Towards Unified Conversational Recommender Systems via Knowledge-Enhanced Prompt Learning

Xiaolei Wang*, Kun Zhou*, Ji-Rong Wen, Wayne Xin Zhao
Renmin University of China

• Background

Conversational recommender system (CRS)

A system aims to provide high-quality recommendations to users via natural language conversations. Typically, it needs to solve two tasks:

- **Recommendation**: predict user-preferred items;
- **Conversation**: generate a proper response for conversing with the user.

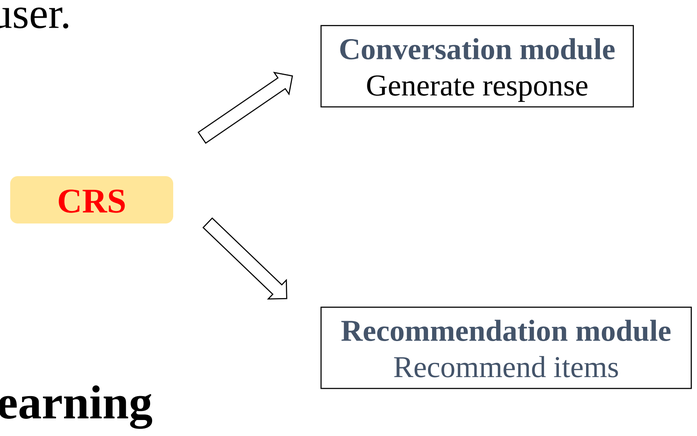

Prompt Learning

Most of PLMs are pre-trained with the objective of language modeling but are fine-tuned on downstream tasks with quite different objectives. To overcome the gap between pre-training and fine-tuning, prompt learning has been proposed, which relies on carefully designed prompts to reformulate the downstream tasks as the pre-training task.

- **Discrete Prompt**: actual text string
- **Continuous Prompt**: described directly in the embedding space of the underlaying PLM

• Motivation

Existing works

A capable CRS should be able to **seamlessly integrate the recommendation module and the conversation module**, because they are highly coupled.

- One line of work incorporates shared knowledge resources and their representations (e.g., knowledge graphs and reviews)

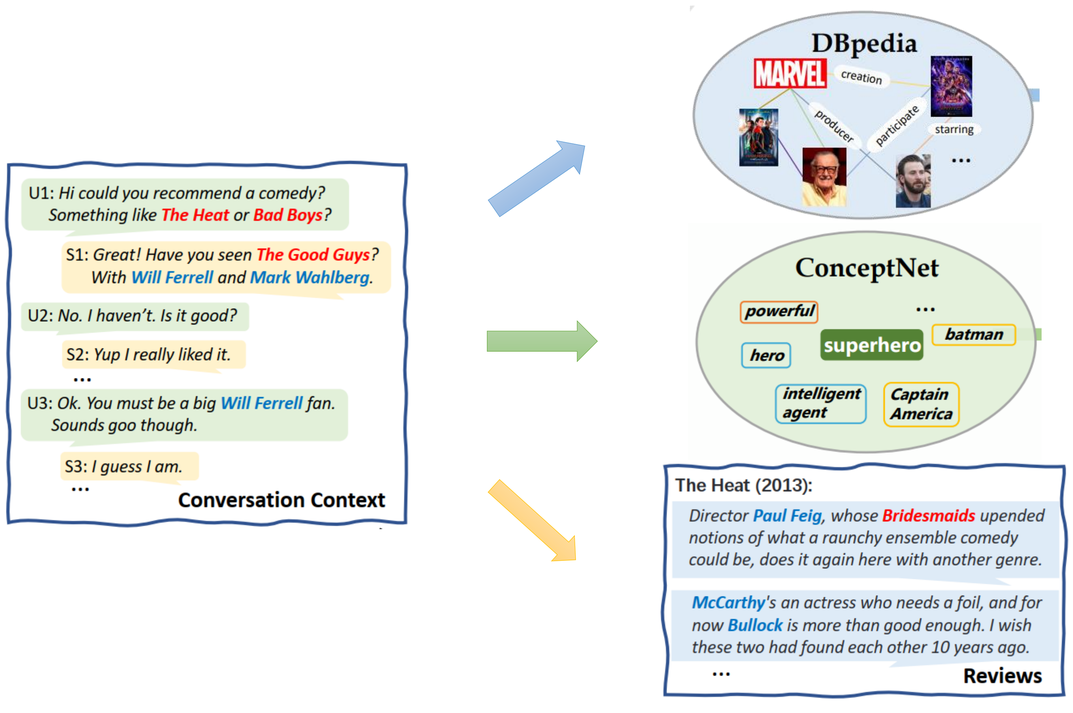

- Another line of work designs special representation alignment strategies, such as pre-training tasks and regularization terms (e.g., mutual information maximization and contrastive learning)

The fundamental issue of **semantic inconsistency** between the recommendation and conversation modules has not been well addressed.

As shown in the following figure, although the recommendation module predicts the movie "Frozen 2 (2019)", the conversation module generates a mismatched response that contains another movie "Pretty Woman (1990)".

USER:	Hello! I am looking for some movies.
HUMAN:	What kinds of movie do you like? I like animated movies such as *Frozen (2013)*.
USER:	I do not like animated films. I would love to see a movie like *Pretty Woman (1990)* starring Julia Roberts. Know any that are similar?
KGSF:	**Recommendation:** Frozen 2 (2019) **Response:** *Pretty Woman (1990)* is a great movie.
OURS:	**Recommendation:** My Best Friend's Wedding (1997) **Response:** Have you seen *My Best Friend's Wedding (1997)*? Julia Roberts also stars in it.
HUMAN:	*Pretty Woman (1990)* was a good one. If you are in it for Julia Roberts you can try *Runaway Bride (1999)*.

Two Reasons:

- Most of these methods develop the two modules with different architectures or techniques;
- Results from one module can-not be perceived and utilized by the other.

• Approach

Overview

Base PLM: DialoGPT (pre-trained with dialogue corpus)

Prompt-augmented Dialogue Context: $\widetilde{C} \rightarrow \underbrace{p_1, \ldots, p_{n_P}}_{\text{prompt tokens}}, \underbrace{w_1 \cdots w_{n_W}}_{\text{word tokens}}$

Semantic Fusion

Due to pre-trained on general dialogue corpus, DialoGPT lacks domain knowledge needed for CRS.

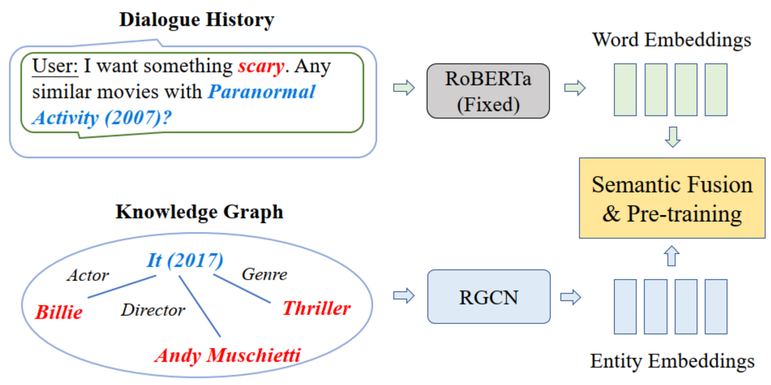

- Encoding: RoBERTa for word tokens, R-GCN for KG entities
- Semantic Fusion: Bi-attention
- Pre-Training: Self-supervised learning

Subtask Prompt Design

- Response Generation $\widetilde{C}_{gen} \rightarrow [\ \widetilde{\mathbf{T}};\ \mathbf{P}_{gen};\ C\ ]$
  - To share intermediate results, DialoGPT generates the response template instead of the complete response.

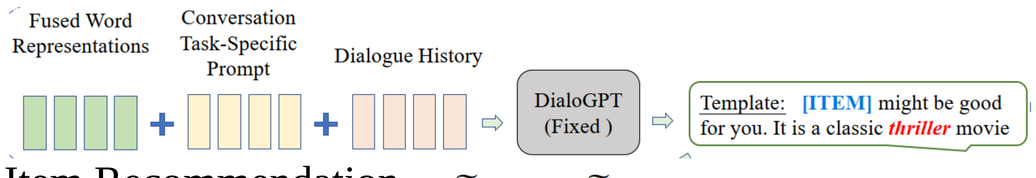

- Item Recommendation $\widetilde{C}_{rec} \rightarrow [\ \widetilde{\mathbf{E}};\ \mathbf{P}_{rec};\ C;\ S\ ]$
  - Include the response templates from the response generation task as part of the prompts.

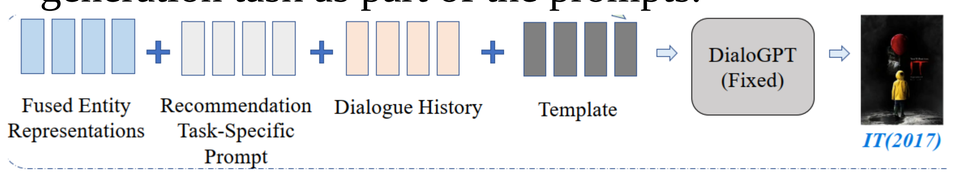

- Finally, the items will be filled into the template as complete responses.