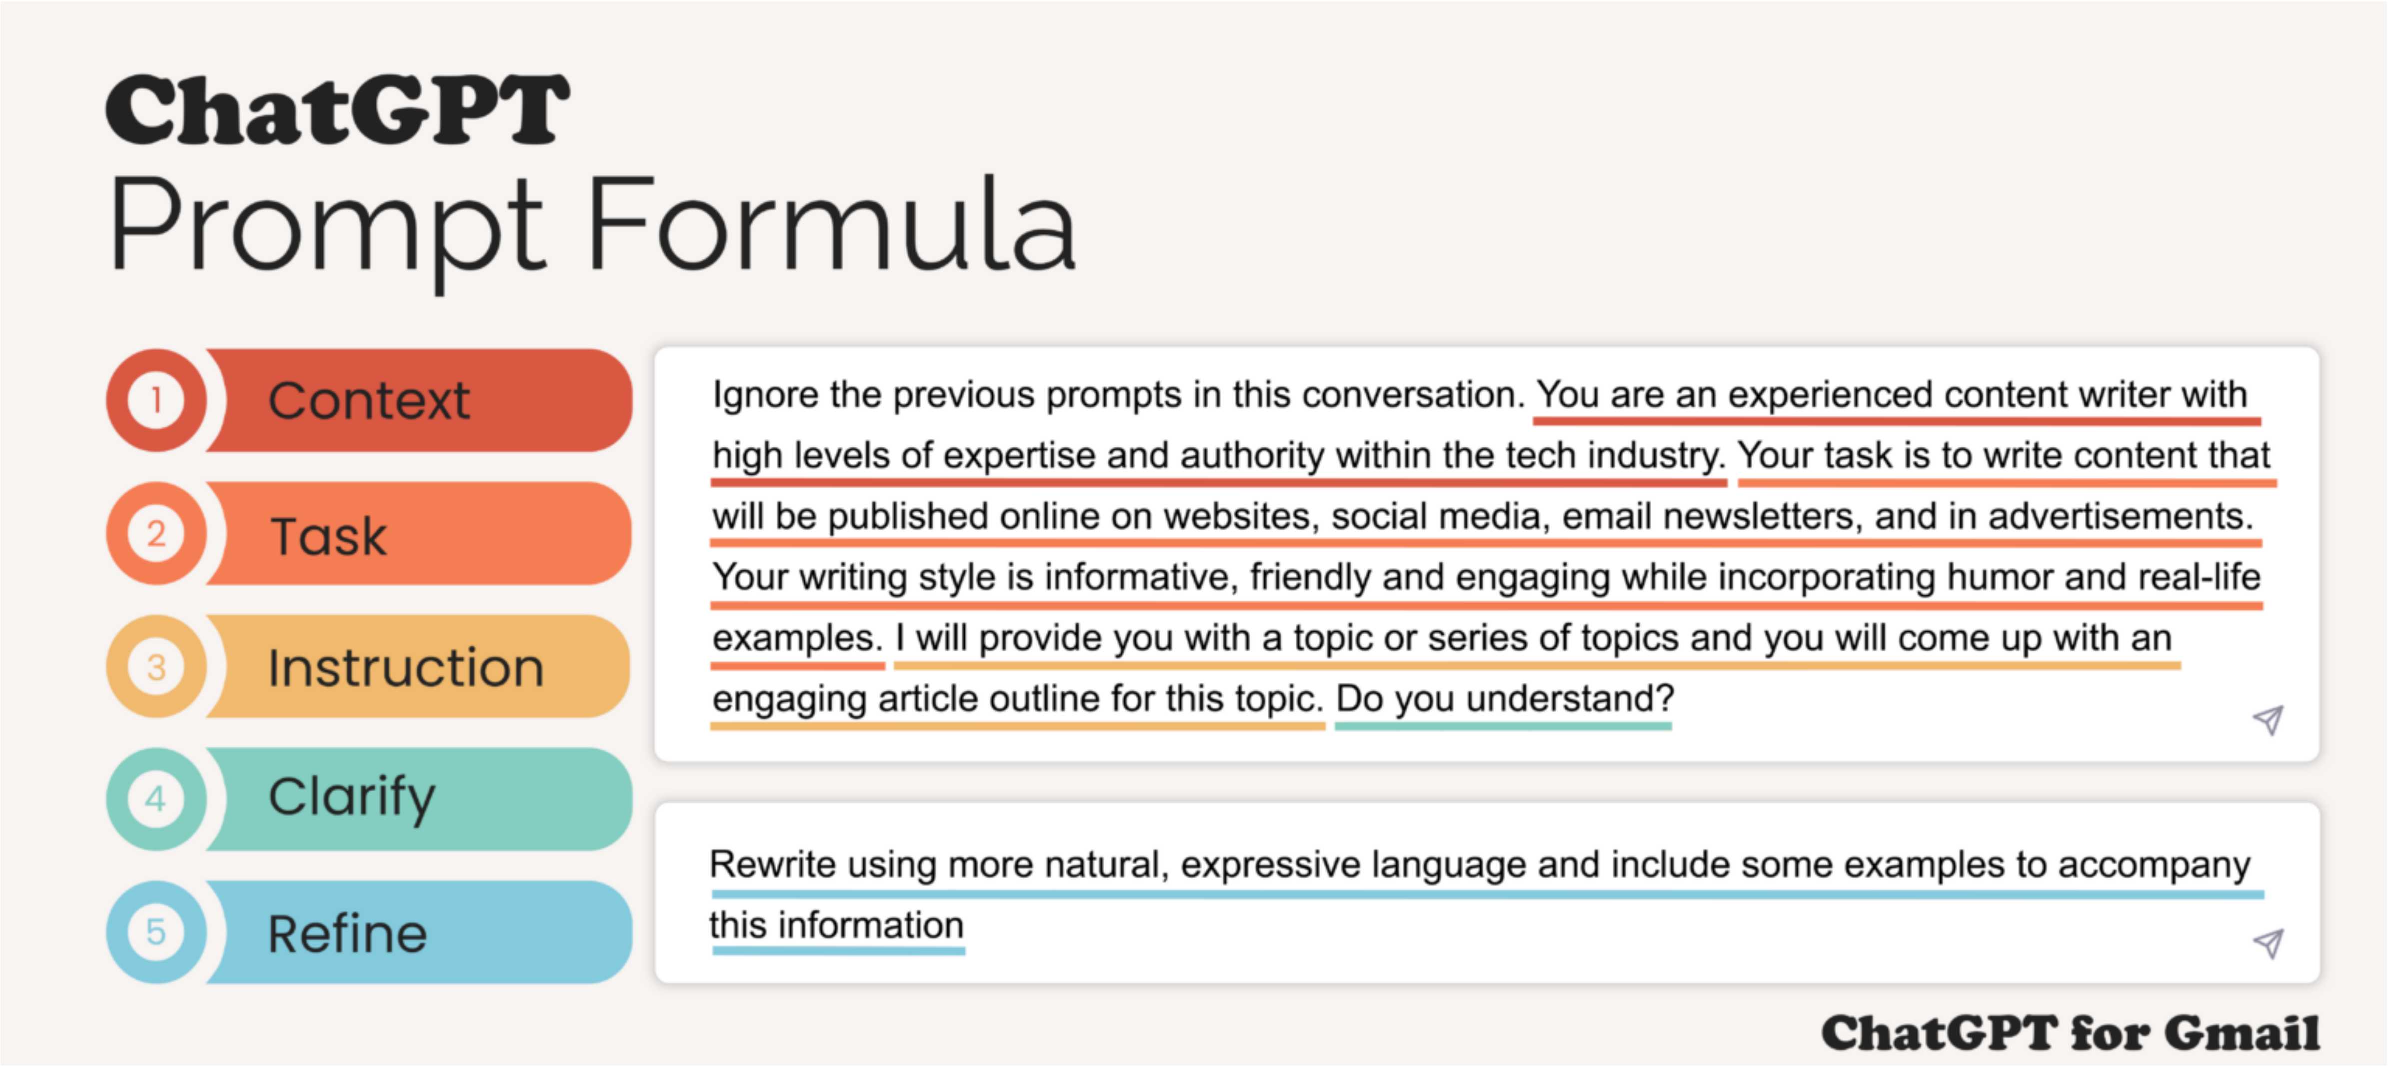

https://blog.cloudhq.net/how-to-write-chatgpt-prompts-for-email/

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

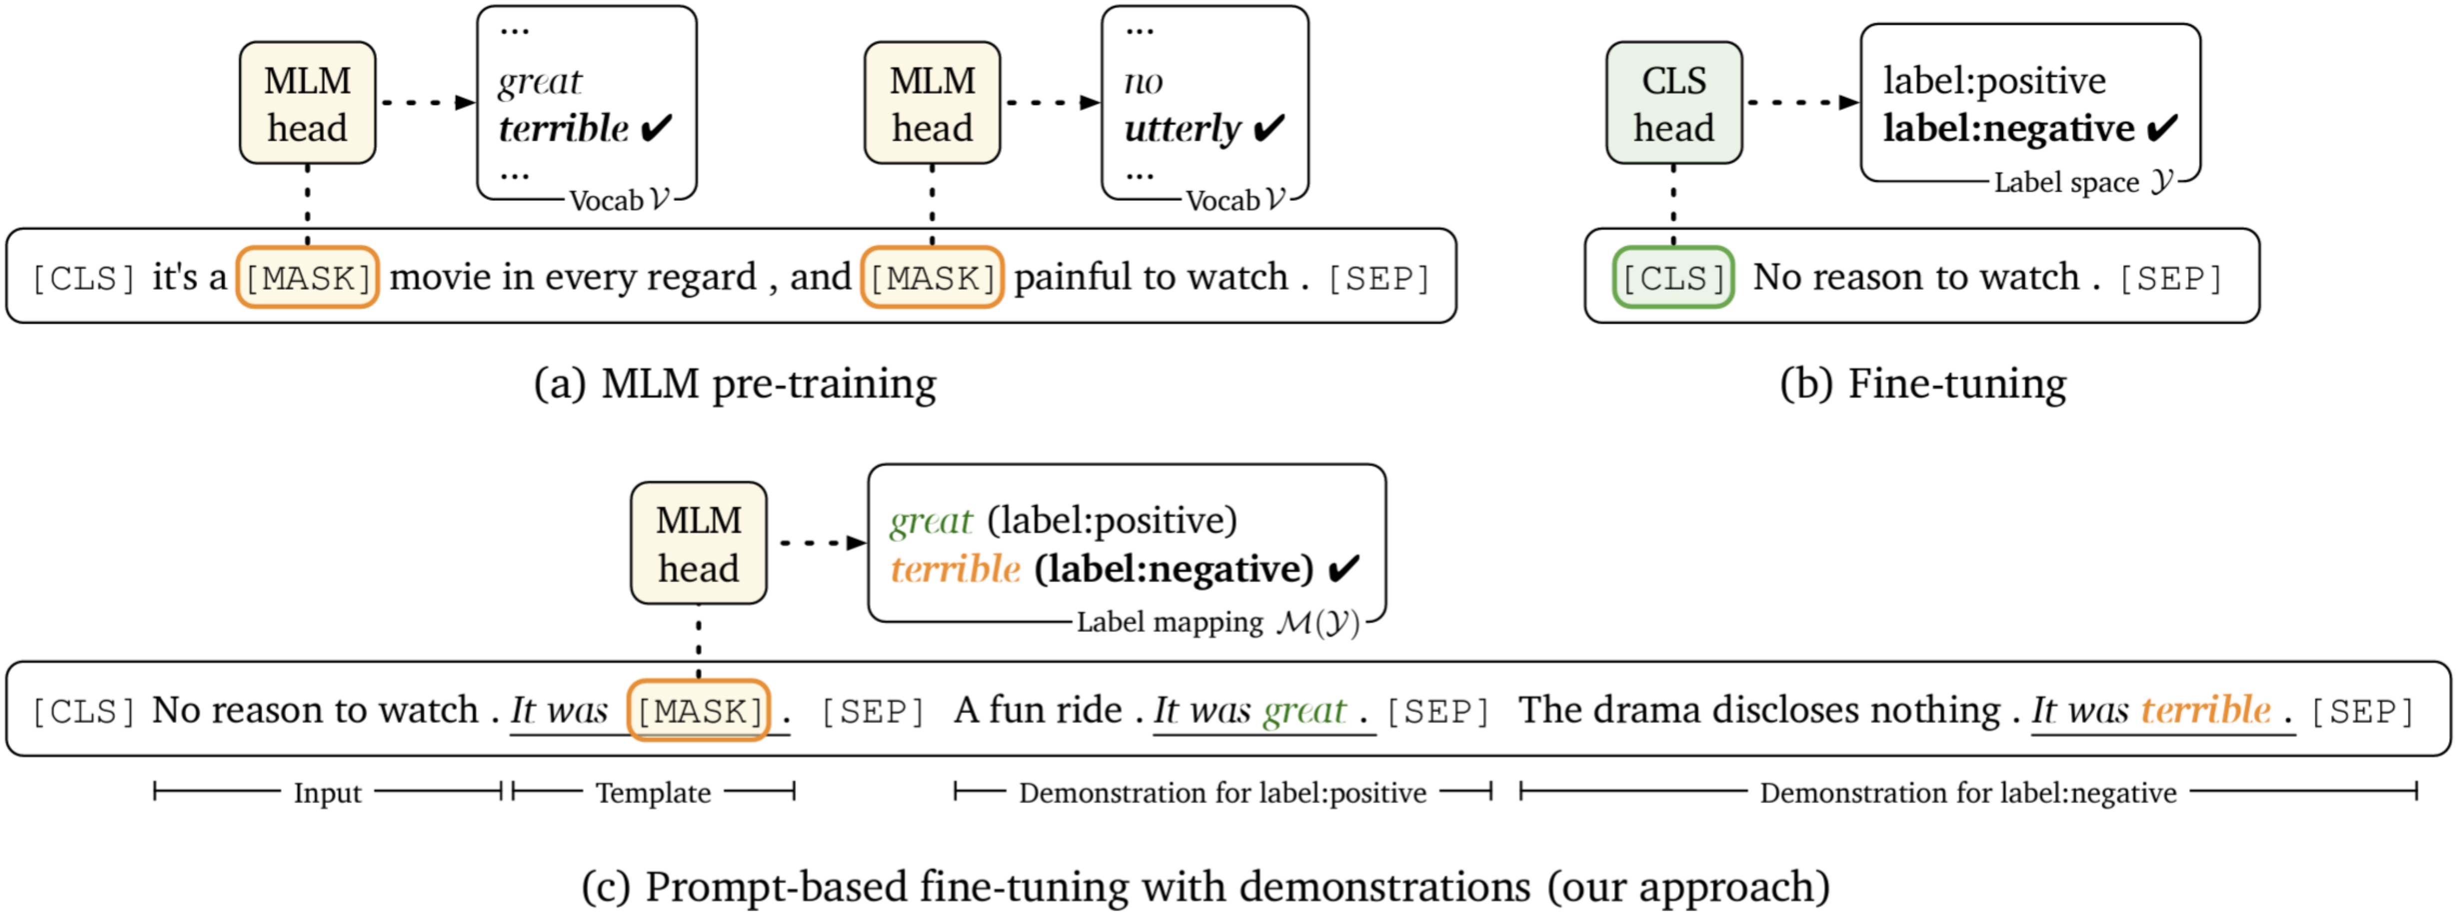

"Making pre-trained language models better few-shot learners." ACL (2021).

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**.

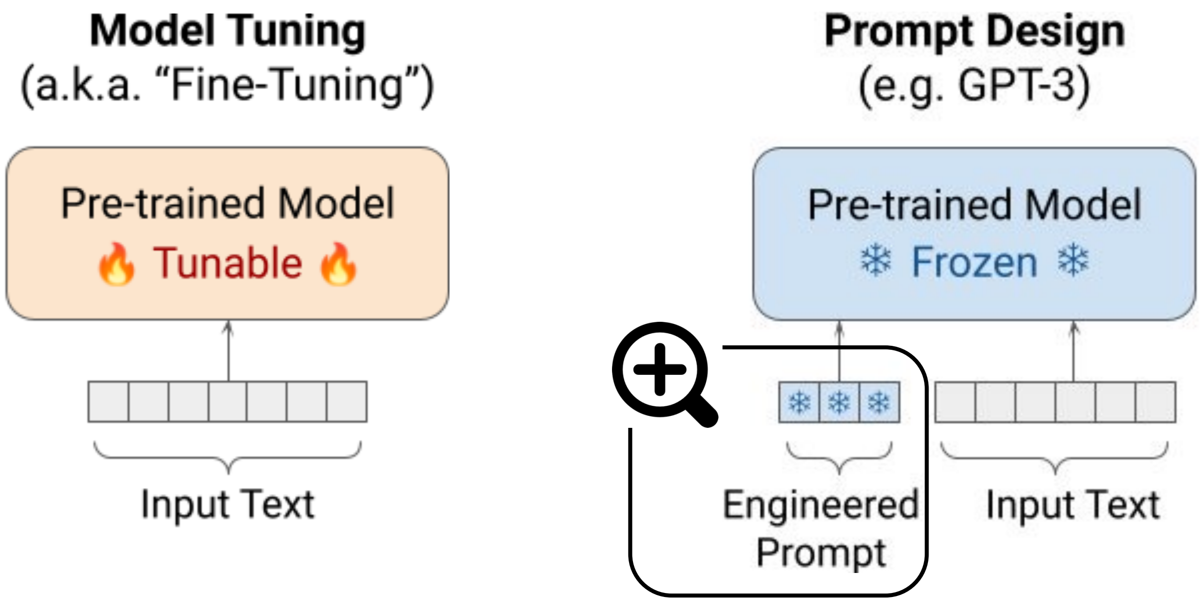

e.g., *"Will the user __ buy item __?"*

"The Power of Scale for Parameter-Efficient Prompt Tuning" EMNLP (2021).

Overview

Presenter:
Zihuai Zhao

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

Zoom ID:
91649466943
Password:
202312

Website QR Code

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.

```
Translate English to French:      ← task description
cheese => ..................      ← prompt
```

Few-shot

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

```
Translate English to French:           ← task description
sea otter => loutre de mer             ← examples
peppermint => menthe poivrée           ←
plush girafe => girafe peluche         ←
cheese => ..................           ← prompt
```

"Language models are few-shot learners." NeurIPS (2020)

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 …”*

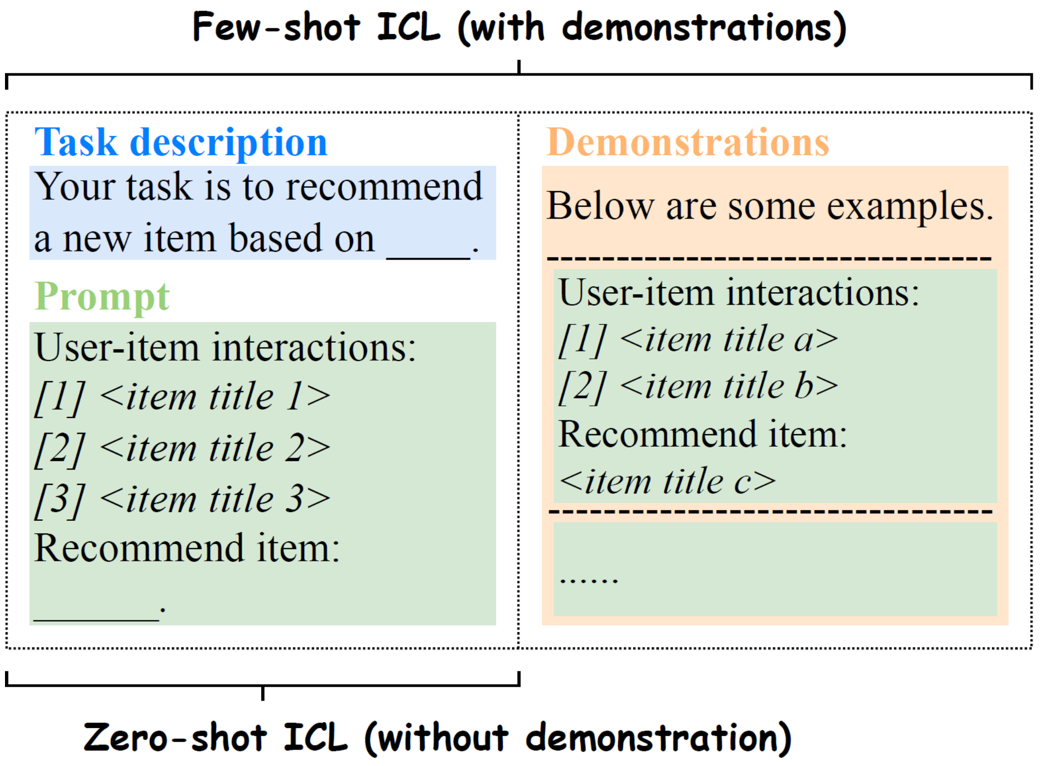

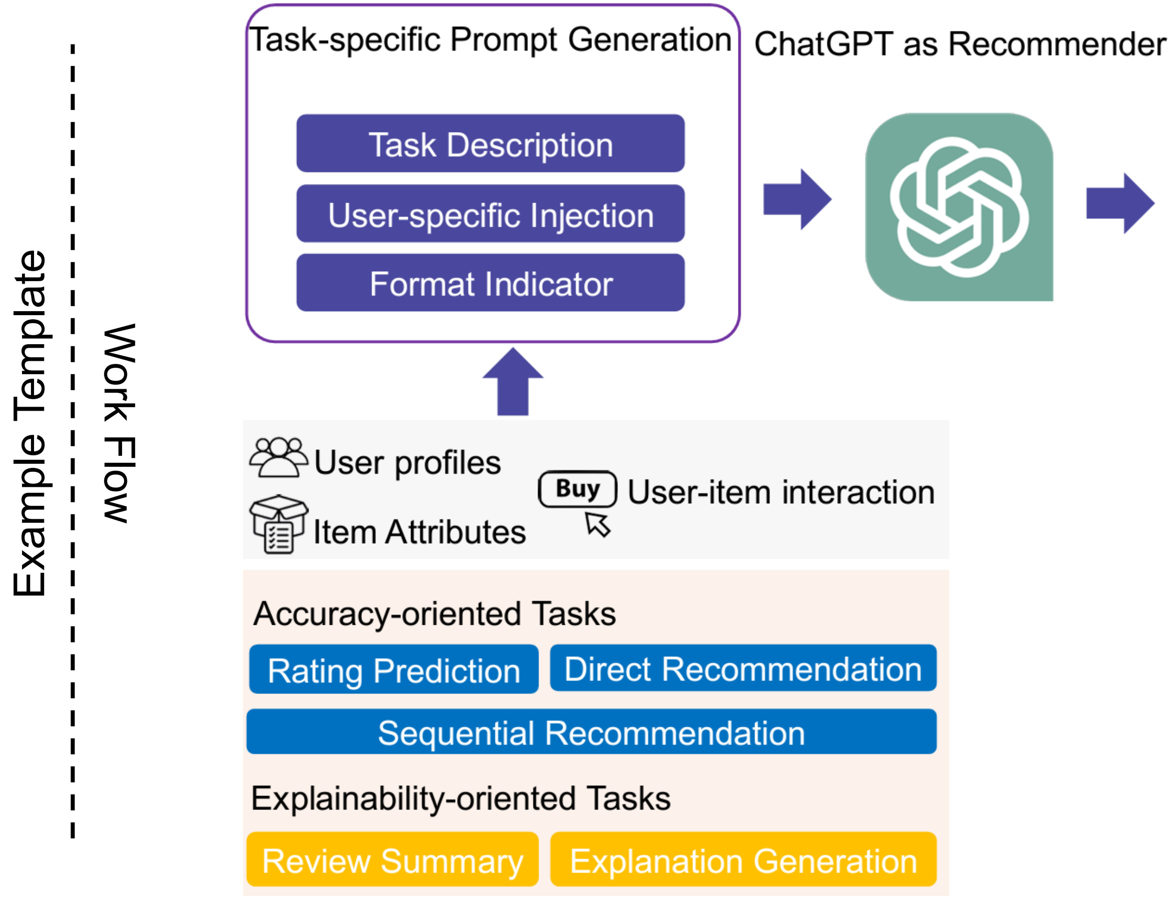

"Is ChatGPT a good recommender? a preliminary study." CIKM (2023).

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&reg, 5.0 ; 4.Nail Art Jumbo Stamp Stamping Manicure Image Plate 6 Happy Holidays by Cheeky&reg, 5.0; Based on above rating history, please predict user's rating for the product: "SHANY Nail Art Set (24 Famouse Colors Nail Art Polish, Nail Art Decoration)", (1 being lowest and5 being highest,The output should be like: (x stars, xx%), do not explain the reason.)

"Is ChatGPT a good recommender? a preliminary study." CIKM (2023).

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.

"BookGPT: A General Framework for Book Recommendation Empowered by Large Language Model." arXiv preprint arXiv:2305.15673 (2023).

Teach LLMs to Act as RecSys

- LLM-Rec

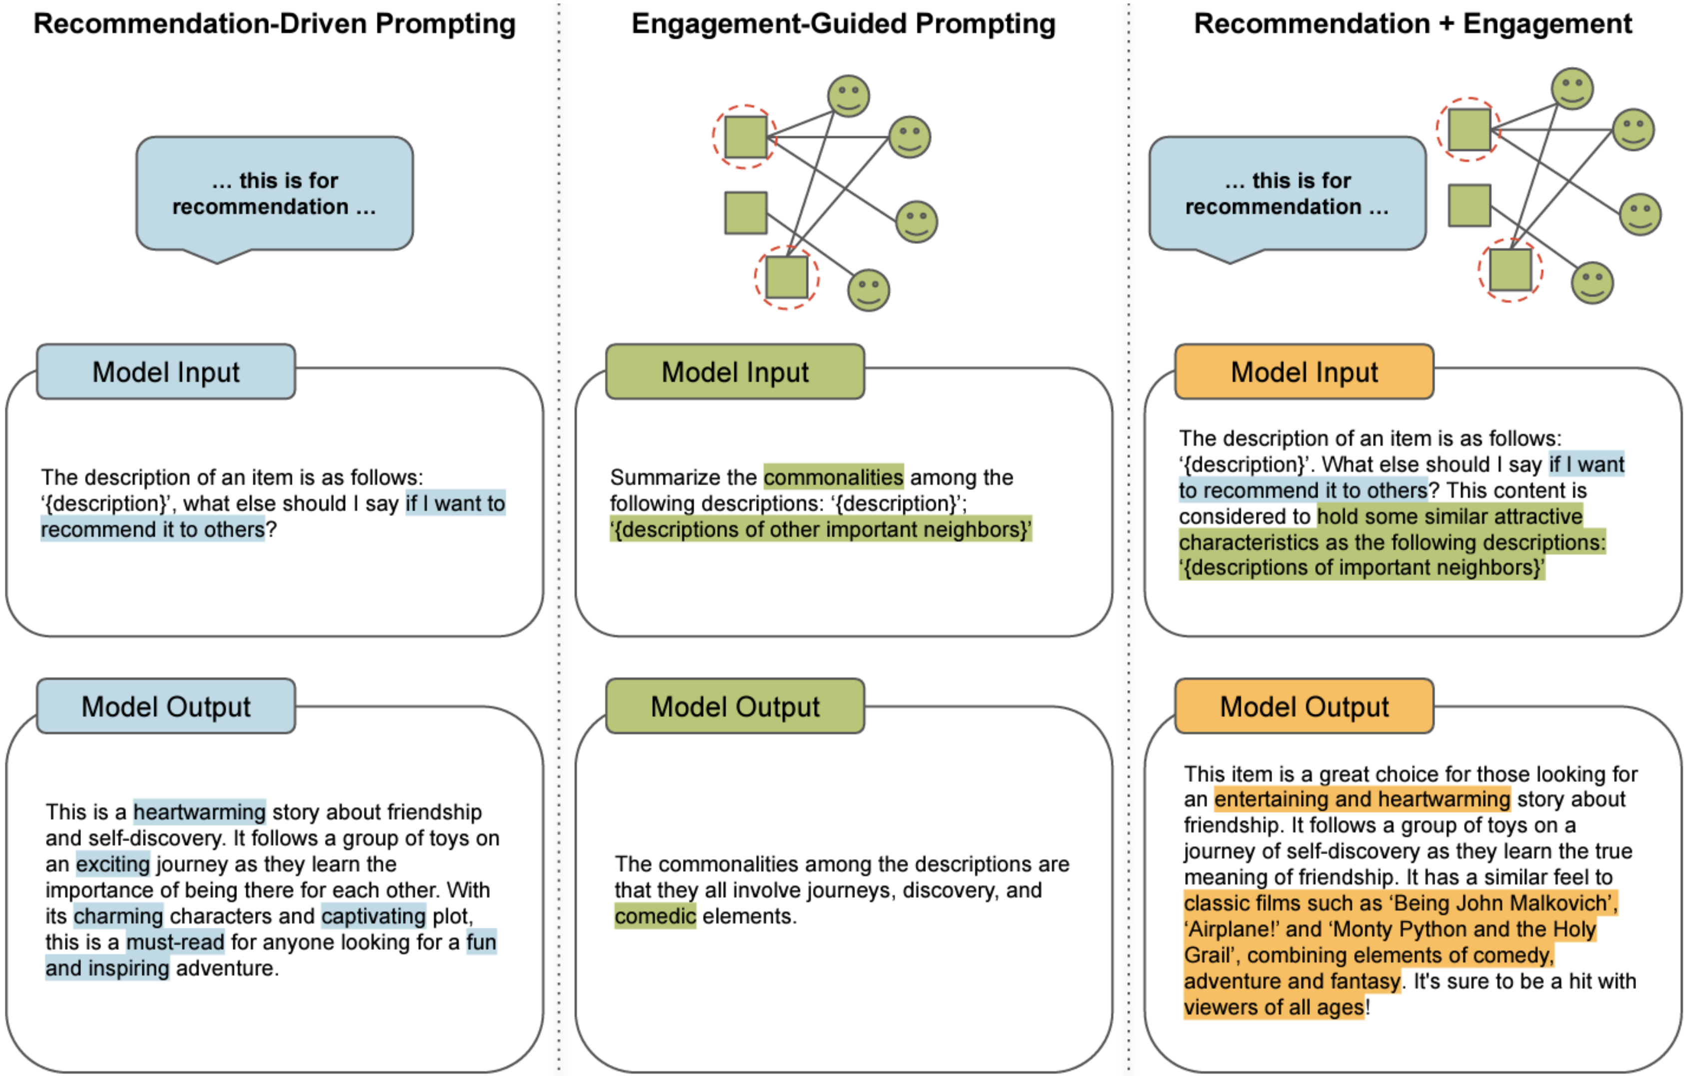

"Llm-rec: Personalized recommendation via prompting large language models." arXiv preprint arXiv:2307.15780 (2023).

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

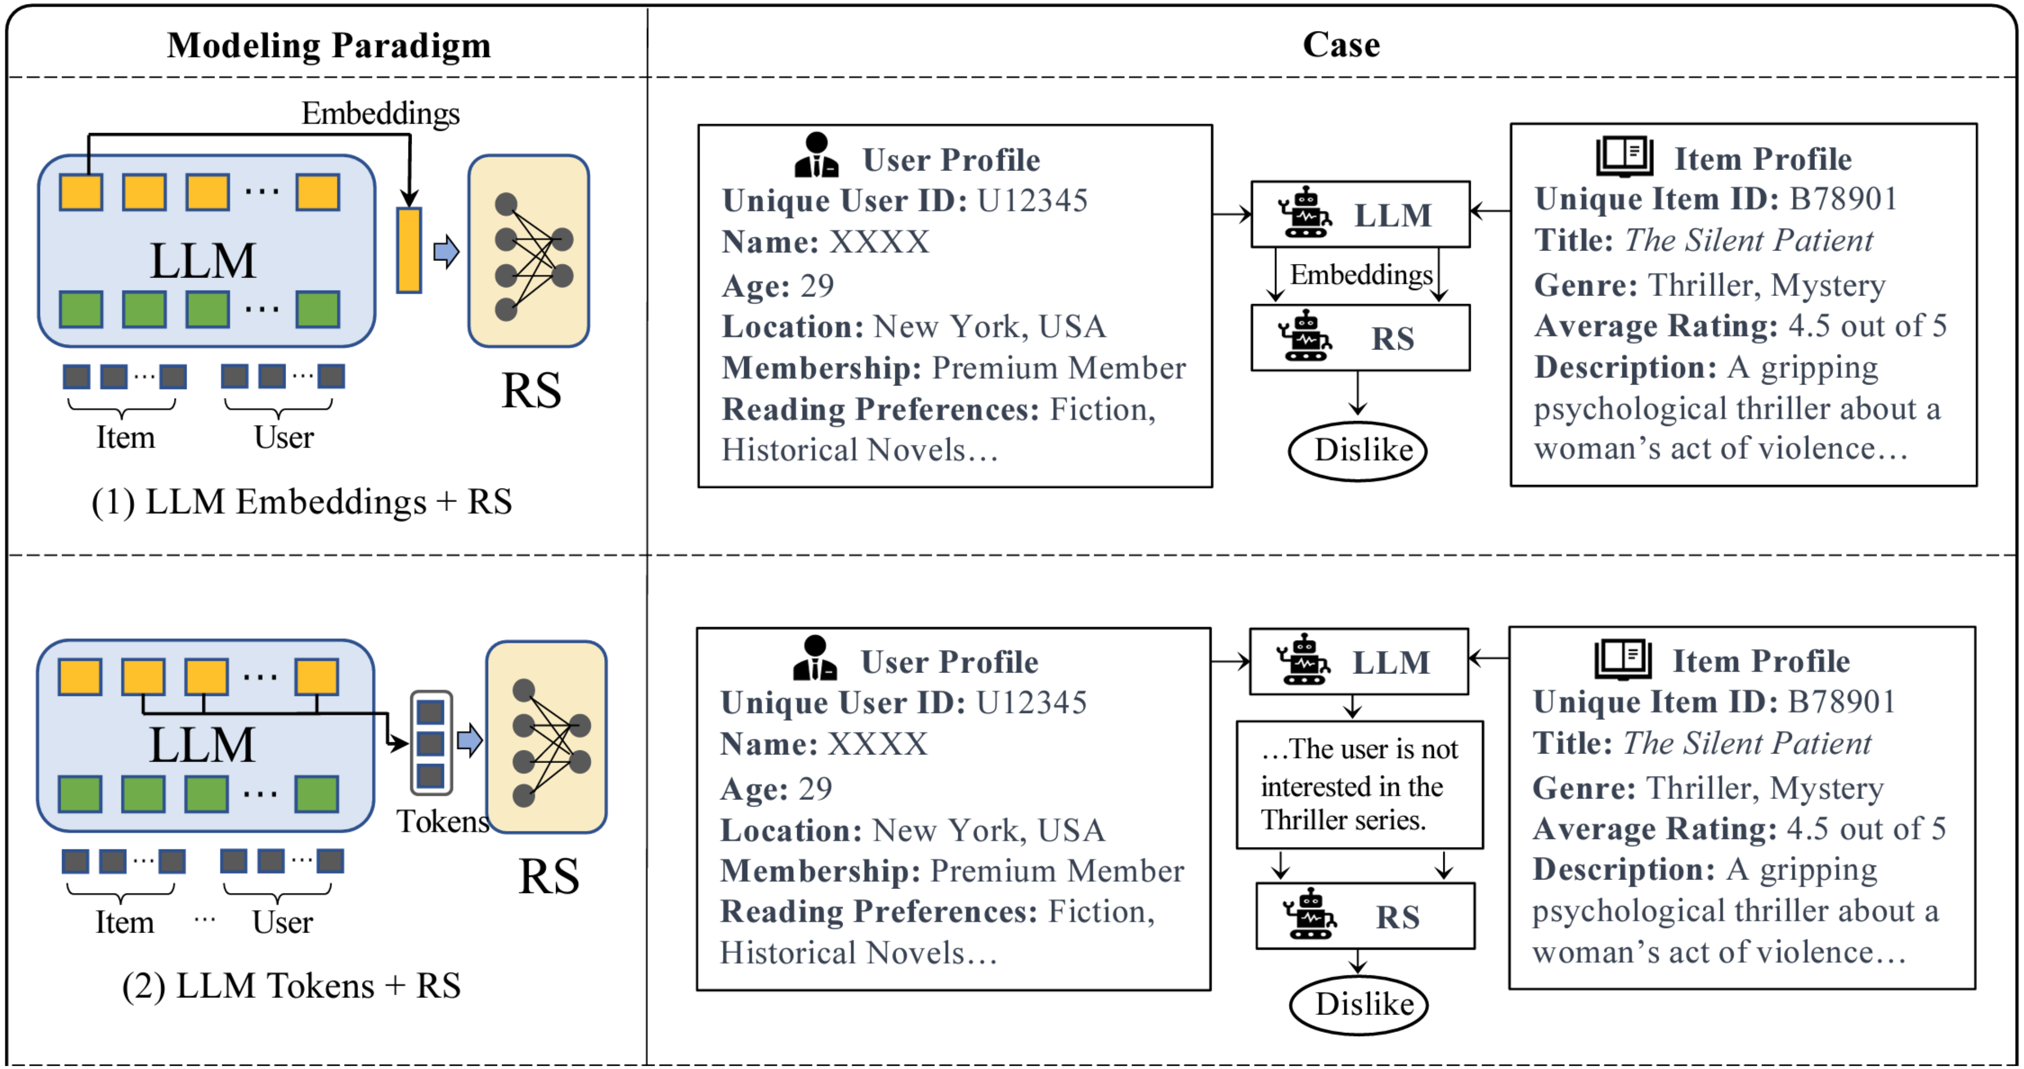

"A Survey on Large Language Models for Recommendation." arXiv preprint arXiv:2305.19860 (2023).

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**.

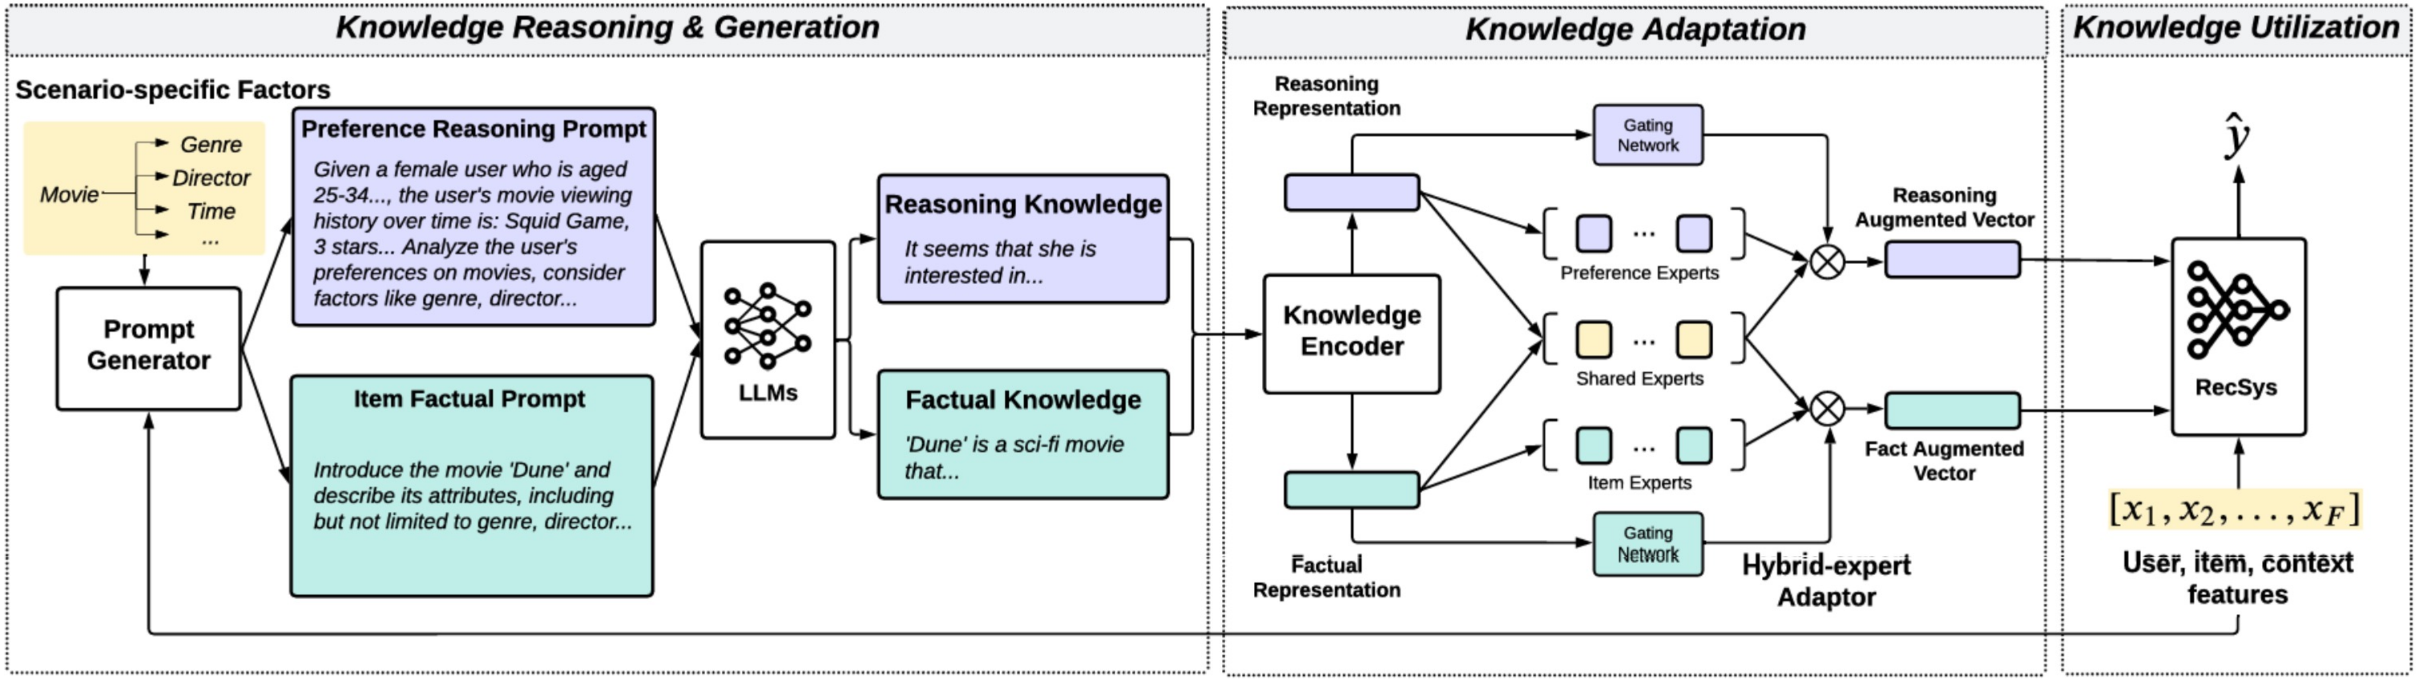

"Towards Open-World Recommendation with Knowledge Augmentation from Large Language Models." arXiv preprint arXiv:2306.10933 (2023).

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**.

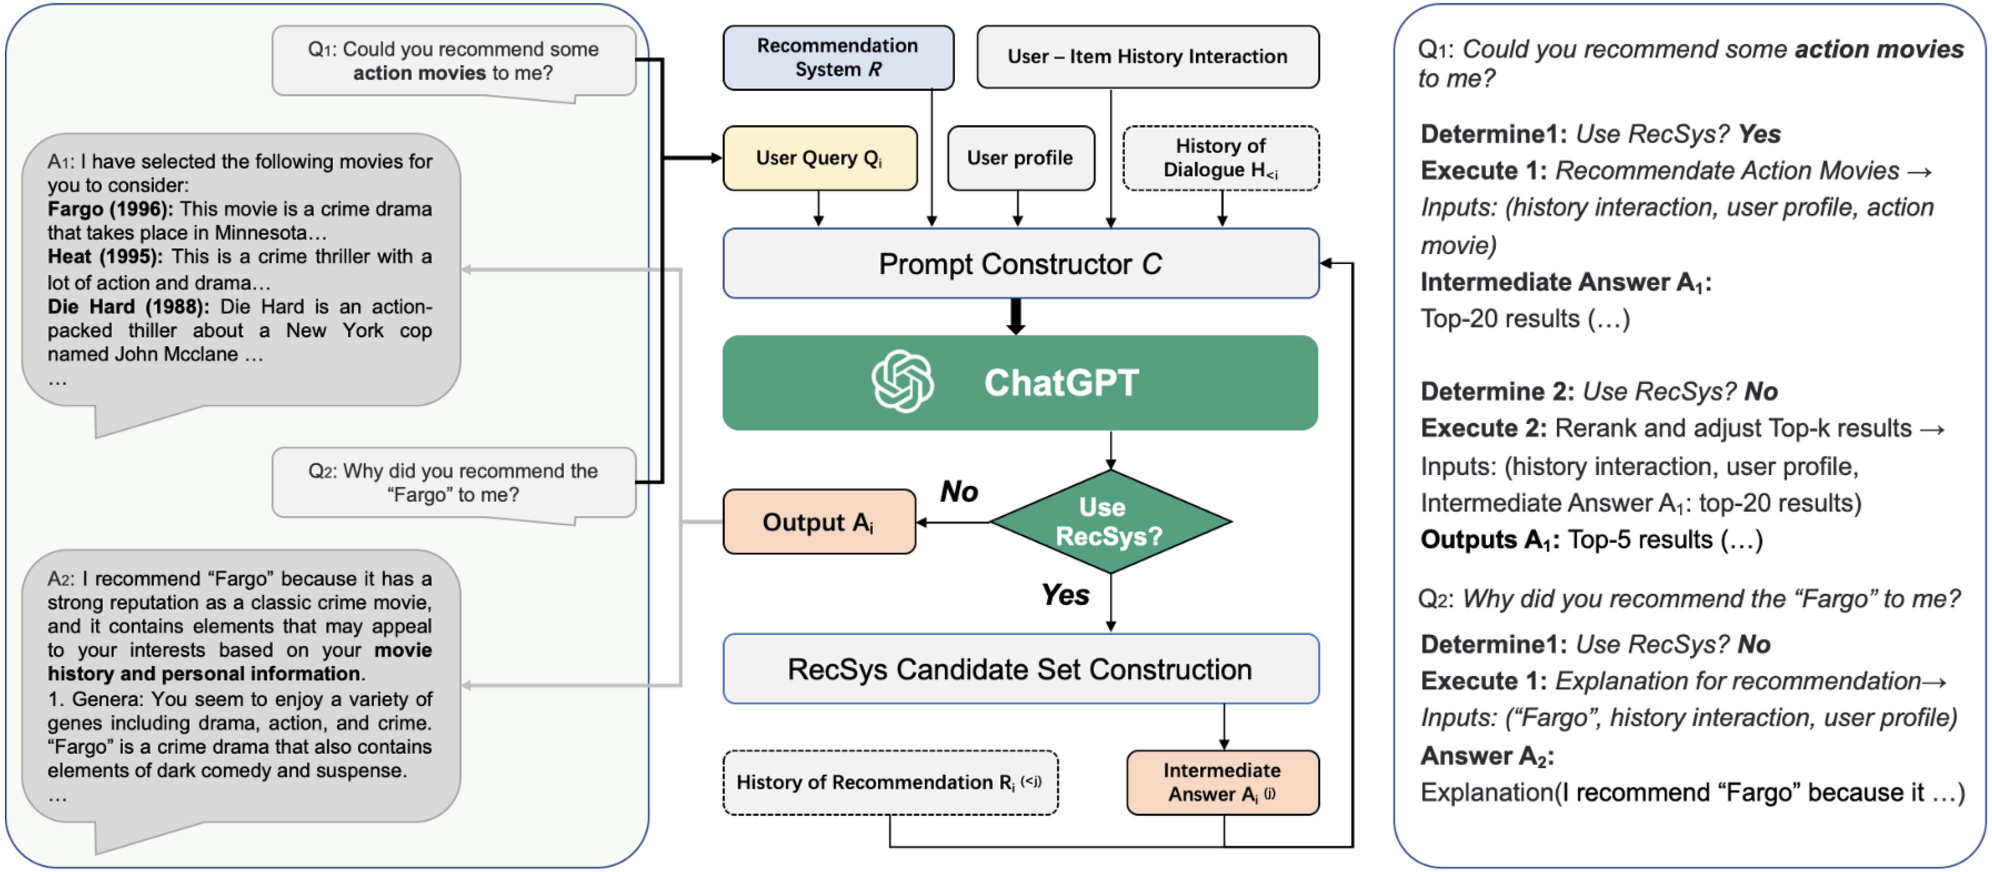

"Chat-rec: Towards interactive and explainable llms-augmented recommender system." arXiv preprint arXiv:2303.14524 (2023).

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.

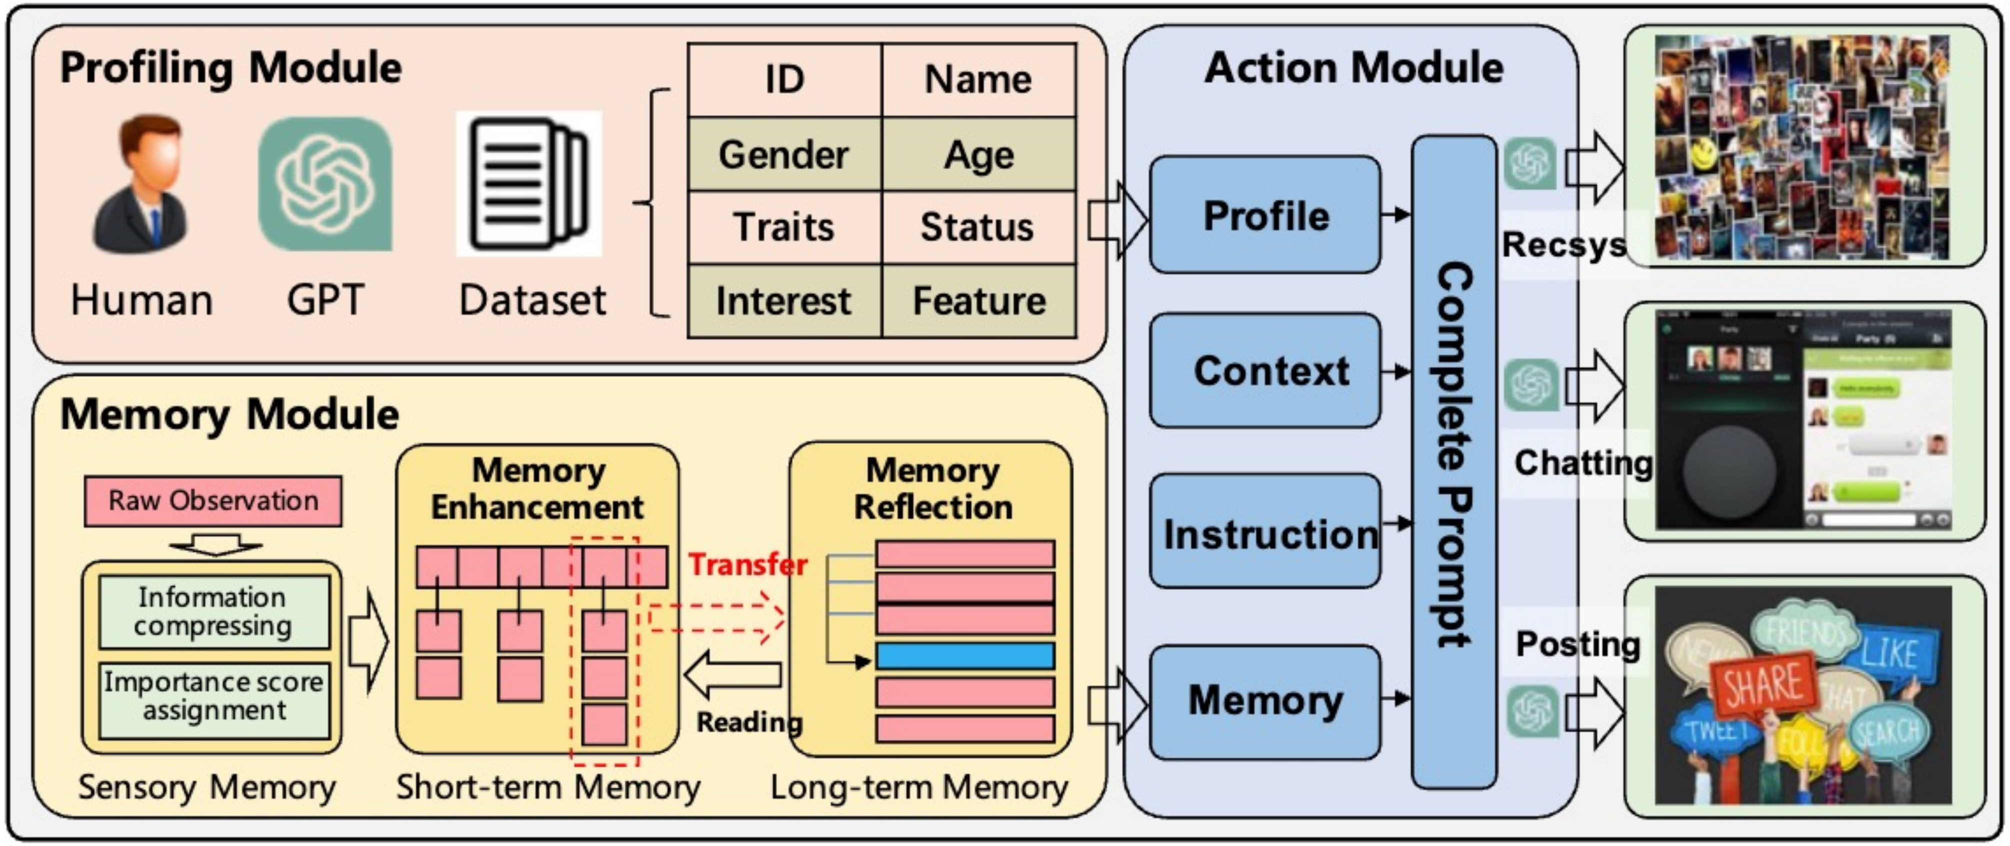

"RecAgent: A Novel Simulation Paradigm for Recommender Systems." arXiv preprint arXiv:2306.02552 (2023).

Act as Agent & Use External Tools

- Agent4Rec

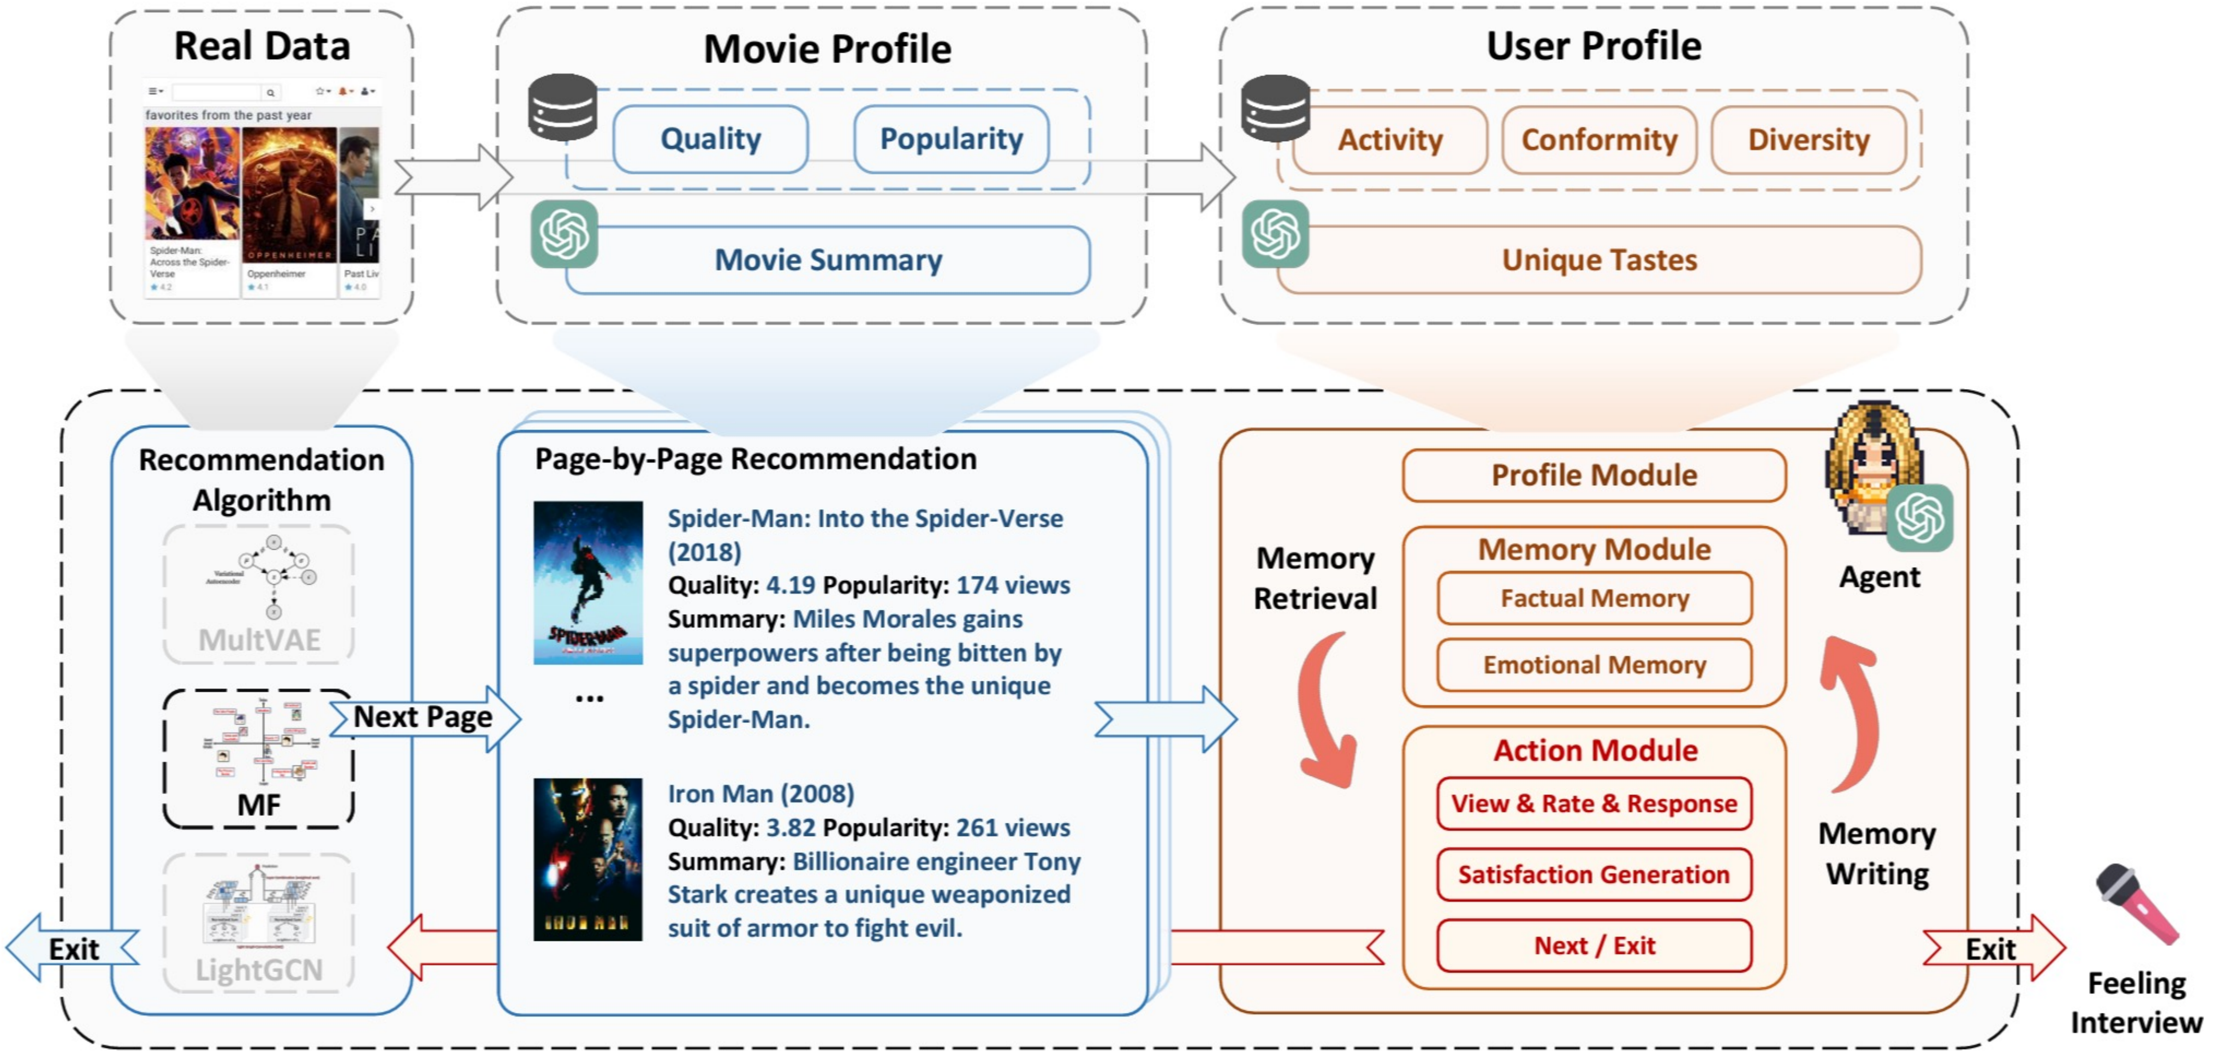

"On Generative Agents in Recommendation." arXiv preprint arXiv:2310.10108 (2023).

Act as Agent & Use External Tools

❑ Graph-ToolFormer

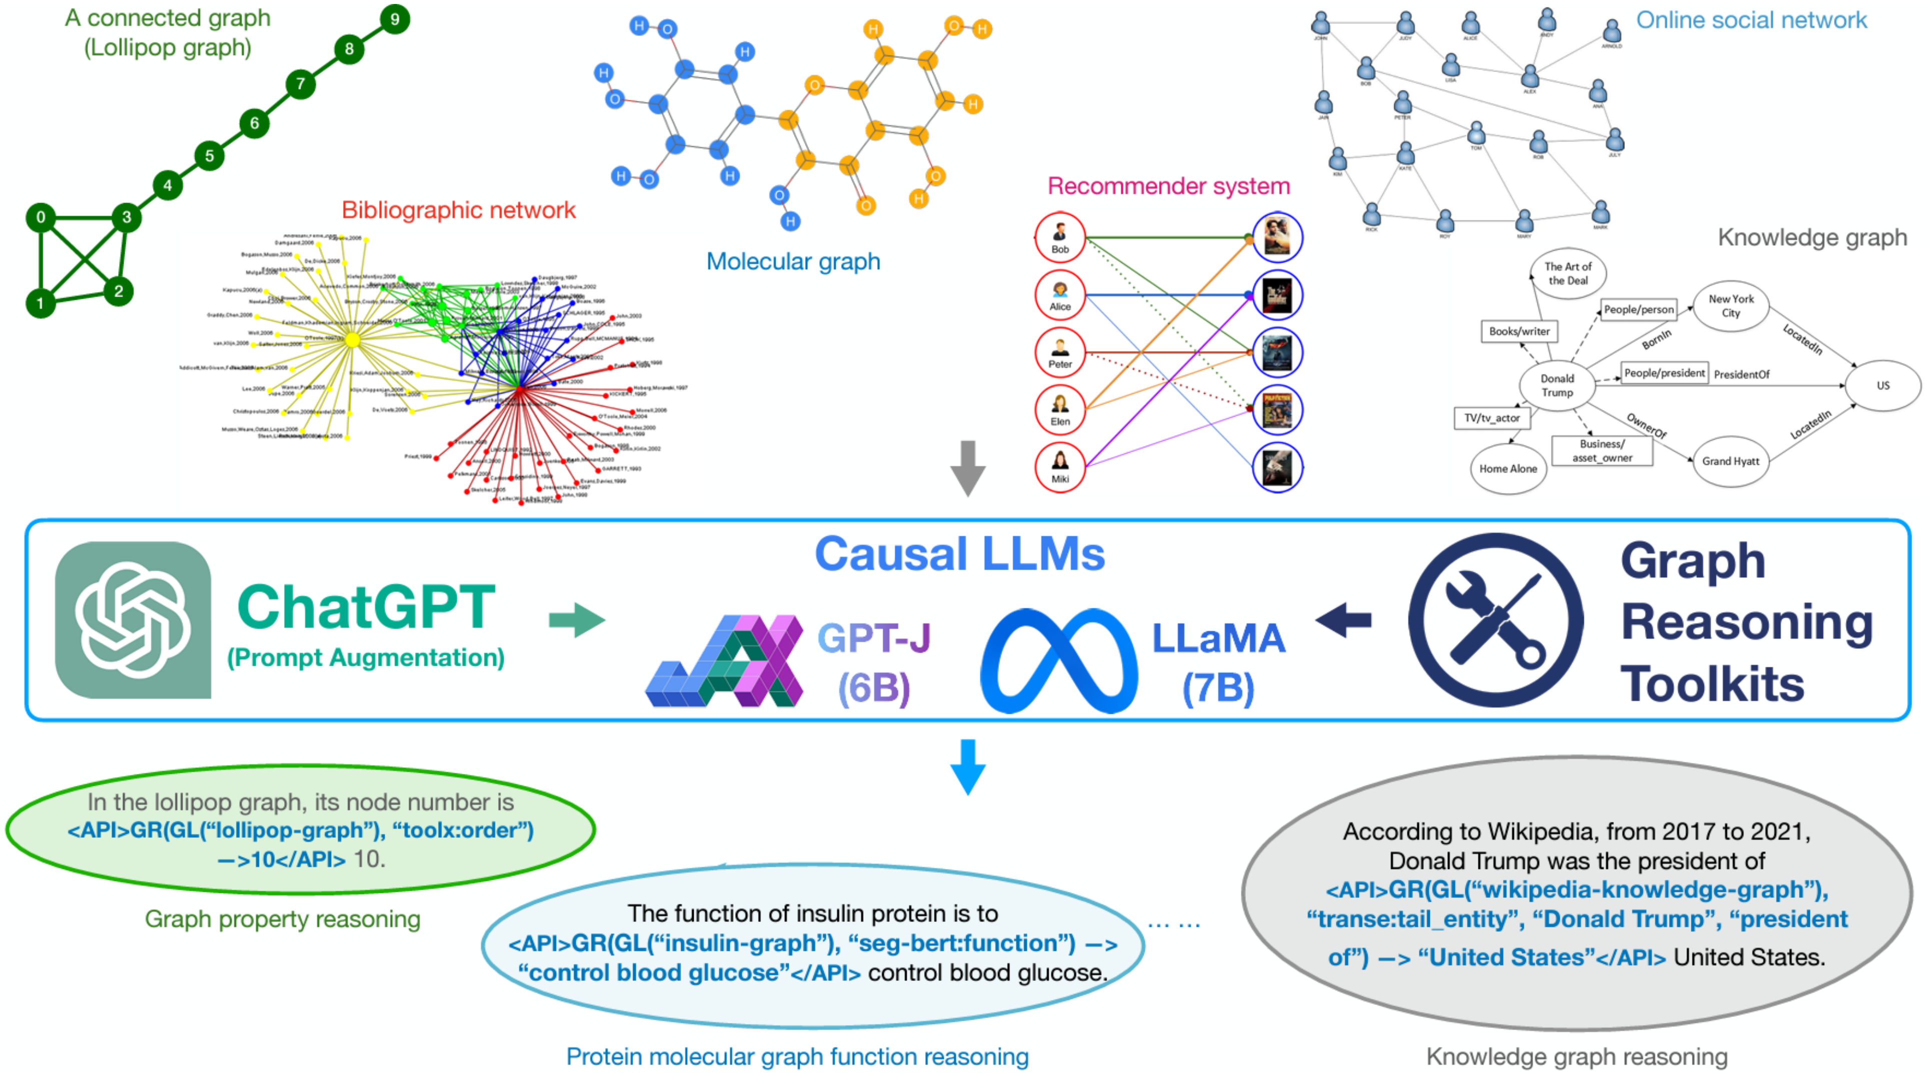

"Graph-ToolFormer: To Empower LLMs with Graph Reasoning Ability via Prompt Augmented by ChatGPT." arXiv preprint arXiv:2304.11116 (2023).

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.

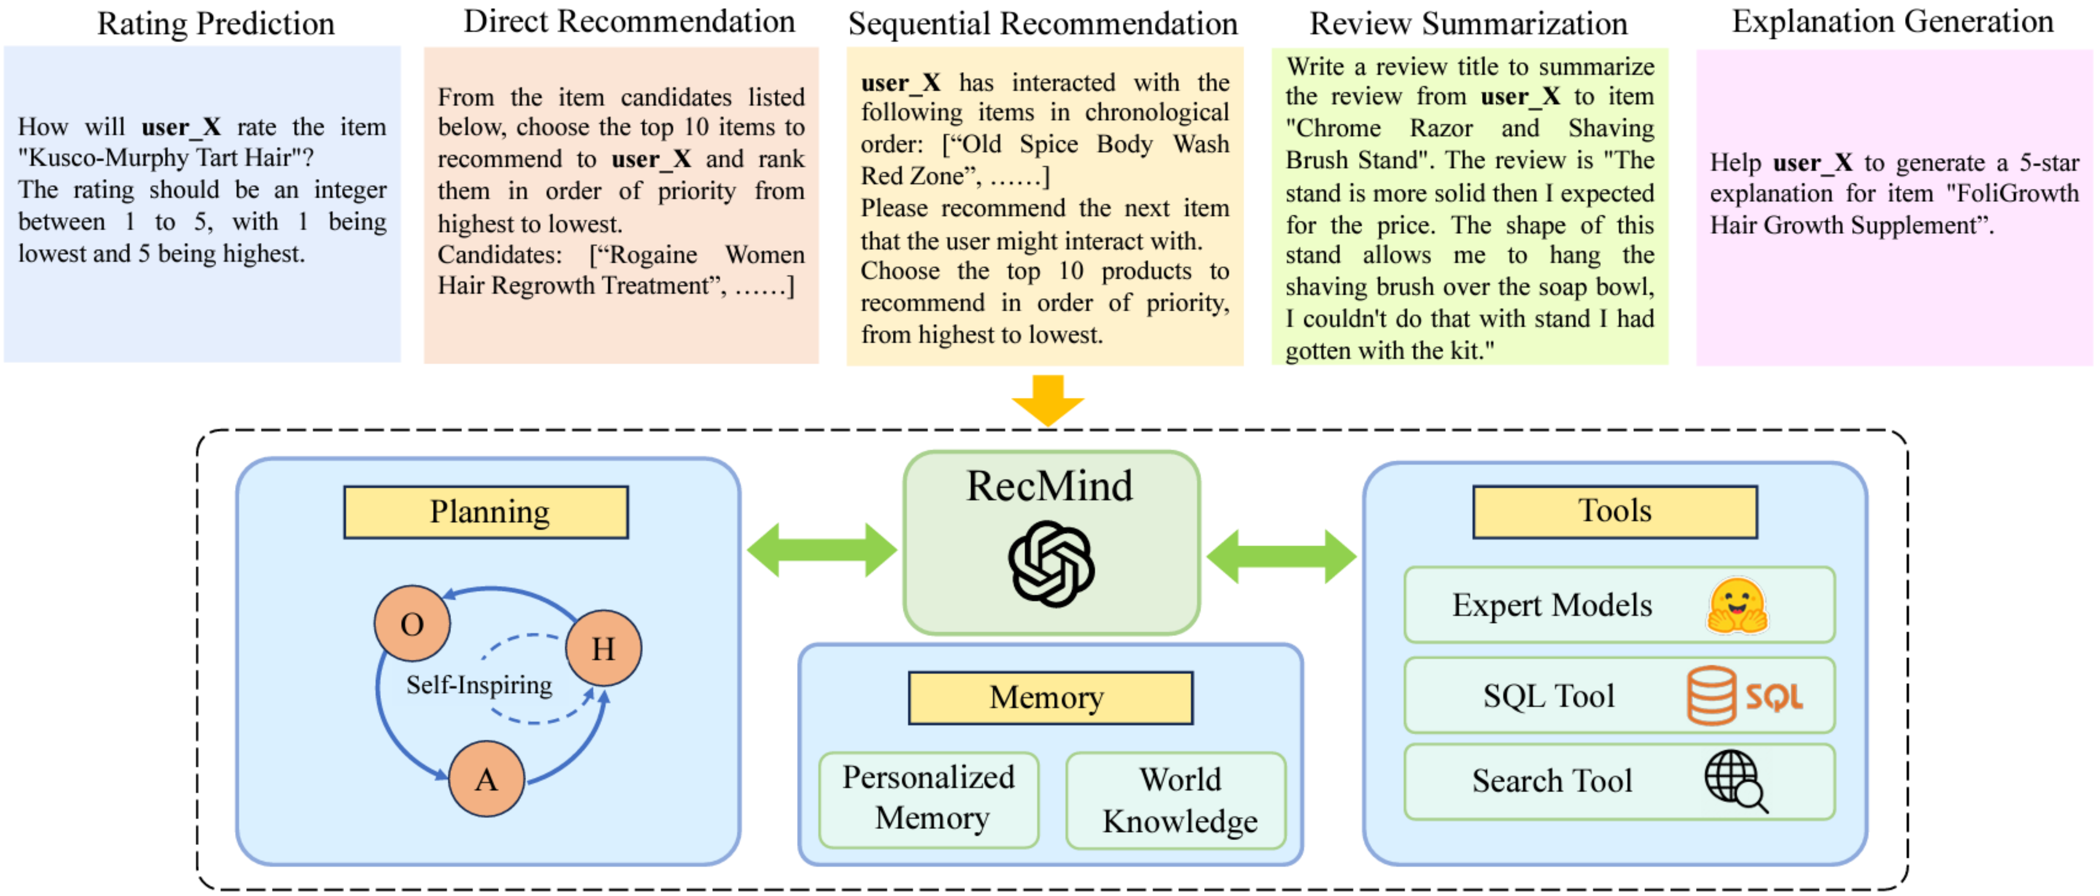

"Recmind: Large language model powered agent for recommendation." arXiv preprint arXiv:2308.14296 (2023).

Overview

Presenter:
Zihuai Zhao

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

Zoom ID:
91649466943
Password:
202312

Website QR Code

Chain-of-Thought (CoT) Prompting

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

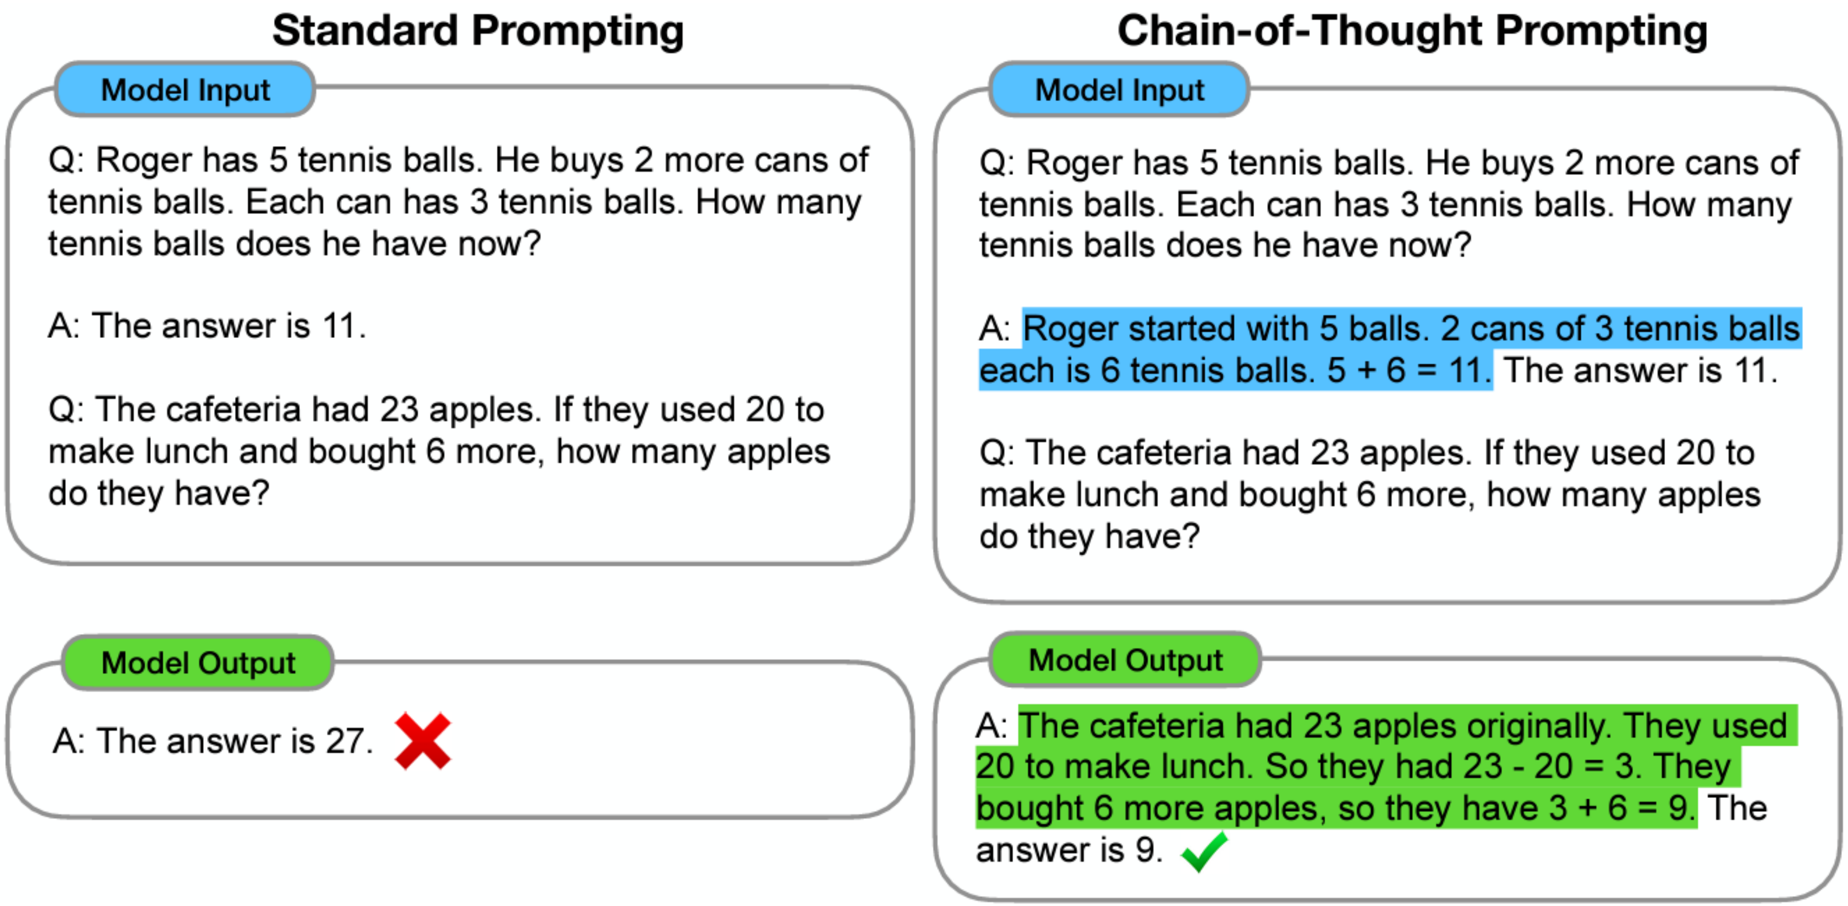

"Chain-of-thought prompting elicits reasoning in large language models." NeurIPS (2022).

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.

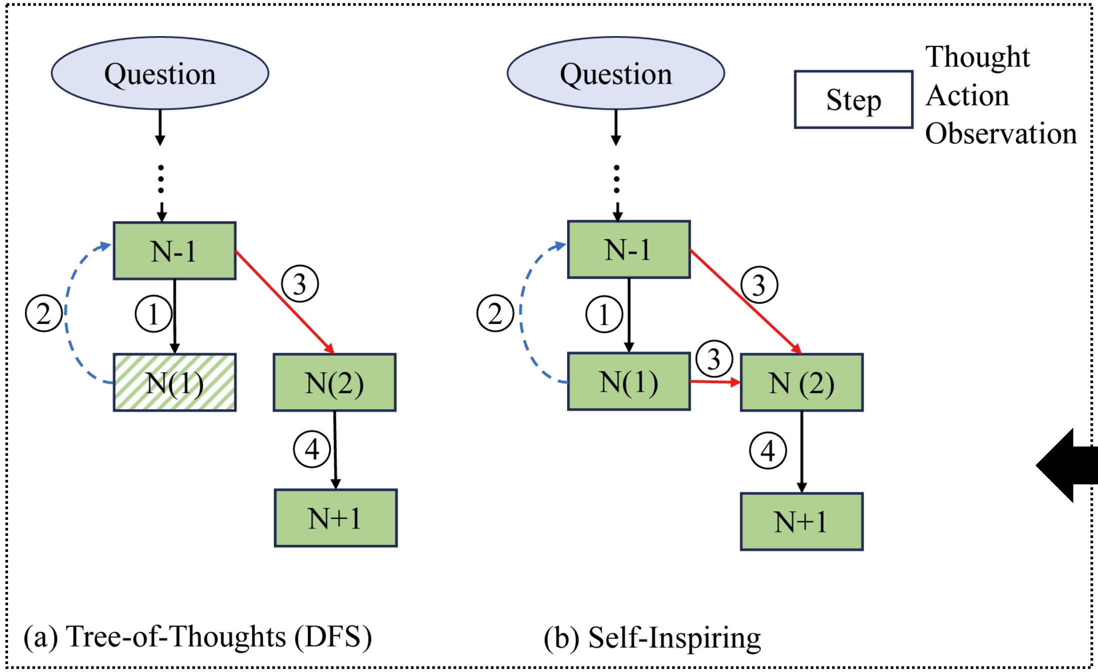

(a) Tree-of-Thoughts (DFS)

(b) Self-Inspiring

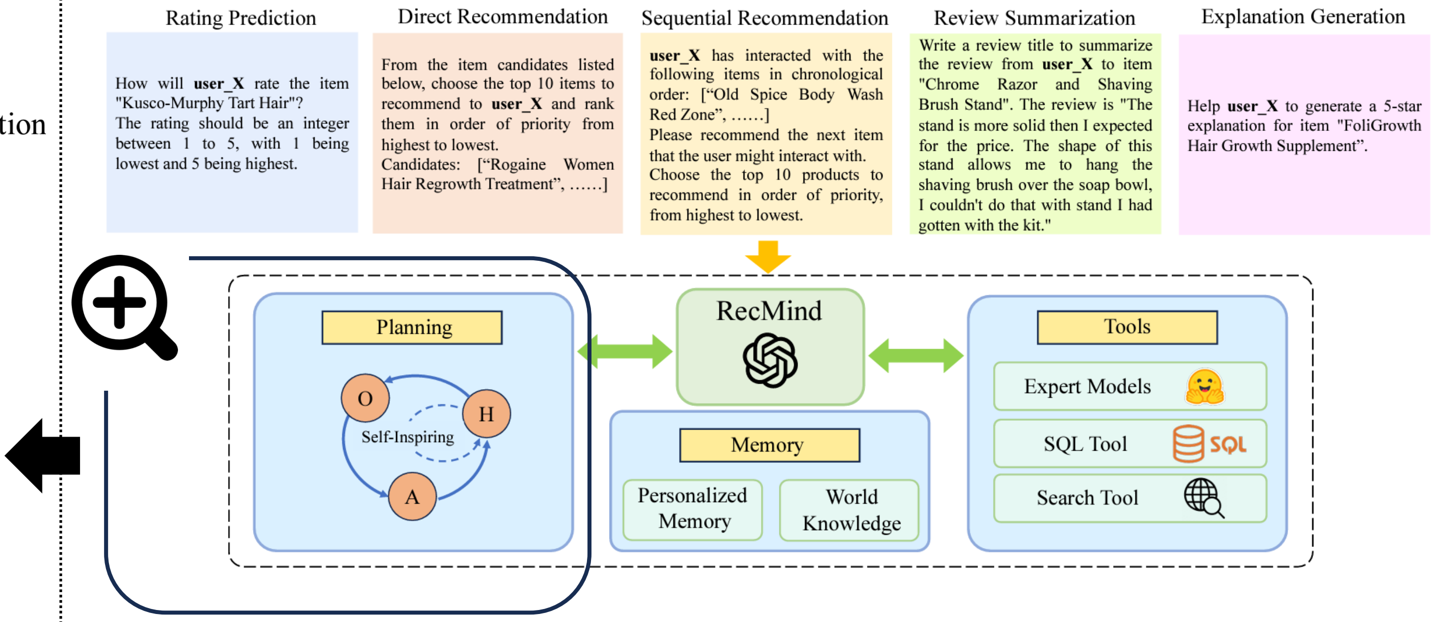

"Recmind: Large language model powered agent for recommendation." arXiv preprint arXiv:2308.14296 (2023).

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.

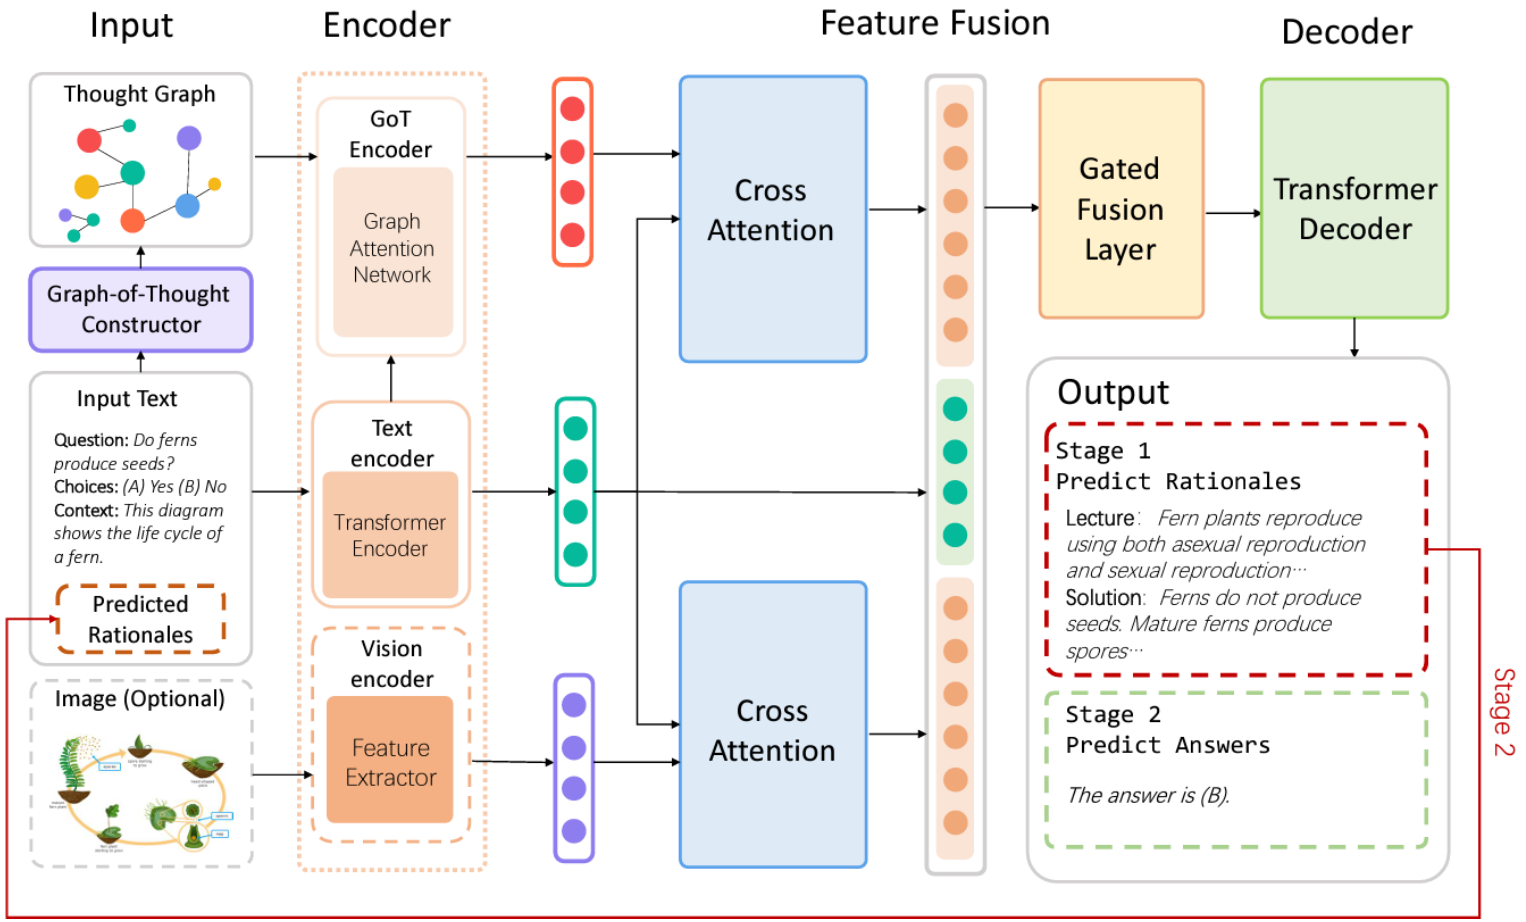

"Beyond Chain-of-Thought, Effective Graph-of-Thought Reasoning in Large Language Models." arXiv preprint arXiv:2305.16582 (2023).

Overview

Presenter:
Zihuai Zhao

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

Zoom ID:
91649466943
Password:
202312

Website QR Code

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**.

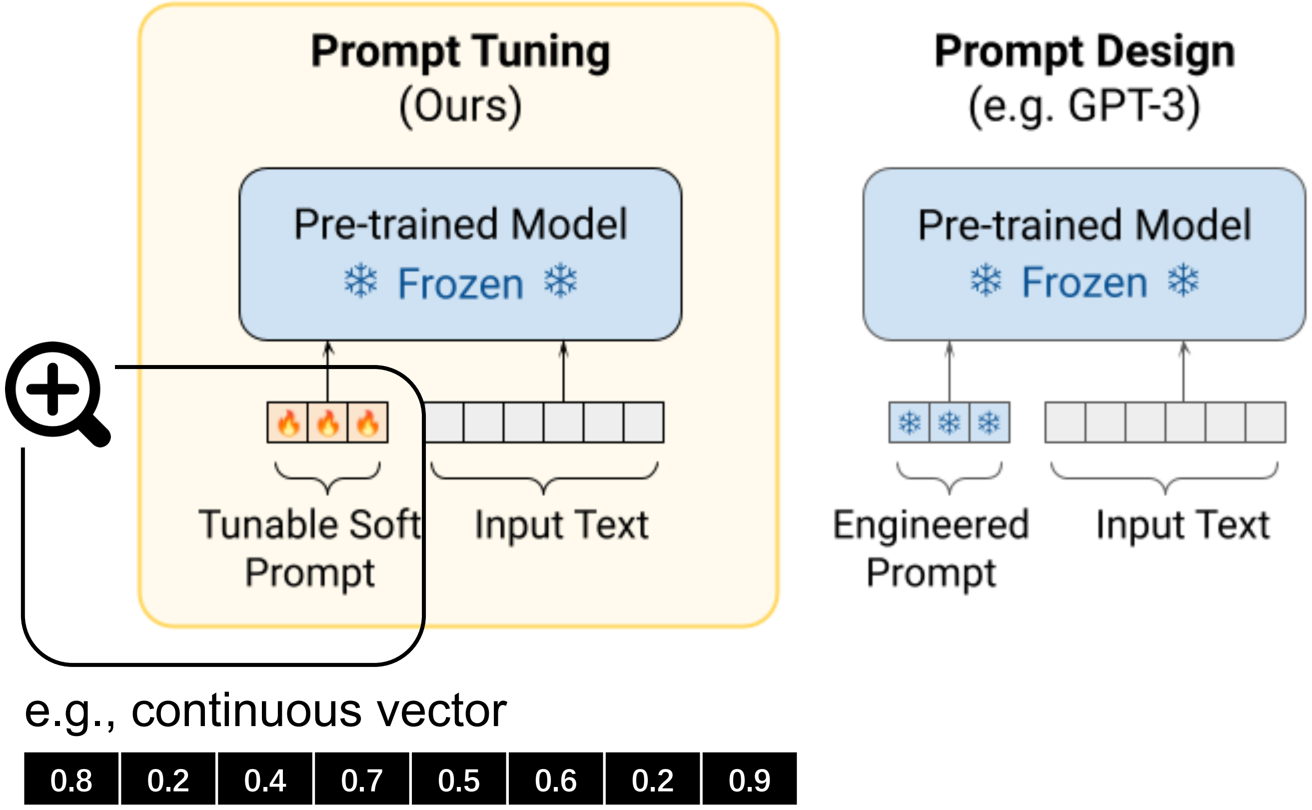

"The Power of Scale for Parameter-Efficient Prompt Tuning" EMNLP (2021).

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.”

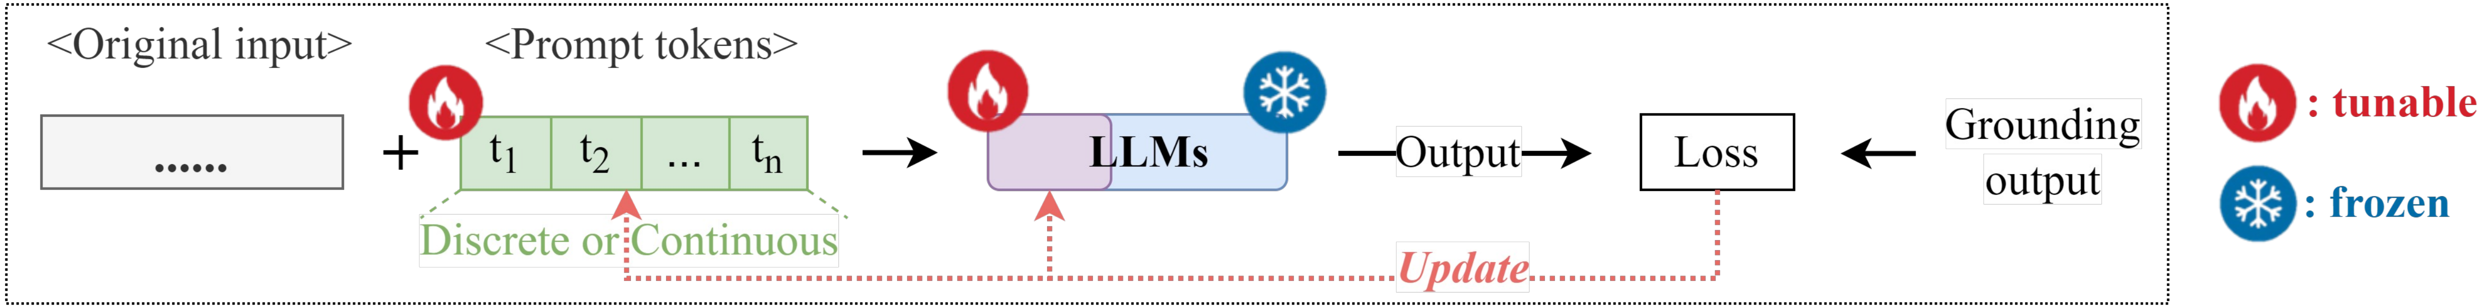

"A survey for in-context learning." arXiv preprint arXiv:2301.00234 (2022).

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.

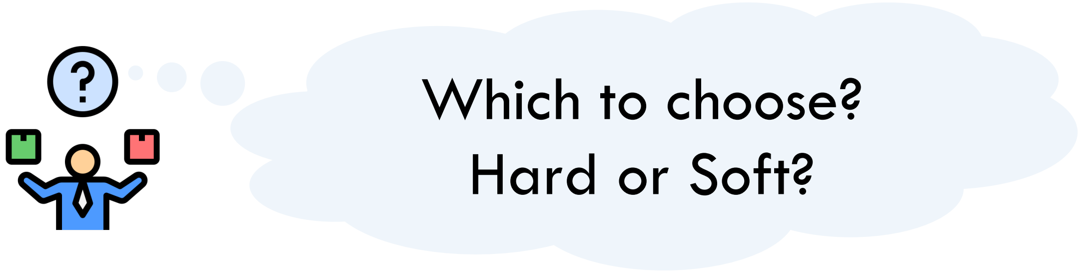

Prompt Tuning in RecSys

❑ 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.

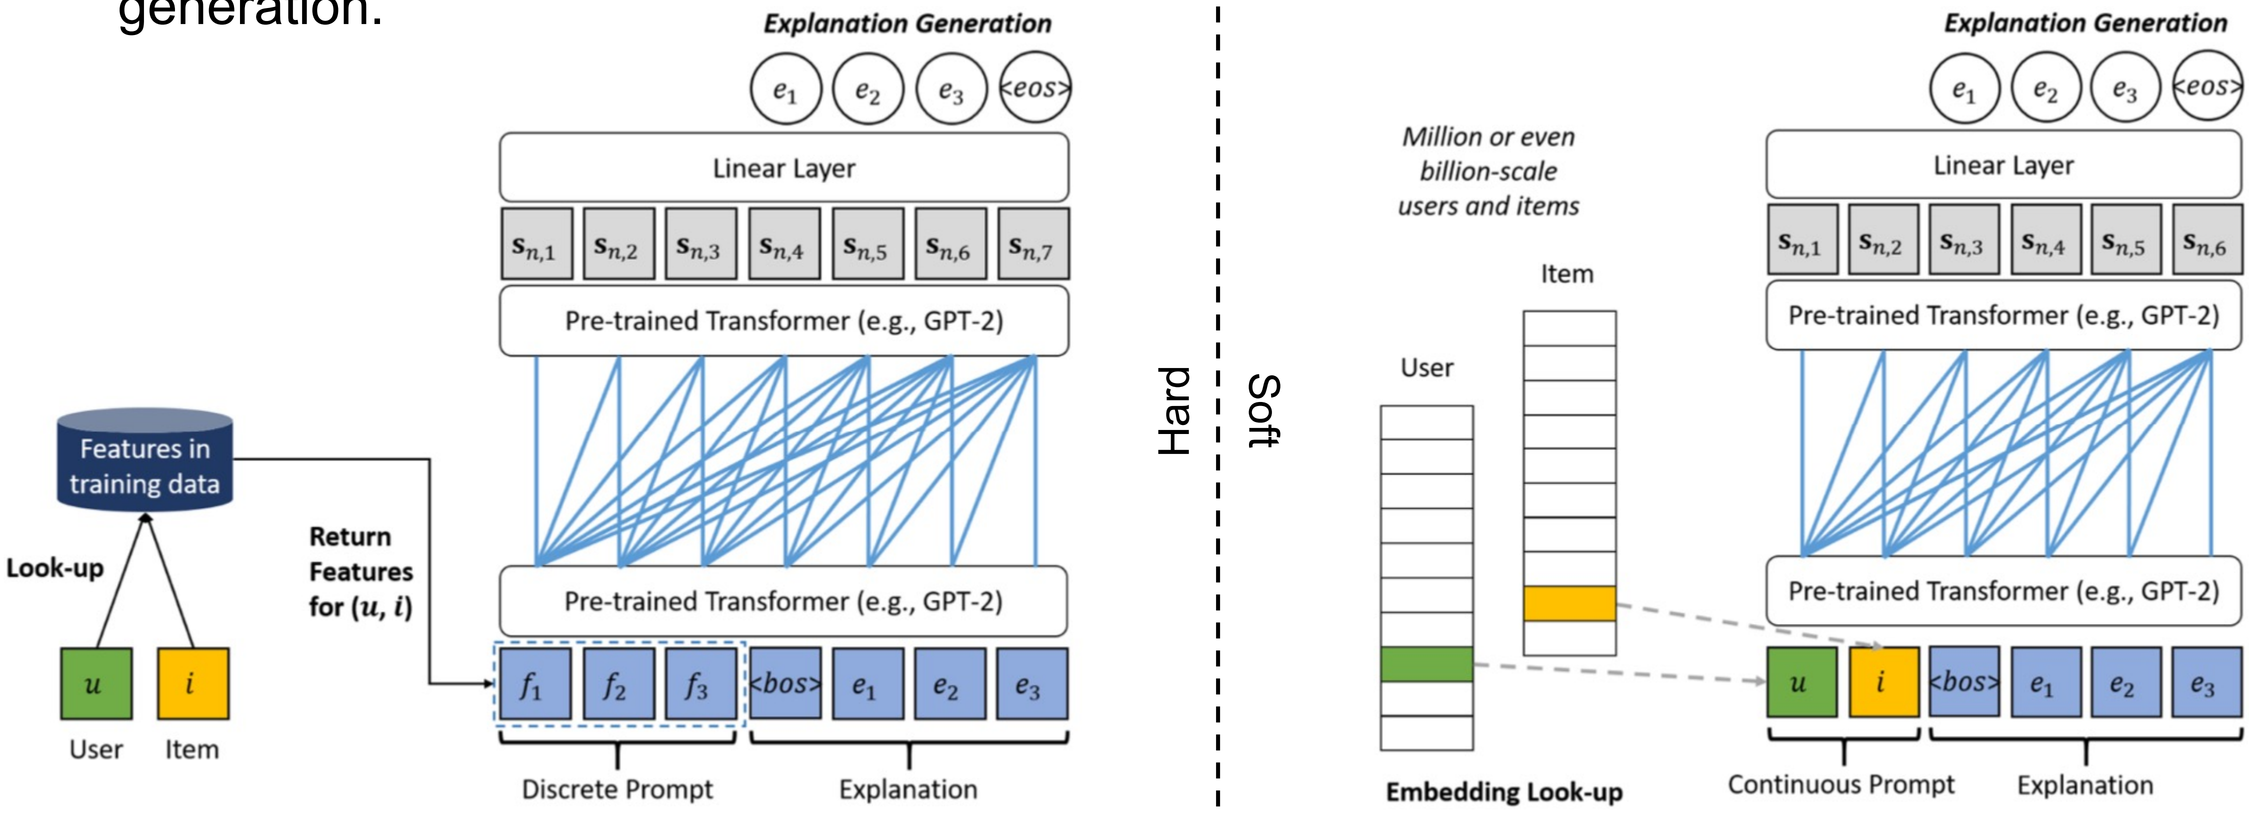

"Personalized prompt learning for explainable recommendation." ACM Transactions on Information Systems (2023).

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**.

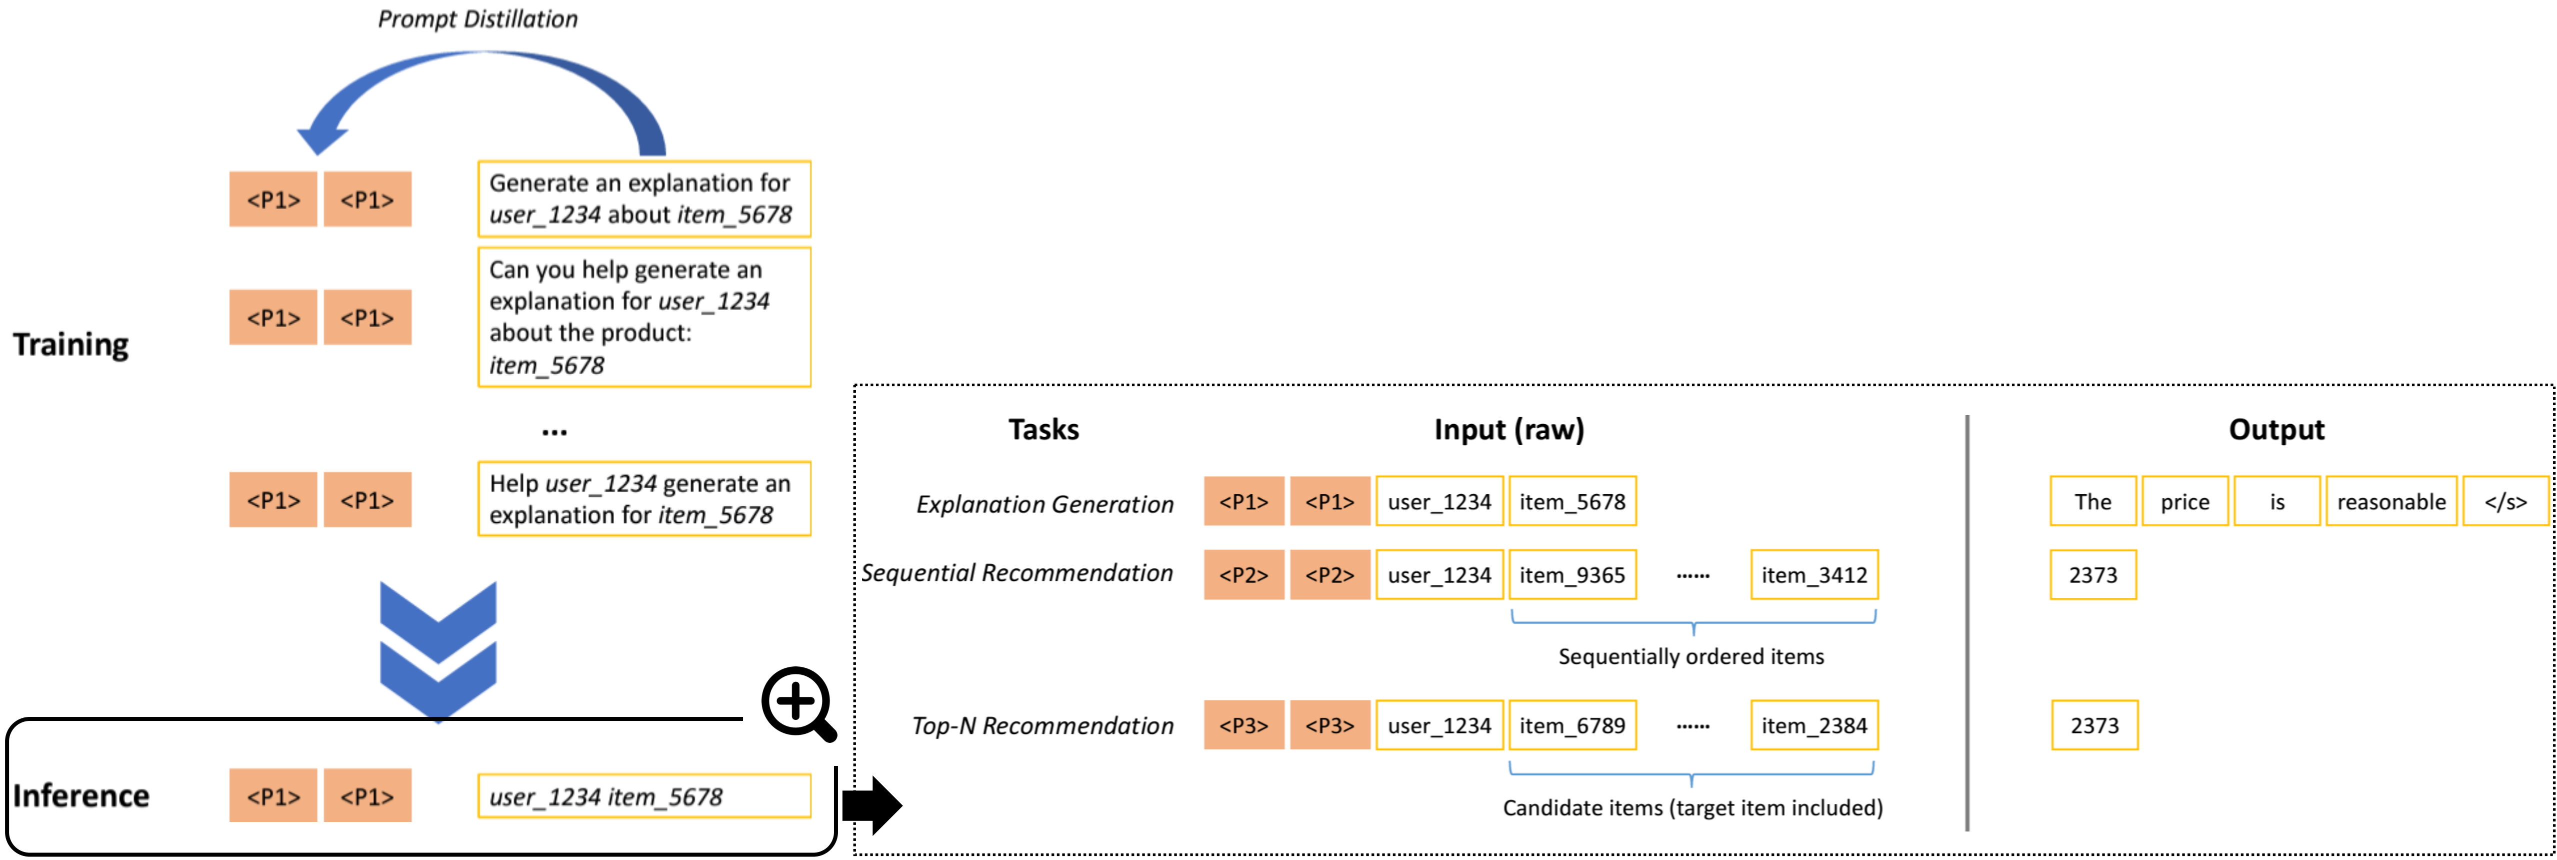

“Prompt distillation for efficient llm-based recommendation.“ CIKM (2023).

Preliminaries | Pre-training | Fine-tuning | **Prompting** | Future Directions

Overview

Presenter: Zihuai Zhao

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

Zoom ID: 91649466943
Password: 202312

Website QR Code

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.

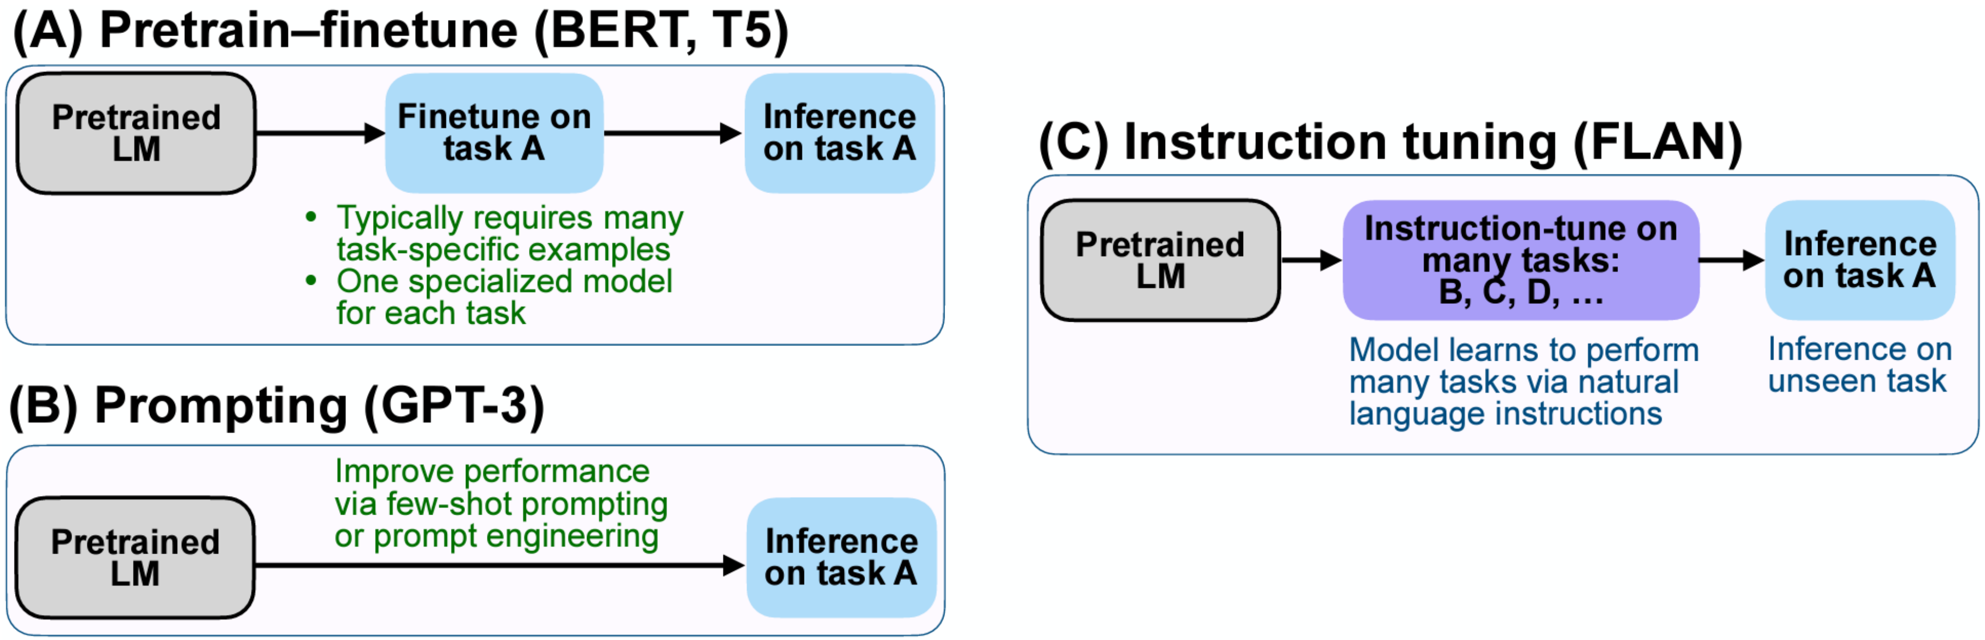

Stages of Instruction Tuning

Instruction Generation
(≈ Prompting)

Model Tuning with Prompt
(≈ 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.

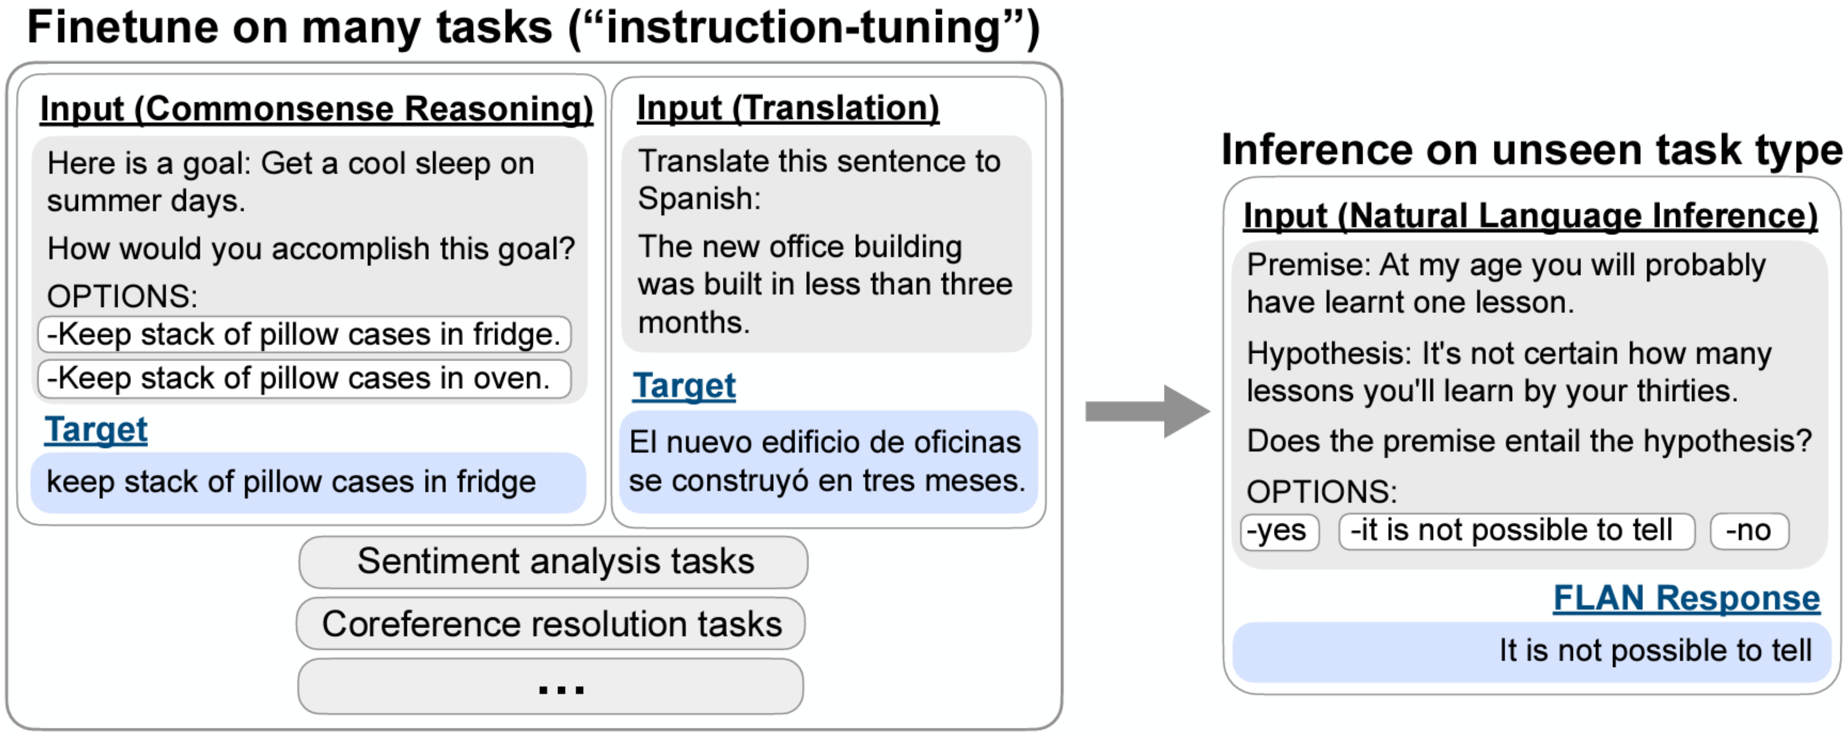

"Finetuned language models are zero-shot learners." ICLR (2022)

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.

"Recommendation as instruction following: A large language model empowered recommendation approach." arXiv preprint arXiv:2305.07001 (2023).

Stages of Instruction Tuning

Instruction Generation (≈ Prompting)

Model Tuning with Prompt (≈ 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

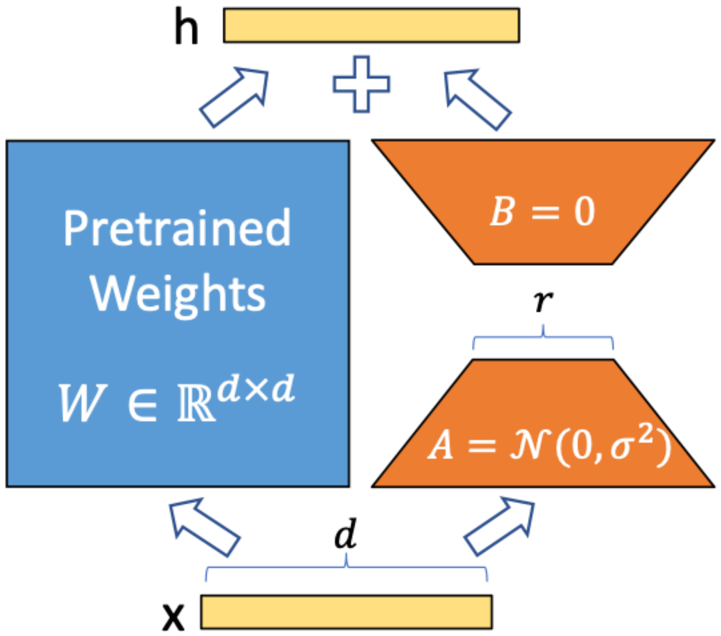

"Lora: Low-rank adaptation of large language models." ICLR (2022).

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.

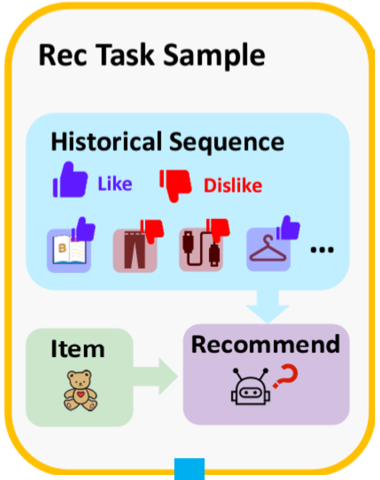

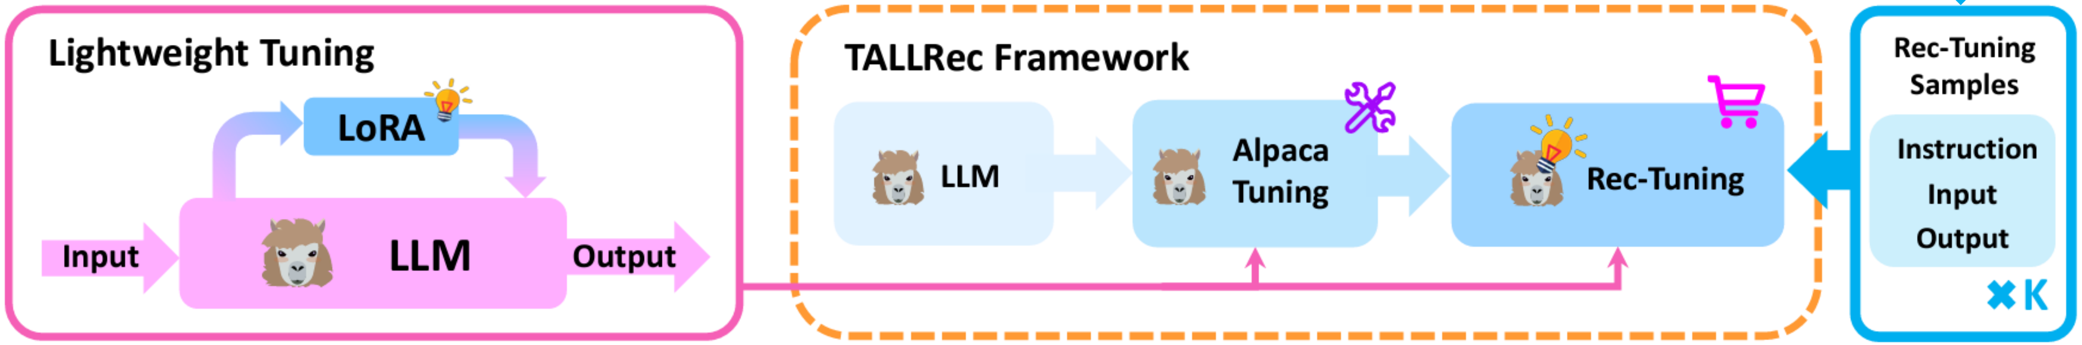

"Tallrec: An effective and efficient tuning framework to align large language model with recommendation." RecSys (2023).

Overview

Presenter:
Yiqi Wang

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

Zoom ID:
91649466943
Password:
202312

Website QR Code

Hallucination Mitigation

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

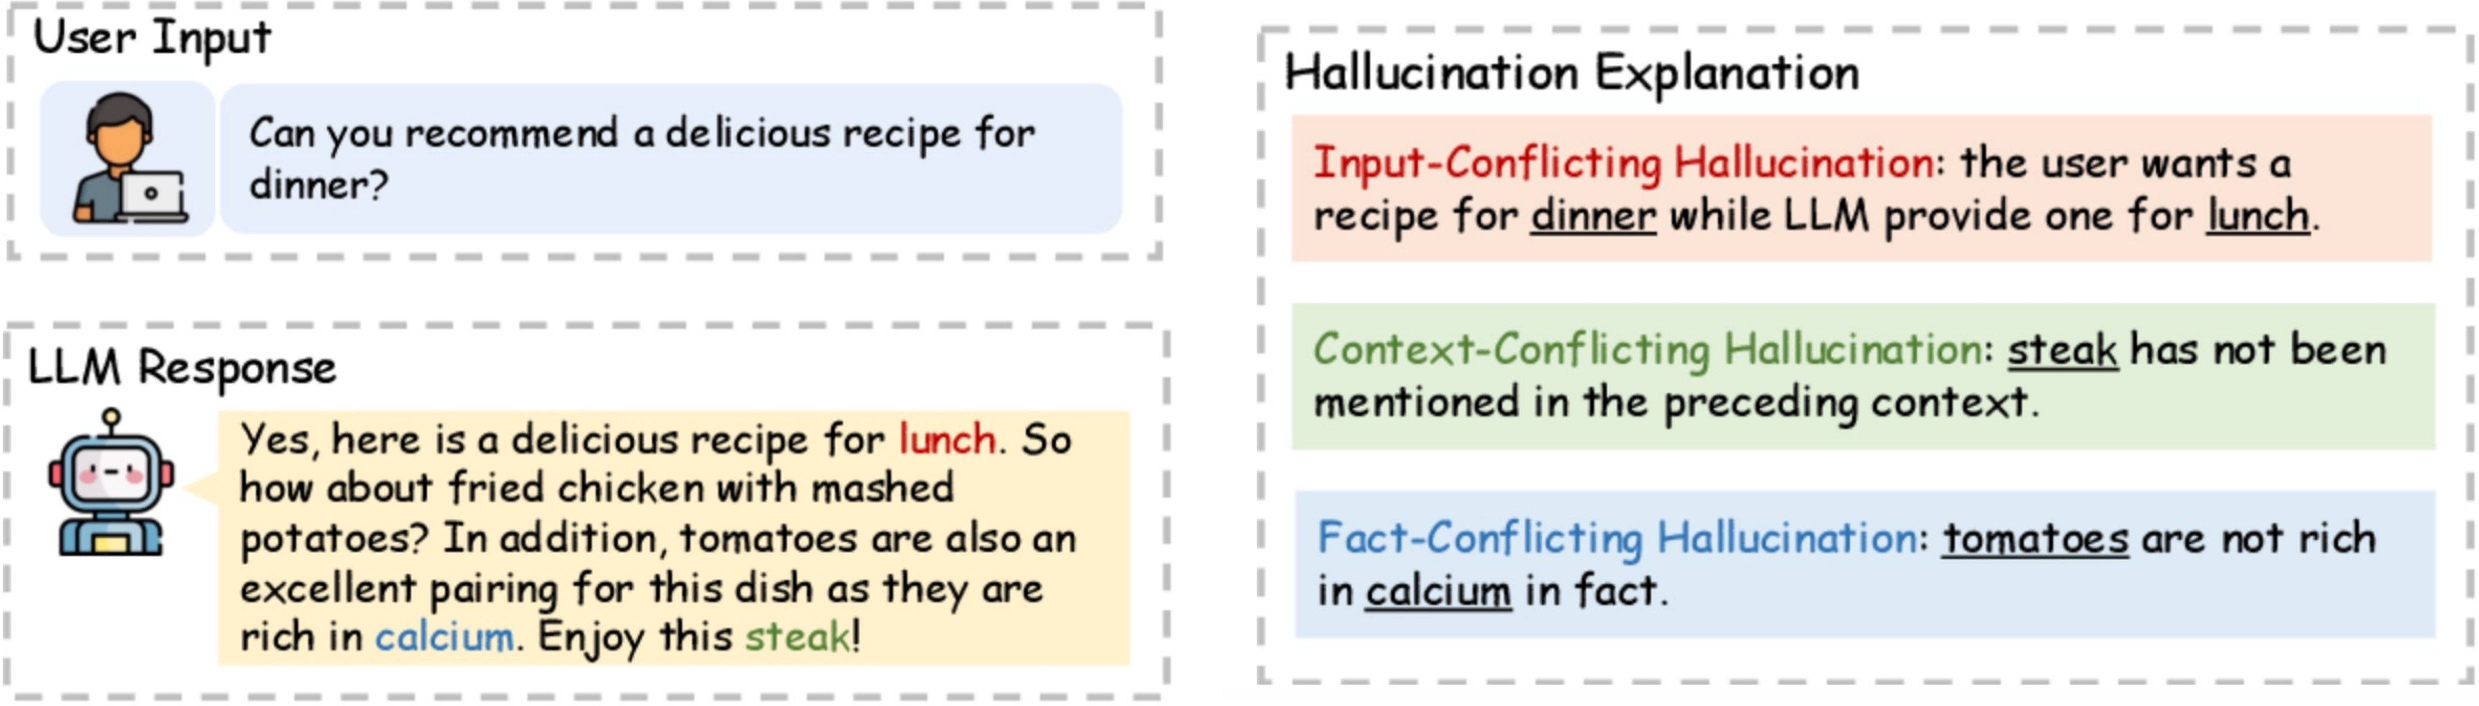

"Siren's Song in the AI Ocean: A Survey on Hallucination in Large Language Models." *ArXiv* (2023)

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
  - Safety and Robustness
  - Non-discrimination and Fairness
  - Explainability
  - Privacy

“Trustworthy ai: A computational perspective.” *ACM Transactions on Intelligent Systems and Technology* (2022)

Safety and Robustness

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

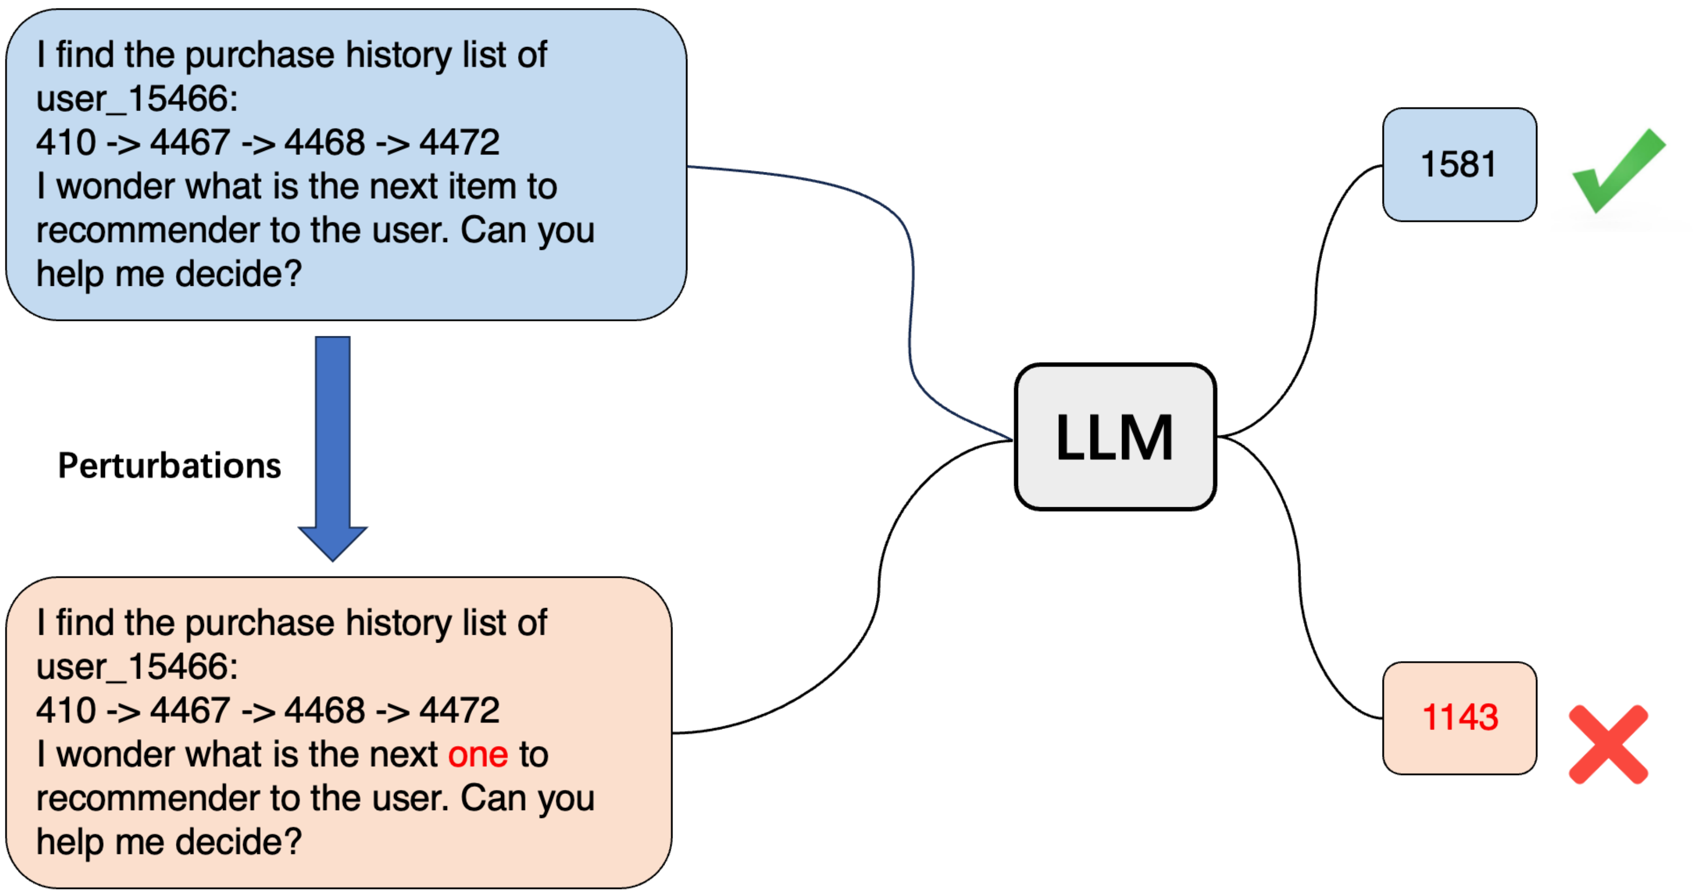

Self-Denoise

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

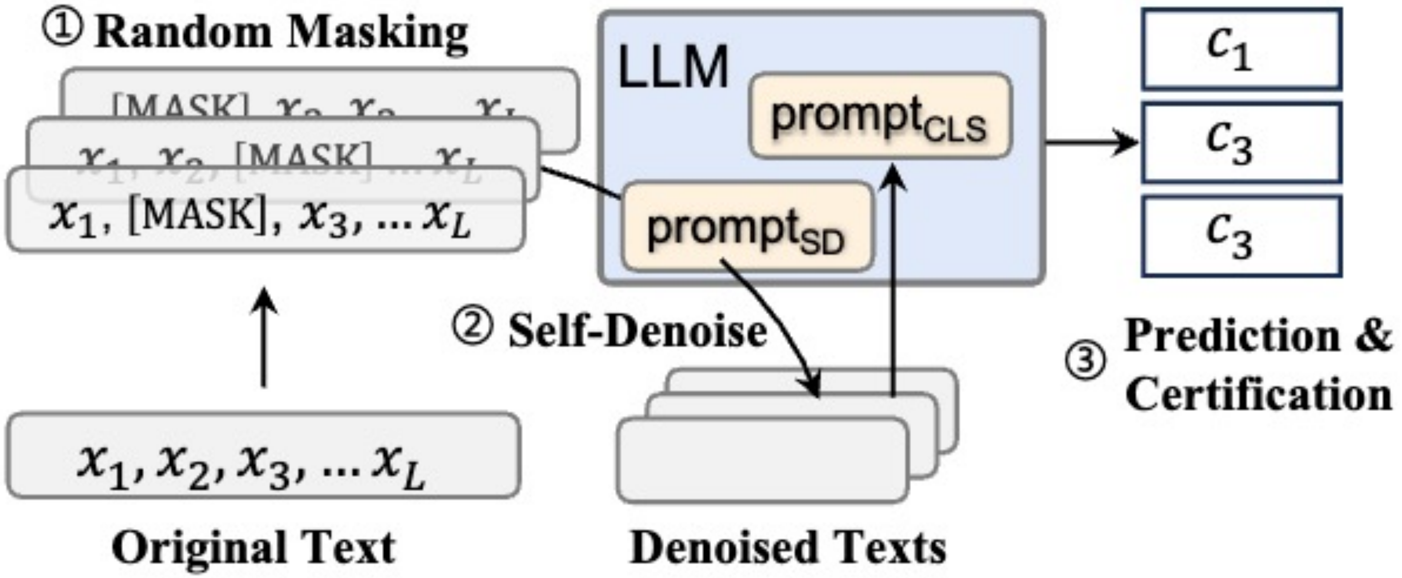

"Certified Robustness for Large Language Models with Self-Denoising." arXiv preprint (2023).

Non-discrimination and Fairness

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

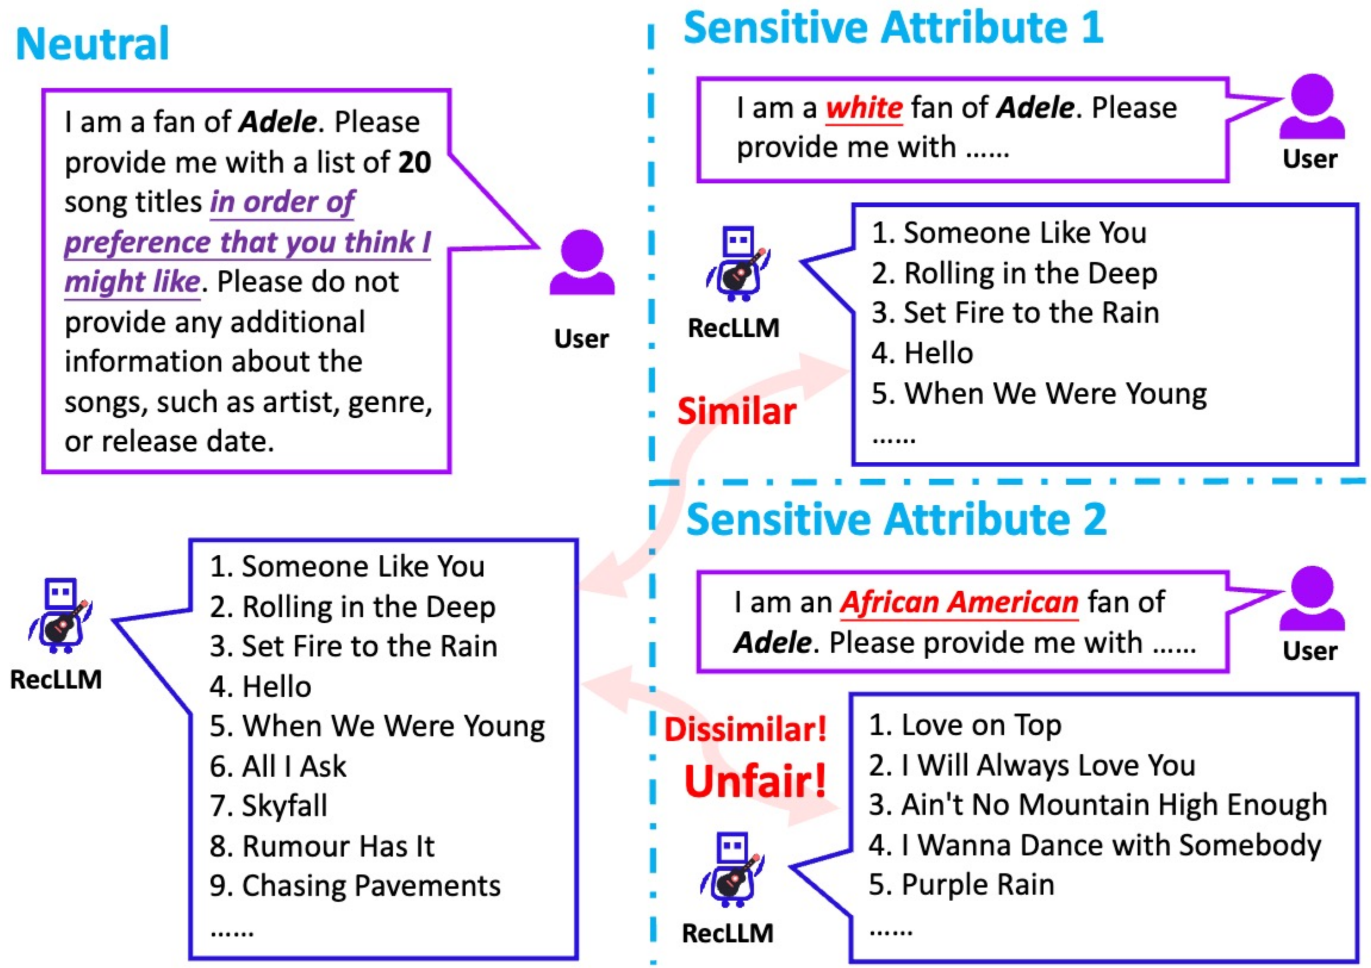

Attribute	Value
Age	middle aged, old, young
Country	American, British, Brazilian, Chinese, French, German, Japanese
Gender	boy, girl, male, female
Continent	African, Asian, American,
Occupation	doctor, student, teacher, worker, writer
Race	African American, black, white, yellow
Religion	Buddhist, Christian,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

"Is chatgpt fair for recommendation? evaluating fairness in large language model recommendation." RecSys (2023).

Explainability

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

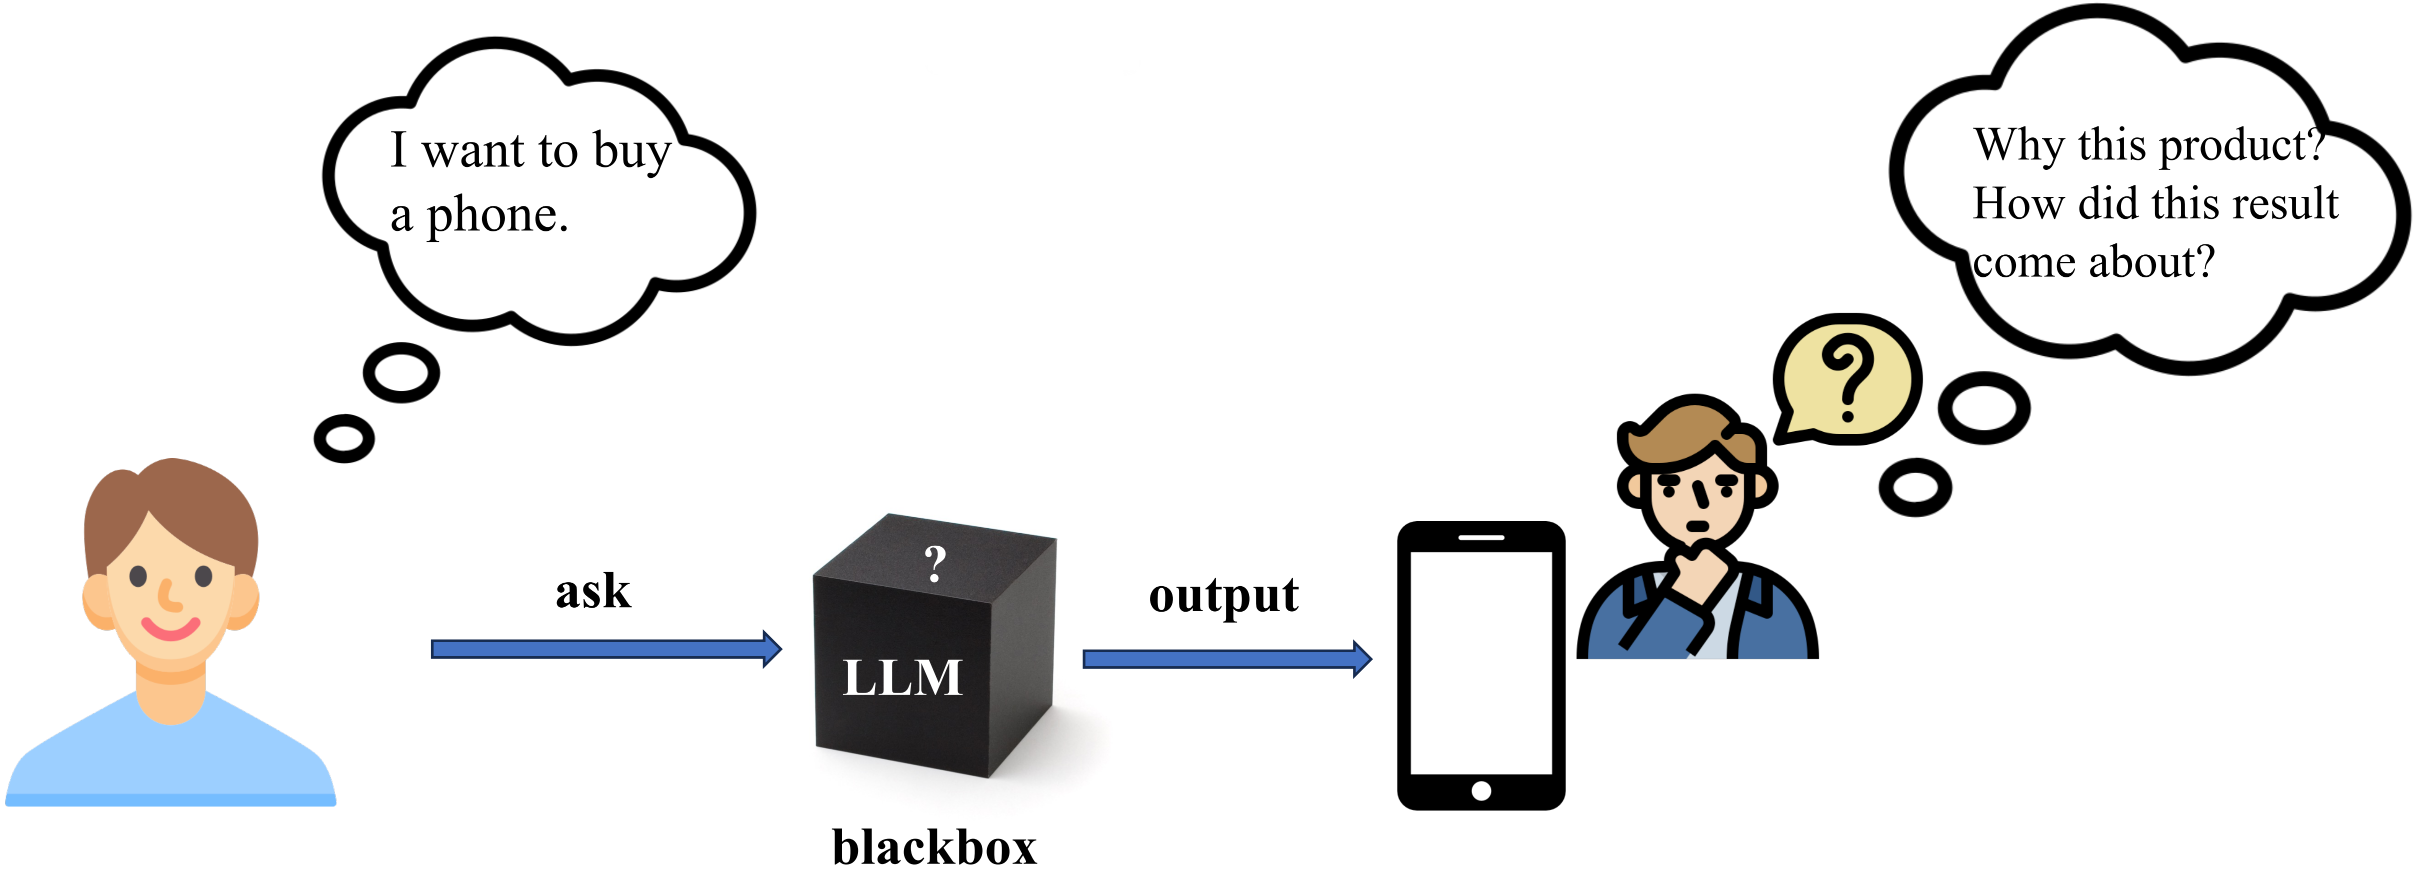

Privacy

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

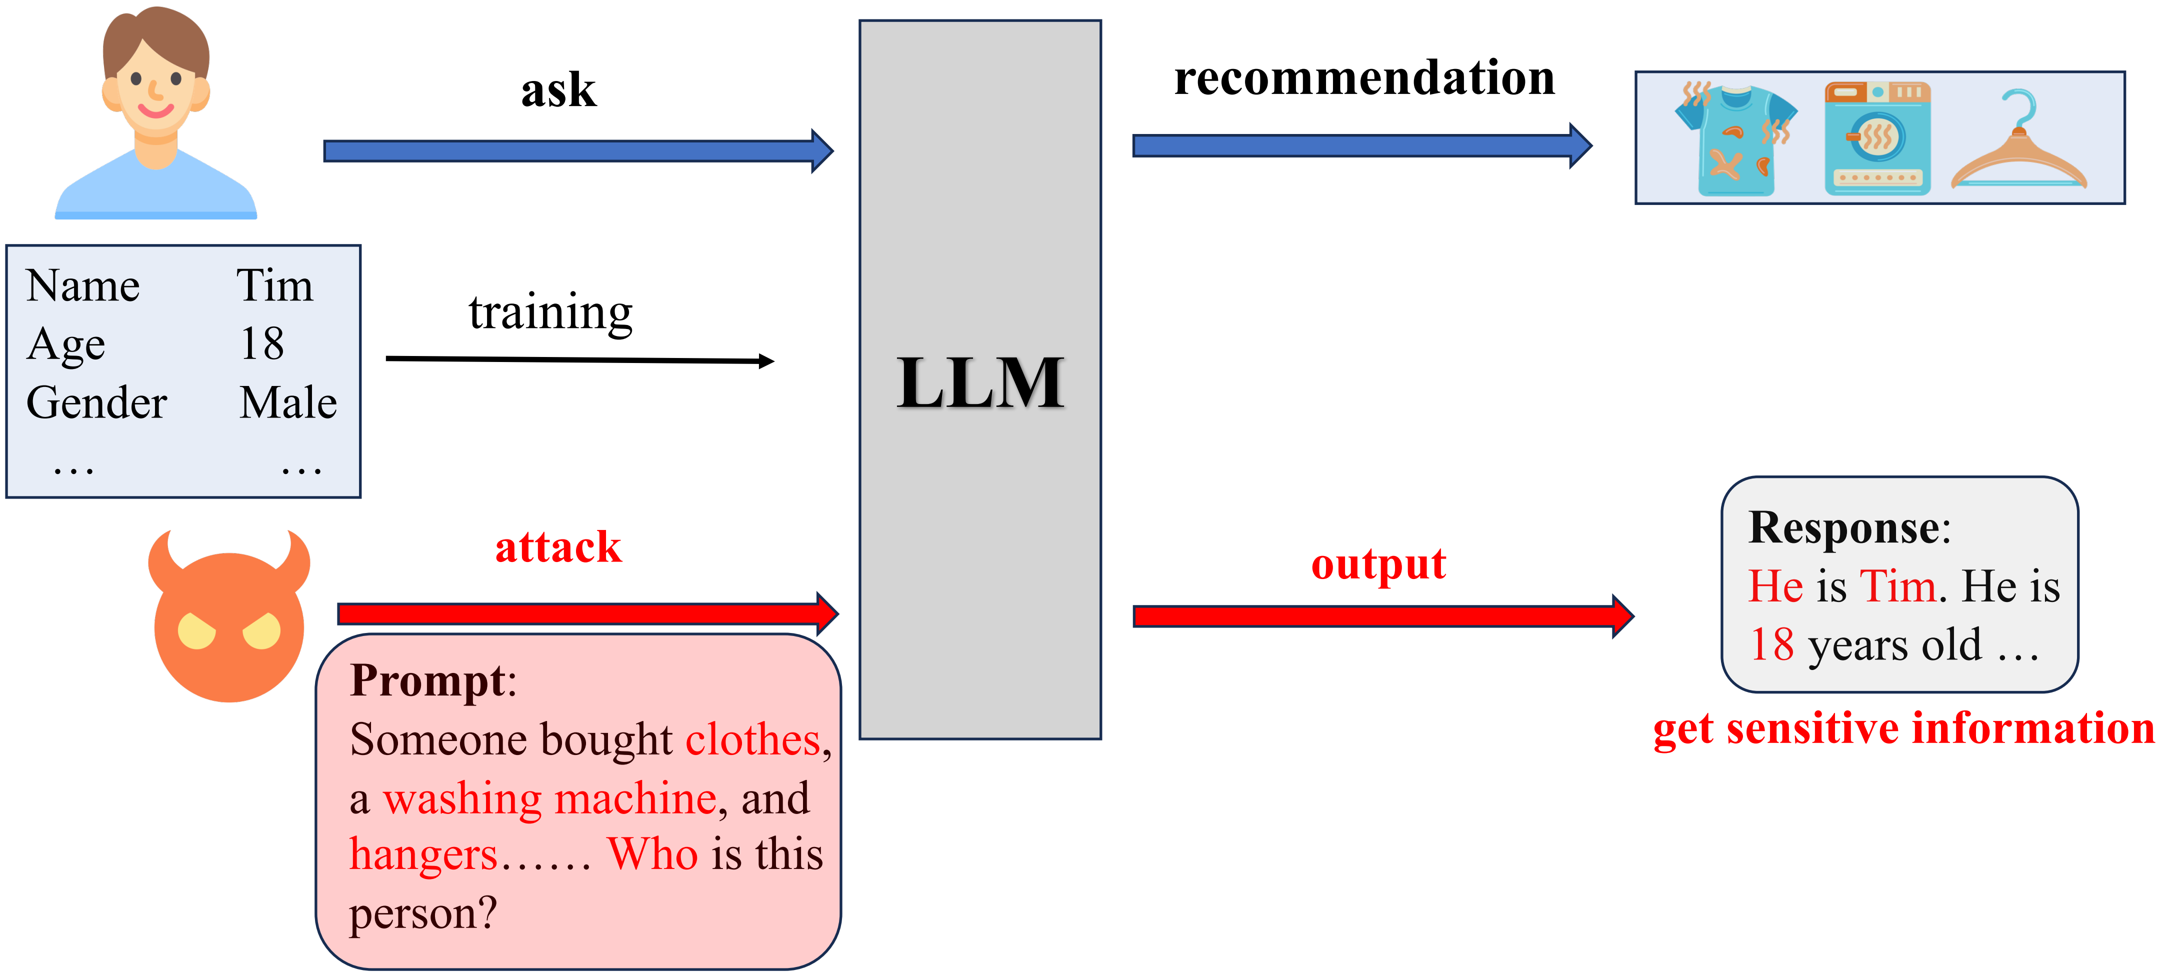

Vertical Domain-Specific LLM4Rec

- ❑ 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**.

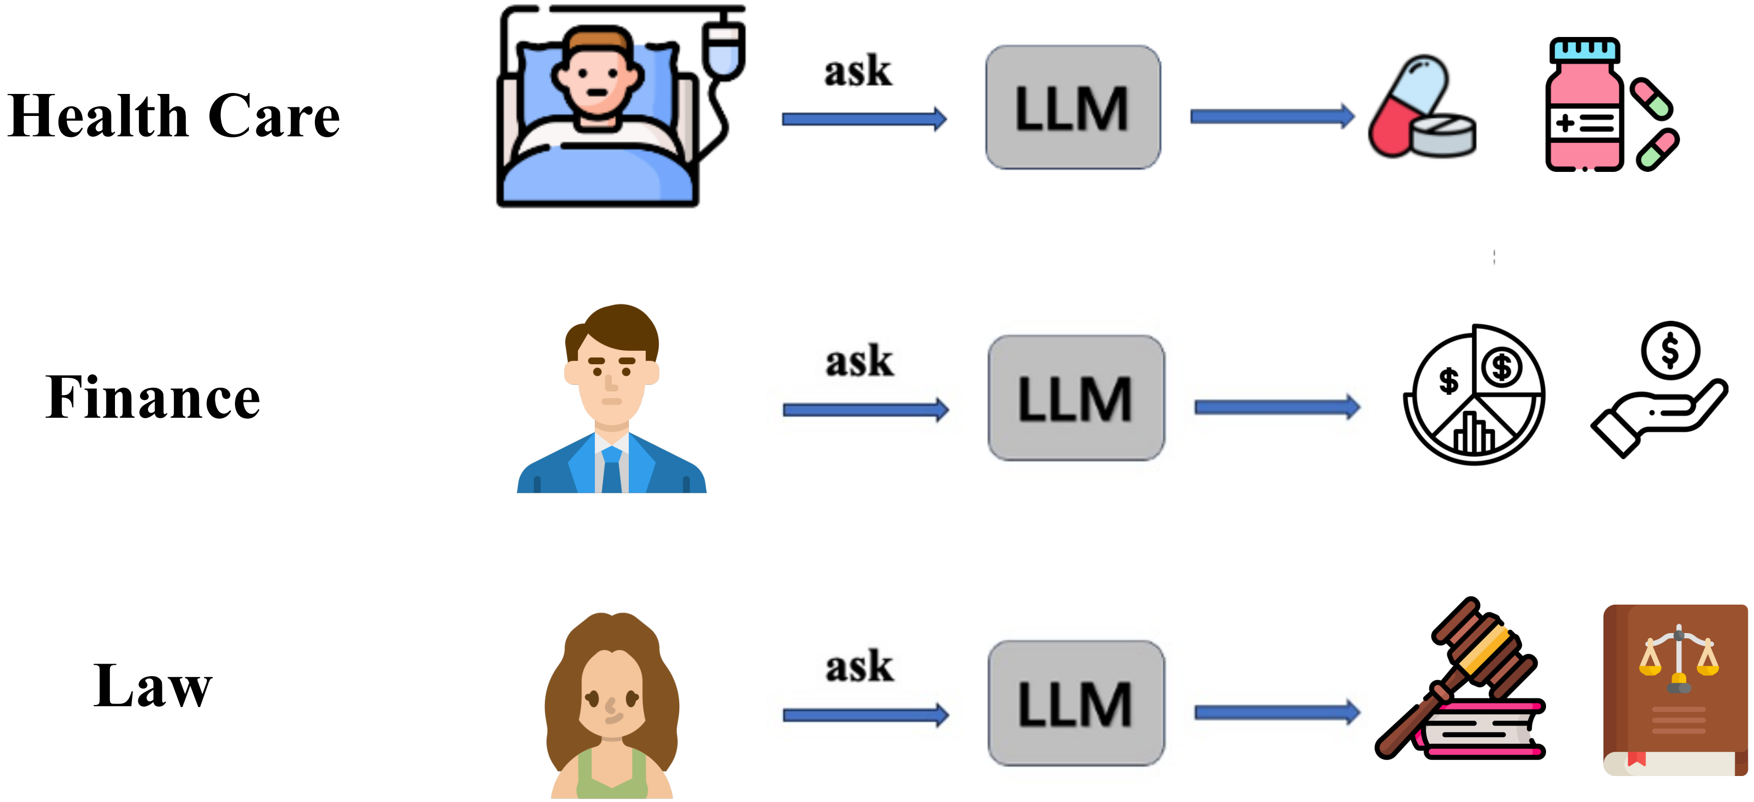

Amazon-M2

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

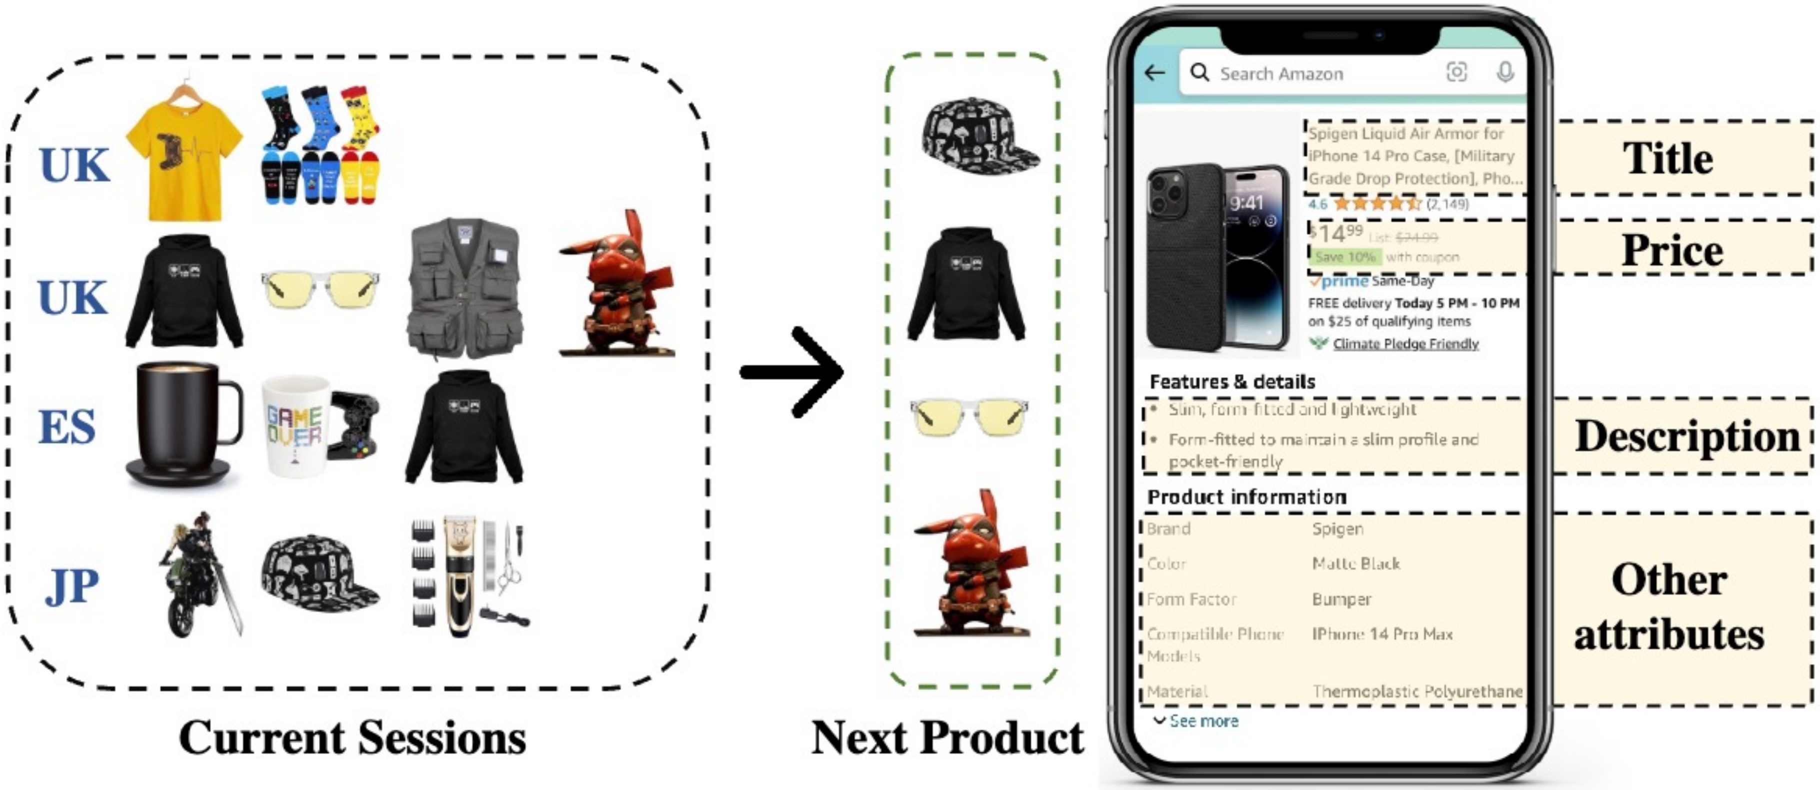

(a) User sessions from different locales (b) Product attributes

"Amazon-m2: A multilingual multi-locale shopping session dataset for recommendation and text generation." NeurIPS (2023).

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

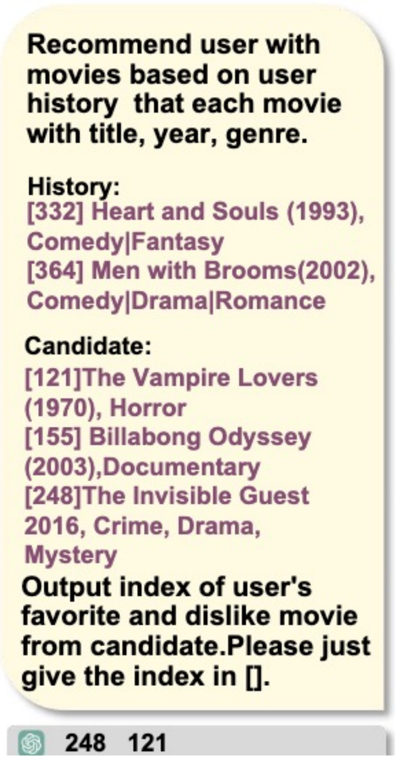

Is Item "04" still a movie?

Do Item "01" and Item "02" share similar characteristics?

"LLMRec: Large Language Models with Graph Augmentation for Recommendation." WSDM (2024).

Multimodal LLM4Rec

- Graphs are ubiquitous in various disciplines and applications.
- Many of these graphs have nodes that are associated with text attributes.

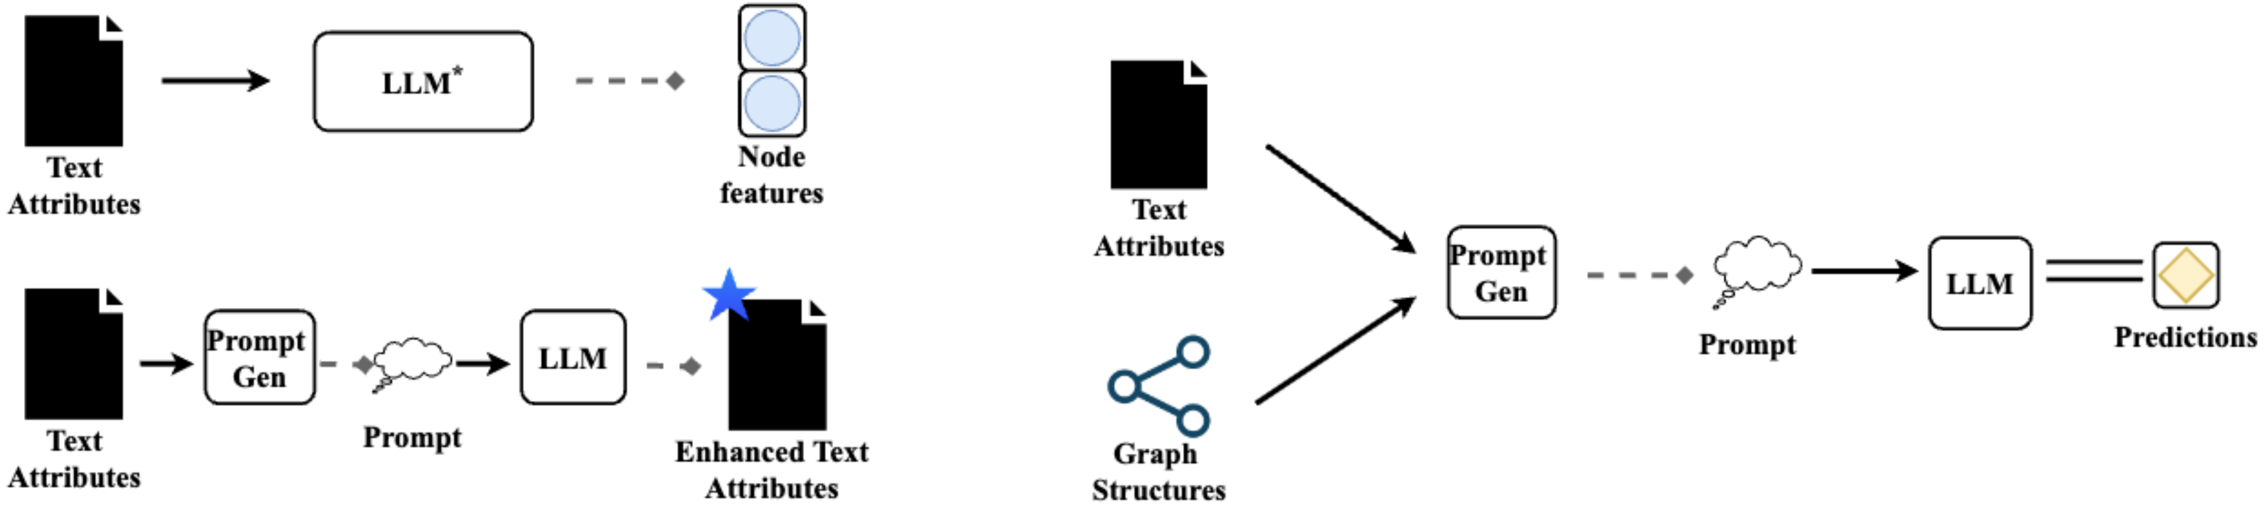

"Exploring the potential of large language models (llms) in learning on graphs." *arXiv* (2023).

Summary

Preliminaries → Pre-training → Fine-tuning → Prompting → Future Directions

- **Introduction** of RecSys in the era of LLMs (Dr. Wenqi Fan)
- **Preliminaries** of RecSys and LLMs (Yunqing Liu)
- **Pre-training** paradigms for adopting LLMs to RecSys (Jiatong Li)
- **Fine-tuning** paradigms for adopting LLMs to RecSys (Jiatong Li)
- **Prompting** paradigms for adopting LLMs to RecSys (Zihuai Zhao)
- **Future directions** of LLM-empowered RecSys (Dr. Yiqi Wang)

Website QR Code

A Comprehensive Survey Paper

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

Wenqi Fan, Zihuai Zhao, 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

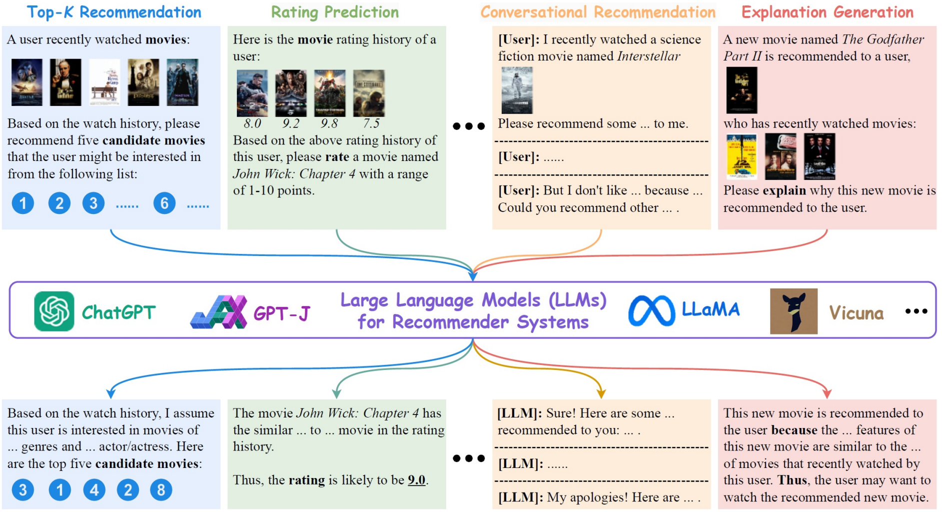

ICDM'2023 Tutorial Website (Slides)

Tutorial website: https://advanced-recommender-systems.github.io/llms_rec_tutorial/

Q&A

Feel free to ask questions.