Neighborhood-Based Collaborative Filtering for Conversational Recommendation

Zhouhang Xie*[1], Junda Wu*[1], Hyunsik Jeon*[1], Zhankui He[2],
Harald Steck[3], Rahul Jha[3], Dawen Liang[3], Nathan Kallus[34], Julian McAuley[1]

[1]UC San Diego [2]Google DeepMind [3]Netflix Research [4]Cornell University

https://github.com/zhouhanxie/neighborhood-based-CF-for-CRS

(1) Introduction

- ***Challenge in Conversational Recommender Systems (CRS)***
  - CRS need to handle complex, knowledge-intensive user queries
  - Existing methods rely on external knowledge bases or LLMs, which can be resource-intensive and have practical limitations
- ***Neighborhood-Based Collaborative Filtering (NBCRS)***
  - We leverage existing dialogue data to recommend items based on similar past queries without external knowledge bases or large models

(2) Proposed Method: NBCRS

- ***Key Ideas***
  - Retrieve similar dialogue contexts from training data
  - Recommend items frequently associated with these contexts
  - (Optional) Train a scoring model $p(item|query)$ to incorporate the scores further

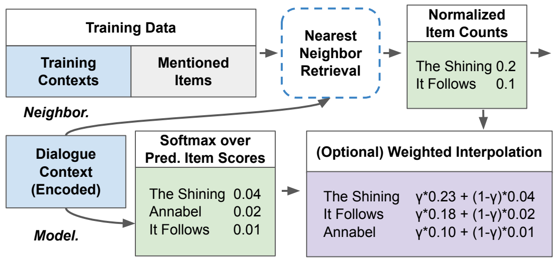

(3) Experiments

- ***Datasets***

Dataset	Total Movies	N. Train Samples	N. Test Samples
Inspired	1506	731	211
Redial	6476	8929	4288
Reddit	29705	39928	19438

Table 1: Statistics of the Datasets

- ***Overall Performance (see the paper for full-table)***

Model	Setting	Inspired Recall@20	Reddit Recall@20	Redial Recall@20
KGSF	Sft+KG	9.17 0.39	8.90 0.49	17.05 0.01
UniCRS	Sft+KG	18.59 0.12	9.79 0.50	27.12 0.28
Popularity	-	11.3 2.19	2.21 0.10	6.01 0.36
FISM	Sft	13.45 0.49	6.51 0.48	8.24 0.46
Gemma-2B	Zero-Shot	4.74 1.47	2.89 0.12	5.78 0.36
	Sft	2.37 1.05	3.49 0.13	5.11 0.33
Vicuna-7B	Zero-Shot	11.37 2.20	6.06 0.17	13.67 0.52
	Sft	10.43 2.11	7.18 0.18	13.67 0.52
NBCRS	Zero-Shot	14.69 2.44	14.08 0.24	16.34 0.56
	NB	**16.11** 2.23	15.52 0.25	16.58 0.54
	MB	1.42 0.81	13.62 0.24	14.80 0.54
	N+M	15.16 0.21	<u>15.58</u> 0.26	**16.86** 0.57

Table 3: Performance of models across datasets with standard errors. The reported numbers are percentages. Best performance excluding and including knowledge-graph-enhanced models are bolded and <u>underlined</u>, respectively.

- ***Further Analysis***

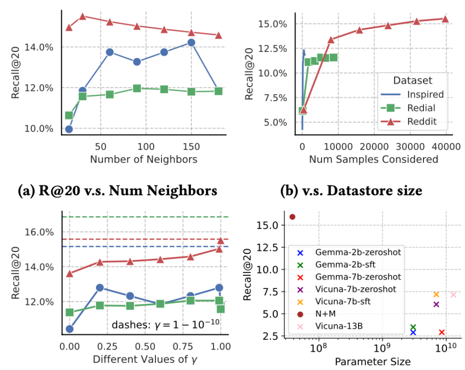

(a) R@20 v.s. Num Neighbors (b) v.s. Datastore size

(c) v.s. Neighborhood Weight (d) v.s. Model Size (on Reddit)

Figure 2: Analysis and ablations

(4) Conclusion

- ***Contributions***
  - Introduced a simple yet effective method for CRS: NBCRS
  - Demonstrated that NBCRS can match or exceed the performance of larger, resource-intensive methods
- ***When to Use NBCRS?***