Learning Personalized End-to-End Goal-Oriented Dialog

Liangchen Luo*[1], Wenhao Huang[2], Qi Zeng[3], Zaiqing Nie[4], Xu Sun[1]

[1]Peking University, [2]Shanghai Discovering Investment, [3]Stony Brook University, [4]Alibaba AI Labs

*luolc@pku.edu.cn

INTRODUCTION

The personalization of dialog systems is a meaningful task, which has not received much attention in the past few years. Here is an example to see what is wrong with conventional content-based goal-oriented dialog systems.

The conversations happen in a restaurant reservation scenario. The first and the second dialogs are with a young male with non-vegetable dietary, and the third one is with an elderly female with vegetable dietary.

These problems in the above example reflect three common issues with current models:

1. The inability to **adjust language style** flexibly.
2. The lack of a **dynamic conversation policy** based on the interlocutor's profile.
3. The incapability of **handling ambiguities** in user requests.

Correspondingly, the goals of personalization in goal-oriented dialog systems are solving these issues.

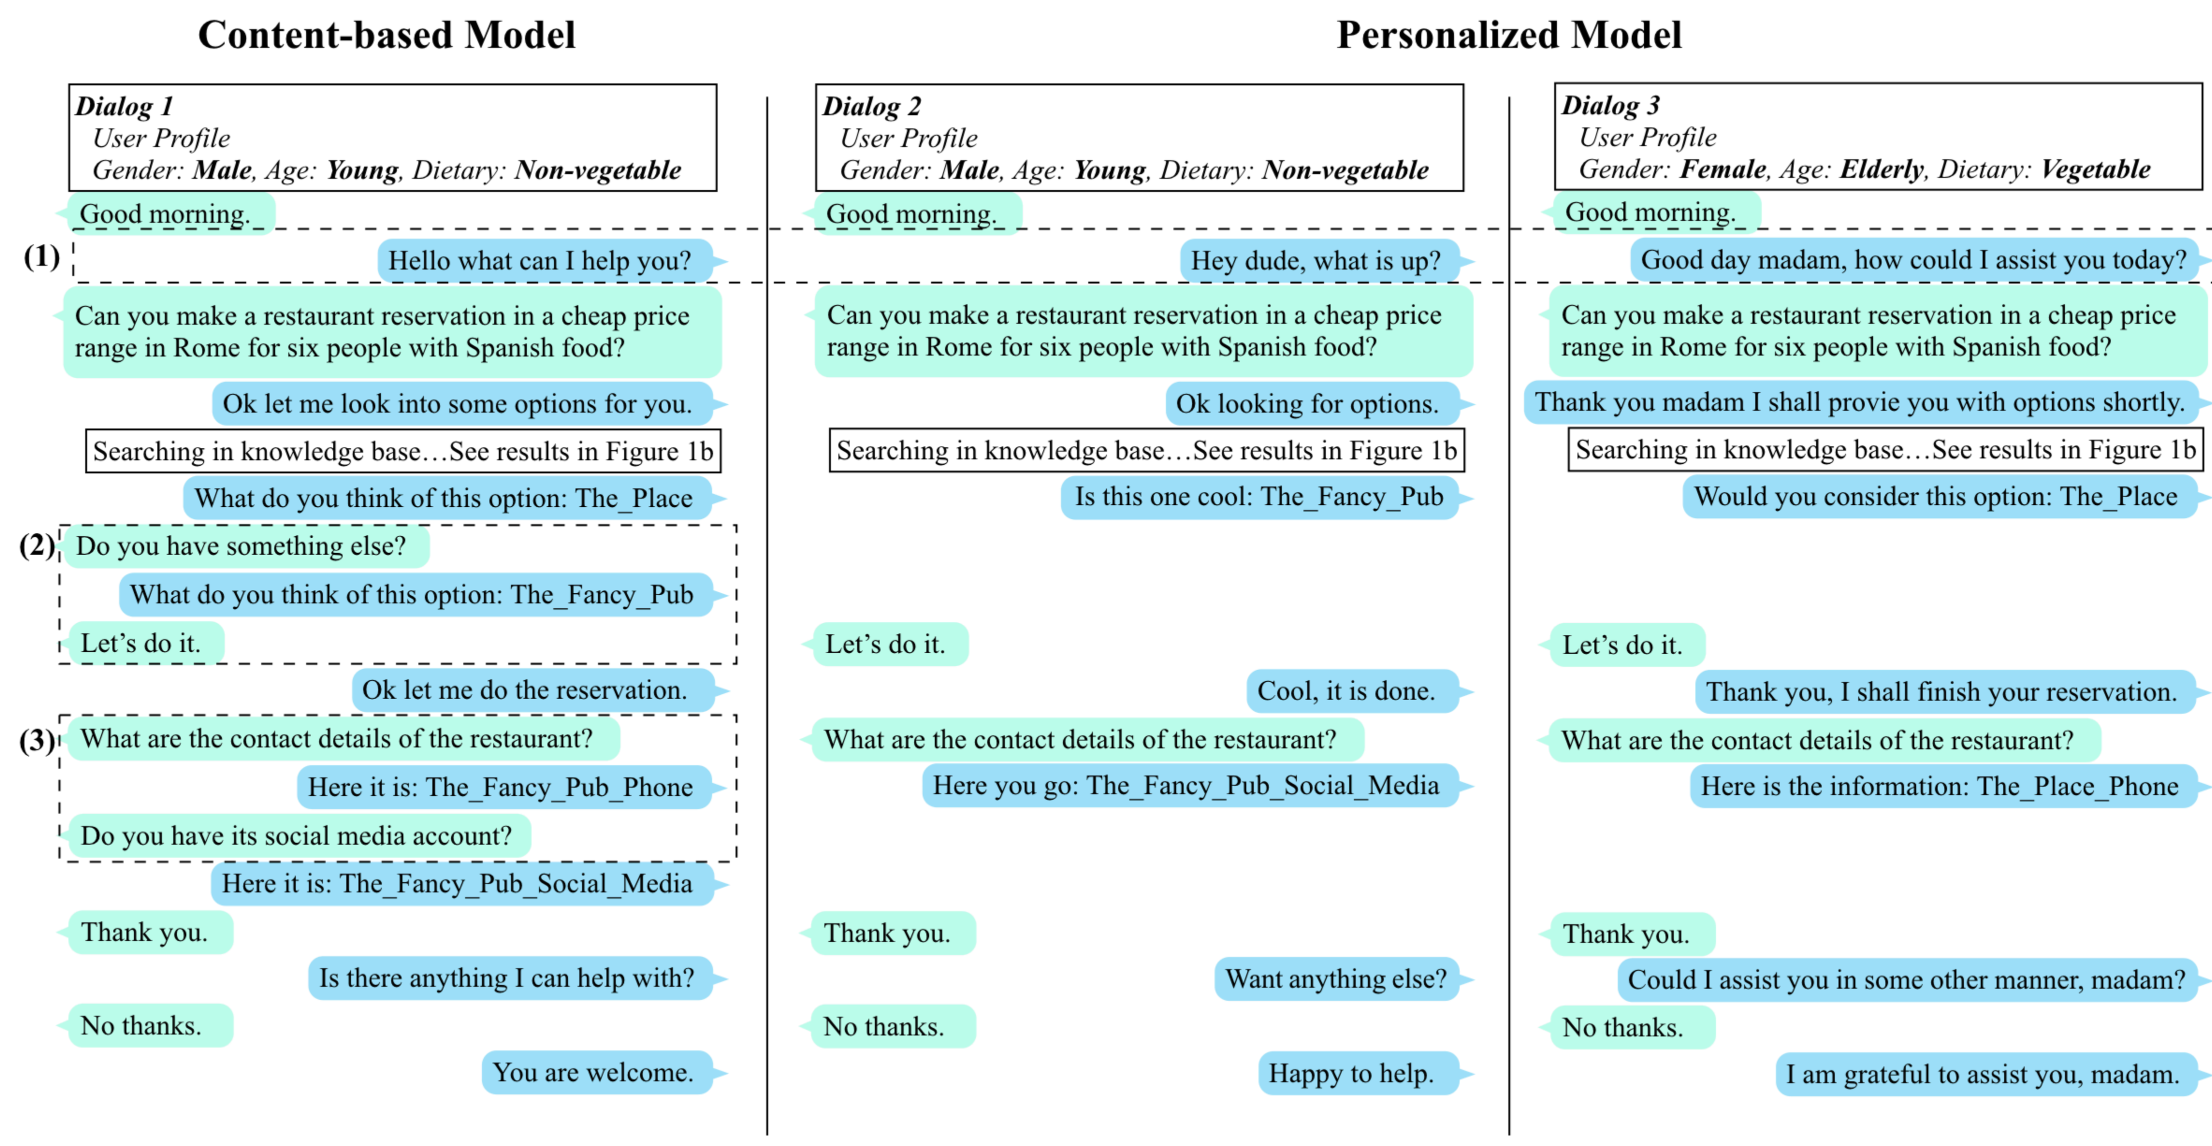

	Price	Location	Number	Cuisine	Phone	Social Media	Type
The_Place	Cheap	Rome	6	Spanish	The_Place_Phone	The_Place_Social_Media	Vegetable
The_Fancy_Pub	Cheap	Rome	6	Spanish	The_Fancy_Pub_Phone	The_Fancy_Pub_Social_Media	Non-vegetable
……							

MODEL & RESULTS

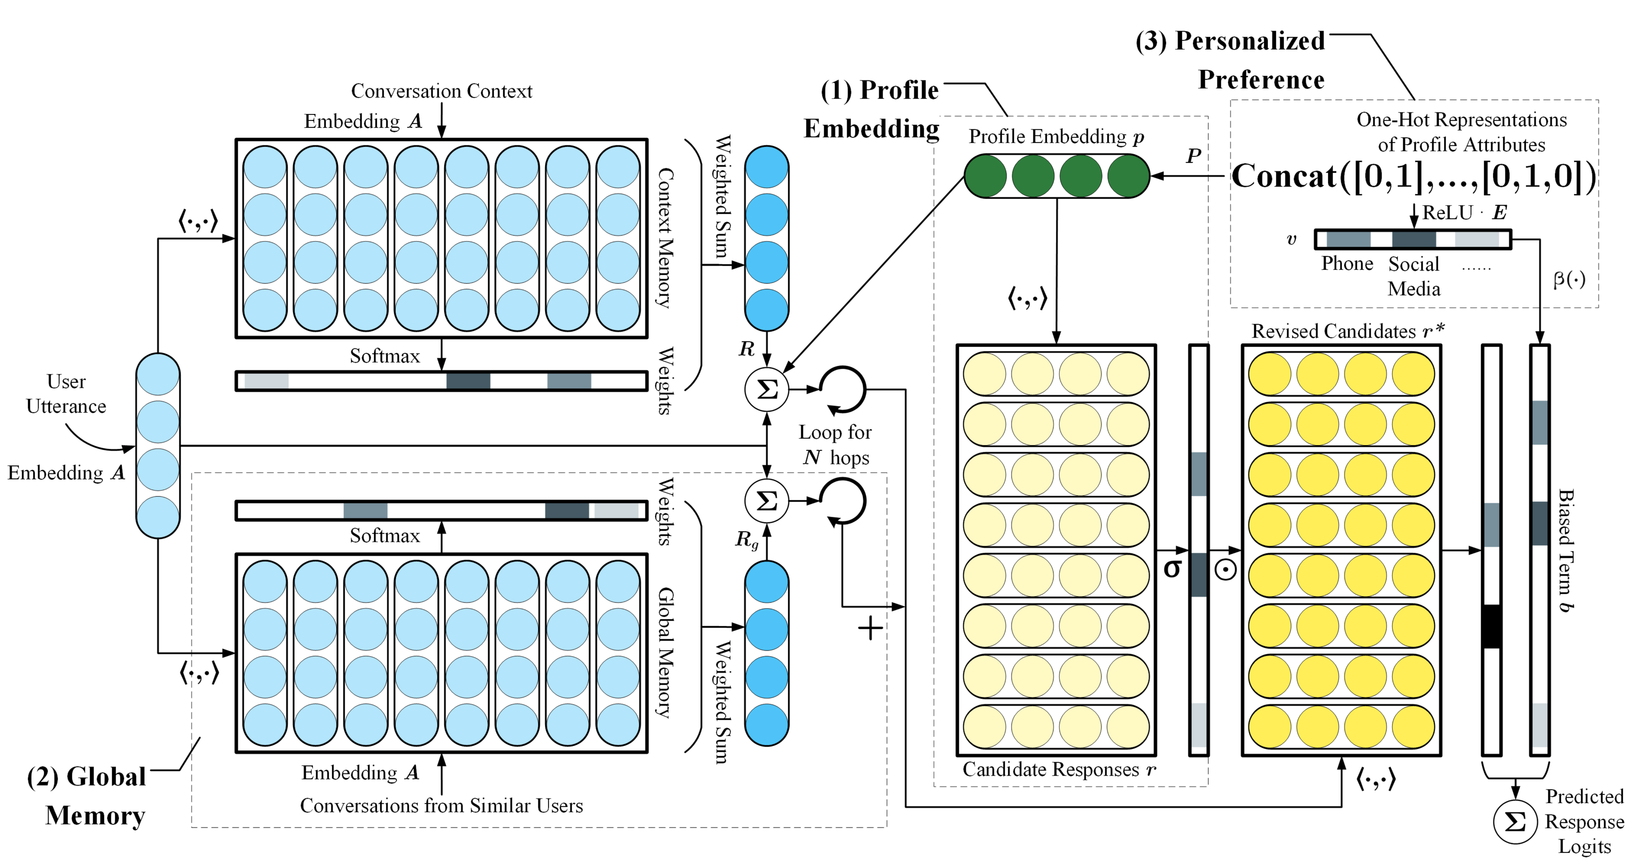

Our model, Personalized MemN2N, is in the vein of the memory network models for goal-oriented dialog, consisting of three main components: **profile embedding**, **global memory** and **personalized preference**.

- The incoming user utterance is embedded into a query vector. The model first reads the memory (at top-left) to find relevant history and produce attention weights. Then it generates an output vector by taking the weighted sum followed by a linear transformation.
- Part (1) is **Profile Embedding**: the profile vector p is added to the query at each iteration, and is also used to revise the candidate responses r.
- Part (2) is **Global Memory**: this component (at bottom-left) has an identical structure as the original MemN2N, but it contains history utterances from other similar users.
- Part (3) is **Personalized Preference**: the bias term is obtained based on the user preference and added to the prediction logits.

We conduct experiments on the personalized bAbI dialog dataset (Joshi et al., 2017) and consider the following baselines: Supervised Embedding Model, Memory Network (Bordes et al., 2017), and Split Memory Network (Joshi et al., 2017).

Models	Issuing API Calls	Updating API Calls	Displaying Options	Providing Information	Full Dialog
Supervised Embeddings	84.37	12.07	9.21	4.76	51.60
MemN2N	99.83	**99.99**	58.94	57.17	85.10
Split MemN2N	85.66	93.42	68.60	57.17	87.28
Profile Embedding	**99.96**	99.96	71.00	57.18	93.83
Global Memory	99.76	99.93	71.01	57.18	91.70
Profile Model	99.93	99.94	71.12	57.18	93.91
Preference Model	99.80	99.97	68.90	81.38	94.97
Personalized MemN2N	99.91	99.95	**71.43**	**81.56**	**95.33**

ANALYSIS

Analysis of Profile Embeddings

We group the candidates by their corresponding user profile. For each profile, we generate tendency weights and collect the average value for each group. The results are visulized by a confusion matrix. The weights on the diagonal are significantly larger than others, which demonstrates the contribution of profile embeddings in candidate selection.

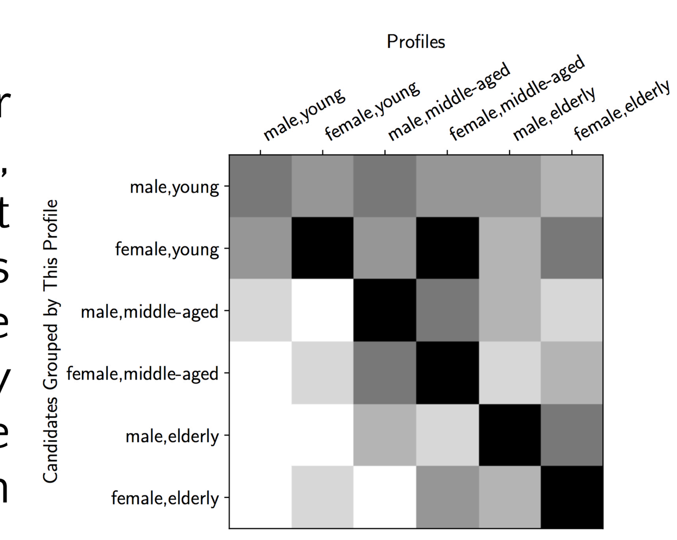

Analysis of Preference

As we use a preference vector to represent the user's preference over the columns in the knowledge base, we can investigate the learned arguments grouped by profile attributes. The model successfully learns the fact that young people prefer social media as their contact information, while middle-aged and elderly people prefer phone number.

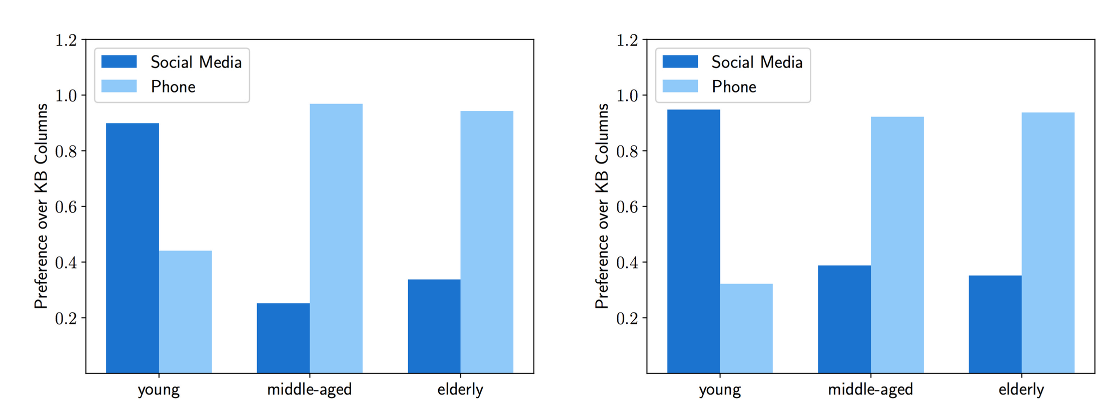

(a) Task 4 (b) Task 5

Human Evaluation

Personalized MemN2N wins the MemN2N baseline with 27.6% and 14.3% higher in terms of task completion rate and satisfaction, with p < 0.03.

MISC

References

1. Bordes A., Boureau Y.-L., and Weston J. 2017. Learning end-to-end goal-oriented dialog. In *Proc. of ICLR 2017*.
2. Joshi C. K., Mi F., and Faltings B. 2017. Personalization in goal-oriented dialog. In *NeurIPS 2017 workshop*.

Links

- arXiv: arxiv.org/abs/1811.04604
- website: www.luolc.com/publications/personalized-goal-oriented-dialog/

※*Scan QR code to view the paper* ↗