UvA-DARE (Digital Academic Repository)

DukeNet: A Dual Knowledge Interaction Network for Knowledge-Grounded Conversation

Meng, C.; Ren, P.; Chen, Z.; Sun, W.; Ren, Z.; Tu, Z.; de Rijke, M.

DOI

[10.1145/3397271.3401097](https://doi.org/10.1145/3397271.3401097)

Publication date

2020

Document Version

Author accepted manuscript

Published in

SIGIR '20

[Link to publication](#)

Citation for published version (APA):

Meng, C., Ren, P., Chen, Z., Sun, W., Ren, Z., Tu, Z., & de Rijke, M. (2020). DukeNet: A Dual Knowledge Interaction Network for Knowledge-Grounded Conversation. In *SIGIR '20: proceedings of the 43rd International ACM SIGIR Conference on Research and Development in Information Retrieval : July 25-30, 2020, virtual event, China* (pp. 1151–1160). Association for Computing Machinery. <https://doi.org/10.1145/3397271.3401097>

General rights

It is not permitted to download or to forward/distribute the text or part of it without the consent of the author(s) and/or copyright holder(s), other than for strictly personal, individual use, unless the work is under an open content license (like Creative Commons).

Disclaimer/Complaints regulations

If you believe that digital publication of certain material infringes any of your rights or (privacy) interests, please let the Library know, stating your reasons. In case of a legitimate complaint, the Library will make the material inaccessible and/or remove it from the website. Please Ask the Library: <https://uba.uva.nl/en/contact>, or a letter to: Library of the University of Amsterdam, Secretariat, Singel 425, 1012 WP Amsterdam, The Netherlands. You will be contacted as soon as possible.

UvA-DARE is a service provided by the library of the University of Amsterdam (<https://dare.uva.nl>)

DukeNet: A Dual Knowledge Interaction Network for Knowledge-Grounded Conversation

Chuan Meng¹ Pengjie Ren^{2*} Zhumin Chen^{1*} Weiwei Sun¹

Zhaochun Ren¹ Zhaopeng Tu⁴ Maarten de Rijke^{2,3}

¹Shandong University, Qingdao, China

²University of Amsterdam, Amsterdam, The Netherlands

³Ahold Delhaize, Zaandam, The Netherlands

⁴Tencent AI Lab, Shenzhen, China

{mengchuan,sunweiweisdu}@mail.sdu.edu.cn,{p.ren,m.derijke}@uva.nl

{chenzhumin,zhaochun.ren}@sdu.edu.cn,tuzhaopeng@gmail.com

ABSTRACT

Today’s conversational agents often generate responses that not sufficiently informative. One way of making them more informative is through the use of external knowledge sources with so-called Knowledge-Grounded Conversations (KGCs). In this paper, we target the *Knowledge Selection* (KS) task, a key ingredient in KGC, that is aimed at selecting the appropriate knowledge to be used in the next response. Existing approaches to KS based on learned representations of the conversation context, that is previous conversation turns, and use Maximum Likelihood Estimation (MLE) to optimize KS. Such approaches have two main limitations. First, they do not explicitly track what knowledge has been used in the conversation nor how topics have shifted during the conversation. Second, MLE often relies on a limited set of example conversations for training, from which it is hard to infer that facts retrieved from the knowledge source can be re-used in multiple conversation contexts, and vice versa.

We propose Dual Knowledge Interaction Network (DukeNet), a framework to address these challenges. DukeNet explicitly models *knowledge tracking* and *knowledge shifting* as dual tasks. We also design Dual Knowledge Interaction Learning (DukeL), an unsupervised learning scheme to train DukeNet by facilitating interactions between knowledge tracking and knowledge shifting, which, in turn, enables DukeNet to explore extra knowledge besides the knowledge encountered in the training set. This dual process also allows us to define rewards that help us to optimize both knowledge tracking and knowledge shifting. Experimental results on two public KGC benchmarks show that DukeNet significantly outperforms state-of-the-art methods in terms of both automatic and human evaluations, indicating that DukeNet enhanced by DukeL can select more appropriate knowledge and hence generate more informative and engaging responses.

*Co-corresponding author.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

SIGIR ’20, July 25–30, 2020, Virtual Event, China

© 2020 Association for Computing Machinery.

ACM ISBN 978-1-4503-8016-4/20/07...\$15.00

<https://doi.org/10.1145/3397271.3401097>

CCS CONCEPTS

• **Computing methodologies** → **Natural language generation; Discourse, dialogue and pragmatics.**

KEYWORDS

Knowledge-grounded conversation; Knowledge selection; Dual learning; Chatbot

ACM Reference Format:

Chuan Meng, Pengjie Ren, Zhumin Chen, Weiwei Sun, Zhaochun Ren, Zhaopeng Tu, and Maarten de Rijke. 2020. DukeNet: A Dual Knowledge Interaction Network for Knowledge-Grounded Conversation. In *Proceedings of the 43rd International ACM SIGIR Conference on Research and Development in Information Retrieval (SIGIR ’20)*, July 25–30, 2020, Virtual Event, China. ACM, New York, NY, USA, 10 pages. <https://doi.org/10.1145/3397271.3401097>

1 INTRODUCTION

Open-domain conversational agents (a.k.a. chatbots) aim to satisfy human needs for information, communication, entertainment and more [1, 12, 13, 17]. The development of such agents has benefited significantly from advances in sequence-to-sequence learning [16, 34, 38]. However, a serious problem of vanilla sequence-to-sequence based conversational models is that they tend to generate dull and non-informative responses [3, 8, 43], such as “I don’t know” and “thank you.” To address this problem of generating uninformative responses, the Knowledge-Grounded Conversation (KGC) task has been introduced; it leverages external knowledge to enhance open-domain conversational models [7, 24, 32]. As shown in Fig. 1, given a conversation context and external knowledge pool (with a large set of knowledge sentences), the goal of KGC is to generate informative and engaging responses by referring to relevant knowledge. KGC can be divided into two sequential subtasks: (1) *Knowledge Selection* (KS): to select the appropriate knowledge at the current turn from a knowledge pool; and (2) *Response Generation* (RG): to generate a natural language response based on the selected knowledge and conversation context. Of the two subtasks, KS is of vital importance, as it decides what to be talked in the response, and selecting inappropriate knowledge will directly result in an inappropriate response [19].

There is a growing number of studies on KS and promising results have been achieved [see, e.g., 7, 9]. In terms of modeling, most approaches try to predict the next knowledge based on representations of the conversation context [19]. In terms of learning, most

Figure 1: An example of a knowledge-grounded conversation from the Wizard of Wikipedia dataset [7].

approaches optimize their models via Maximum Likelihood Estimation (MLE) based on conversations encountered in the training set [14]. Previous work fails to address two important characteristics of KS for conversations.

- (1) *knowledge tracking* (ground the knowledge that has been talked about to the conversation context) and *knowledge shifting* (select the knowledge to be talked about next) are not explicitly modeled. Explicitly modeling knowledge tracking and shifting allows us to better capture the interaction between the knowledge at adjacent turns. Consider, for example, in Fig. 1, for the knowledge K2 “YouTube was created by three former PayPal employees...” and the knowledge K3 “Google bought the site in November 2006 for US\$1.65 billion...” If we know that K2 has been used in the conversation context, it is natural to use K3 in the next response.
- (2) Unlike Question Answering (QA) tasks in which each query has only one unique answer in most cases [30, 31], Kim et al. [14] show that one-to-many mapping between context and knowledge is common in KGC, i.e., for a given conversation context, we can choose different knowledge sentences to form different responses. In fact, there is also many-to-one mapping between context and knowledge, i.e. the same knowledge sentence can be used in different conversation contexts. Thus, this forms a many-to-many mapping. For instance, given the context in Fig. 1, it is also appropriate to select other knowledge besides K3, e.g., the most popular video or the most famous video creator. Conversely, it is also appropriate to use K3 in Fig. 1 in other contexts such as “I wonder what the relationship between Google and Youtube is” or “Do you know the history of Youtube.” Unfortunately, for a given conversation context, only one knowledge sentence is encountered in existing datasets [7, 25], because it is really hard to collect all possible knowledge sentences for a certain context when constructing a KGC dataset. Existing studies into KGC rely on MLE to train their models, while MLE only considers the demonstrated knowledge as ground truth, which is restrictive.

To address the issues listed above, we propose a novel framework, Dual Knowledge Interaction Network (DukeNet), that explicitly models *knowledge tracking* and *knowledge shifting* in conversations. Besides, we further formulate knowledge tracking and knowledge shifting as dual tasks and devise Dual Knowledge Interaction Learning (DukeL) to better facilitate interaction between them, as shown in Fig. 1. DukeL enables knowledge tracking and shifting to teach

each other in an unsupervised learning way without external supervision. We alternate these two dual processes until convergence. During the two dual processes, DukeL will explore and reward extra appropriate knowledge that is not manifest in the training set, but which help address the many-to-many mapping phenomenon in conversations.

However, there are incompatible dual processes between training and inference. Specifically, during inference we can only execute knowledge tracking and shifting in order, thus knowledge tracking cannot get the shifted knowledge as input. To alleviate this problem, we further distinguish knowledge tracking as *prior and posterior knowledge tracking*, where the former only takes context as input, while the latter additionally takes the shifted knowledge as input. During training, besides optimizing the posterior knowledge tracking and knowledge shifting based on the closed loop between them, we force the prior knowledge tracking to approximate the output of posterior knowledge tracking to get benefit. During inference, we only execute prior knowledge tracking and shifting.

Experiments on the Wizard of Wikipedia [7] and Holl-E [25] datasets show that DukeNet can select more appropriate knowledge and hence generate more informative and engaging responses by explicitly modeling knowledge tracking and knowledge shifting, and formulating their interactions as dual learning.

The contributions of this paper are summarized as follows:

- We propose a novel framework, DukeNet, which explicitly models *knowledge tracking* and *knowledge shifting* as dual tasks to promote KS.
- We devise DukeL, which introduces an unsupervised learning scheme for KS.
- We conduct automatic and human evaluations on two benchmark datasets, which shows that DukeNet outperforms recent state-of-the-art methods, and can select more appropriate knowledge to generate more informative and engaging responses.

2 RELATED WORK

We survey two types of related work: Knowledge-Grounded Conversations (KGCs) and dual learning.

2.1 Knowledge-grounded conversation

Existing methods on KGC can be divided into two categories: *structured knowledge based* and *unstructured knowledge based*. The former conditions response generation on knowledge triples [21, 22, 41, 52], while the latter conditions on free knowledge text [28]. Unstructured knowledge based approaches to KGC can be further divided into *document based* (where they are given whole documents, e.g., Reddit articles) [18, 25, 29, 53] and *sentence based* (where they are given separate sentences, e.g., Foursquare tips) [7, 9, 14, 19]. In this paper, we focus on sentence based KGC. Next, we briefly introduce recent advances in this direction.

Ghazvininejad et al. [9] regard Foursquare tips as knowledge sentences and propose MemNet (a variant of Memory Network [36]) which stores the latent representations of knowledge sentences in a memory module such that a SeqSeq model can attentively select useful knowledge from it to generate responses. Dinan et al. [7] collect a large sentence based benchmark dataset, Wizard of Wikipedia, which retrieves Wikipedia articles and then flattens all

Figure 2: An overview of DukeNet. Section 3 contains a walkthrough of the model.

the articles into separate sentences and clearly labels the ground truth knowledge sentence used in each response. They also propose Transformer MemNet (TMemNet) that improves MemNet by replacing RNN with Transformers [42] and introduce a KS loss to supervise the KS process. Lian et al. [19] propose Posterior Knowledge Selection (PostKS), which uses response and context to predict a distribution over knowledge and regards this distribution as posterior knowledge distribution. They use that as pseudo-label to guide KS during training process. Kim et al. [14] propose Sequential Knowledge Transformer (SKT), which sequentially models the history of KS of previous turns via a sequential latent variable model [35] to promote KS at the current turn. Although SKT models the knowledge used in conversation history, it does not explicitly distinguish knowledge tracking and shifting and still uses MLE to train the model based on the limited demonstrated examples in training set.

Unlike SKT, we explicitly distinguish knowledge tracking and shifting and regard them as dual tasks, leveraging the unsupervised dual interaction between them to further improve KS.

2.2 Dual learning

Dual learning has been successfully applied to many tasks, such as machine translation [11, 45], question answering [37, 40], conversation [4, 50], text style transfer [23] as well as image-to-image translation [49]. The core idea of dual learning is to take advantage of the closed loop between the forward agent f (mapping from domain X to domain Y) and backward agent g (mapping from domain Y to X) to improve the performances of each other.

There are many paradigms in dual learning. He et al. [11] first propose dual learning and apply it to machine translation, which executes dual learning by maximizing the likelihood of data reconstruction. Concretely, $x \in X$ is first fed to the forward agent f to output $\hat{y} \in Y$, and $\hat{y}$ is then fed to the backward agent g to output $\hat{x} \in X$. The distortion between x and $\hat{x}$ is used as a reconstruction reward to optimize the forward agent f . Similarly, the reconstruction reward between y and $\hat{y}$ can be used to optimize the backward agent g . Xia et al. [47] propose dual supervised learning, which introduces a duality regularization term in their loss function. The special term reflects the probabilistic correlation between two agents to better guide the training. Xia et al. [46] argue that existing

work only considers the duality during the training process and further propose a dual inference framework, which enables dual agents to improve each other during the inference process without re-training. Xia et al. [48] propose model-level dual learning, which shares the model parameters playing similar roles in the two agents. Wang et al. [44] propose multi-agent dual learning, which introduces more agents in the two directions respectively to maximize the likelihood of data reconstruction. They show that more agents can lead to more reliable and robust reconstruction reward.

Unlike the work listed above, we regard knowledge tracking and shifting in KGC as two dual tasks and propose DukeL to let the two tasks teach each other by maximizing the likelihood of knowledge reconstruction. Especially, we present prior and posterior knowledge tracking in DukeL to solve the incompatible dual processes issue during training and inference. To the best of our knowledge, we are the first introduce the idea of dual learning into KGC.

3 METHOD

3.1 KGC formulation

We use $D = \{(X_\tau, Y_\tau)\}_{\tau=1}^{|D|}$ to denote the set of conversation turns, where (X_τ, Y_τ) is a conversation turn from two distinct speakers. At turn τ , given the conversation context $C_\tau = (X_{\tau-1}, Y_{\tau-1}, X_\tau)$ (we use a single previous turn and the current turn from the first speaker X_τ as conversation context), and corresponding knowledge pool $K_\tau = \{K_{\tau,j}\}_{j=1}^{|K_\tau|}$ with $|K_\tau|$ knowledge sentences (which are retrieved w.r.t. the conversation context), the task of KGC is to generate a response $Y_\tau = (y_{\tau,1}, y_{\tau,2}, \dots, y_{\tau,|Y_\tau|})$ with $|Y_\tau|$ tokens. With sequence-to-sequence modeling, this can be formulated as follows:

$$P(Y_\tau | C_\tau, K_\tau) = \prod_{t=1}^{|Y_\tau|} P(y_{\tau,t} | y_{\tau,<t}, C_\tau, K_\tau), \quad (1)$$

where $y_{\tau,t}$ is the t -th token; $y_{\tau,<t}$ are the tokens up to the $(t-1)$ -th decoding step.

3.2 Overview of DukeNet

As shown in Fig. 2, DukeNet consists of three layers: (1) an *encoding layer*, (2) a *dual knowledge interaction layer*, and (3) a *decoding layer*.

The encoding layer employs a BERT encoder to encode context C_τ and knowledge pool K_τ into context and knowledge representations $\mathbf{H}_p^{C_\tau} = (\mathbf{h}_p^{X_{\tau-1}}, \mathbf{h}_p^{Y_{\tau-1}}, \mathbf{h}_p^{X_\tau})$ and $\mathbf{H}_p^{K_\tau} = \{\mathbf{h}_p^{K_{\tau,j}}\}_{j=1}^{|K_\tau|}$, respectively.

The dual knowledge interaction layer contains a *Prior Knowledge Tracker* (Pri), a *Knowledge Shifter* (Shi), and a *Posterior Knowledge Tracker* (Pos). The prior knowledge tracker takes context $\mathbf{h}_p^{X_{\tau-1}}$, $\mathbf{h}_p^{Y_{\tau-1}}$ as inputs to predict a prior knowledge tracking distribution $P(K_{\tau-1}|pri)$ over knowledge pool $K_{\tau-1}$, based on which we can get the tracked knowledge representation $\mathbf{h}_p^{K_{\tau-1,j}}$. The knowledge shifter takes context $\mathbf{h}_p^{X_\tau}$ and tracked knowledge $\mathbf{h}_p^{K_{\tau-1,j}}$ as inputs to predict knowledge shifting distribution $P(K_\tau|K_{\tau-1,j}, shi)$ over knowledge pool K_τ , based on which we can get the shifted knowledge representation $\mathbf{h}_p^{K_{\tau,j}}$. The posterior knowledge tracker takes the shifted knowledge $\mathbf{h}_p^{K_{\tau,j}}$ as input to predict posterior knowledge tracking distribution $P(K_{\tau-1}|K_{\tau,j}, pos)$ over knowledge pool $K_{\tau-1}$. Note that the posterior knowledge tracker is only used during training.

The *decoding layer* contains a Transformer decoder to generate response Y_τ token by token based on the context representation and the shifted knowledge representation.

During training, we devise an unsupervised learning scheme, DukeL, that regards the posterior knowledge tracker and knowledge shifter as dual tasks. Specifically, for optimizing knowledge shifter, we feed the ground truth tracked knowledge $K_{\tau-1,l}$ (l denotes ground truth) to it, and then the shifted knowledge $K_{\tau,s}$ (s denotes sampling) that is sampled from the knowledge shifting distribution $P(K_\tau|K_{\tau-1,j}, shi)$ can be fed back into the posterior knowledge tracker to recover the tracked knowledge $K_{\tau-1,l}$ (the original input of knowledge shifter). We regard the recovering probability as a reward to optimize the knowledge shifter. The posterior knowledge tracker can be optimized in a similar manner. The above process forms a closed loop to alternatively train the posterior knowledge tracker and knowledge shifter. Meanwhile, we force the prior knowledge distribution $P(K_{\tau-1}|pri)$ (from the prior knowledge tracker) to get close to the posterior knowledge distribution $P(K_{\tau-1}|K_{\tau,j}, pos)$ (from the posterior knowledge tracker) via Kullback-Leibler Divergence Loss (KLDivLoss) such that the prior knowledge tracker can benefit from the above dual learning process even if it is not involved in the closed loop.

During inference, we only execute the prior knowledge tracker and knowledge shifter to do knowledge tracking and shifting, respectively. Hence the issue of incompatible dual processes that arises during training and inference can be solved effectively. And based on the shifted knowledge, we generate the next response. Next, we introduce the three layers and the learning scheme.

3.3 Encoding layer

We encode the conversation context $C_\tau = (X_{\tau-1}, Y_{\tau-1}, X_\tau)$ into hidden representations $\mathbf{H}_p^{C_\tau} = (\mathbf{h}_p^{X_{\tau-1}}, \mathbf{h}_p^{Y_{\tau-1}}, \mathbf{h}_p^{X_\tau})$ using BERT_{base} and an average pooling operation [2]:

$$\begin{aligned} \mathbf{H}^{X_{\tau-1}} &= \text{BERT}(X_{\tau-1}) \in \mathbb{R}^{|X_{\tau-1}| \times d}, \mathbf{h}_p^{X_{\tau-1}} = \text{p}(\mathbf{H}^{X_{\tau-1}}) \in \mathbb{R}^d, \\ \mathbf{H}^{Y_{\tau-1}} &= \text{BERT}(Y_{\tau-1}) \in \mathbb{R}^{|Y_{\tau-1}| \times d}, \mathbf{h}_p^{Y_{\tau-1}} = \text{p}(\mathbf{H}^{Y_{\tau-1}}) \in \mathbb{R}^d, \\ \mathbf{H}^{X_\tau} &= \text{BERT}(X_\tau) \in \mathbb{R}^{|X_\tau| \times d}, \mathbf{h}_p^{X_\tau} = \text{p}(\mathbf{H}^{X_\tau}) \in \mathbb{R}^d, \end{aligned} \quad (2)$$

Algorithm 1 Dual knowledge interaction learning algorithm

```

1: Warm-up training (Phase 1) ▷ See Eq. 12
2: for each iteration  $i$  in dual interaction training do (Phase 2)
3:   Sample an example from training set;
4:   ▷ Start to train knowledge shifter
5:   Feed  $\mathbf{h}_p^{K_{\tau-1,l}}$  to knowledge shifter;
6:   Get knowledge shifting distribution  $P(K_\tau | K_{\tau-1,l}, shi)$ ;
7:   Sample  $K_{\tau,s}$  from  $P(K_\tau | K_{\tau-1,l}, shi)$ ;
8:   Feed  $\mathbf{h}_p^{K_{\tau,s}}$  to posterior knowledge tracker to get reward;
9:   Update  $\theta_{shi}$  using reward;
10:  ▷ Start to train posterior knowledge tracker
11:  Feed  $\mathbf{h}_p^{K_{\tau,l}}$  to posterior knowledge tracker;
12:  Get posterior knowledge distribution  $P(K_{\tau-1} | K_{\tau,l}, pos)$ ;
13:  Sample  $K_{\tau-1,s}$  from  $P(K_{\tau-1} | K_{\tau,l}, pos)$ ;
14:  Feed  $\mathbf{h}_p^{K_{\tau-1,s}}$  to knowledge shifter to get reward;
15:  Update  $\theta_{pos}$  using reward;
16:  ▷ Start reduce the distance between the prior and posterior
    knowledge distribution  $P(K_{\tau-1}|pri)$  and  $P(K_{\tau-1}|K_{\tau,l}, pos)$ 
17:  Update  $\theta$  using KLDivLoss and MLE.
18: end for

```

where d stands for hidden size and p refers to pooling operation. Similarly, we encode the knowledge sentences in knowledge pool $K_\tau = \{K_{\tau,j}\}_{j=1}^{|K_\tau|}$ into representations $\mathbf{H}_p^{K_\tau} = \{\mathbf{h}_p^{K_{\tau,j}}\}_{j=1}^{|K_\tau|} \in \mathbb{R}^{|K_\tau| \times d}$.

3.4 Dual knowledge interaction layer

3.4.1 Prior knowledge tracker. Given context representations $\mathbf{h}_p^{X_{\tau-1}}$ and $\mathbf{h}_p^{Y_{\tau-1}}$, the prior knowledge tracker predicts the prior knowledge tracking distribution $P(K_{\tau-1}|pri)$ over the knowledge pool $K_{\tau-1}$, which is estimated as follows:

$$\begin{aligned} P(K_{\tau-1}|pri) &= \text{softmax}(\mathbf{Q}_{pri} \mathbf{K}_{pri}^T) \in \mathbb{R}^{1 \times |K_{\tau-1}|} \\ \mathbf{Q}_{pri} &= \text{mlp}([\mathbf{h}_p^{X_{\tau-1}}; \mathbf{h}_p^{Y_{\tau-1}}]) \in \mathbb{R}^{1 \times d} \\ \mathbf{K}_{pri} &= \text{mlp}(\mathbf{H}_p^{K_{\tau-1}}) \in \mathbb{R}^{|K_{\tau-1}| \times d}, \end{aligned} \quad (3)$$

where $\text{mlp}(\cdot) = \cdot \mathbf{W} + \mathbf{b}$ is a Multilayer Perceptron (MLP) and $;$ denotes the vector concatenation operation.

3.4.2 Knowledge shifter. Given context representation $\mathbf{h}_p^{X_\tau}$ and tracked knowledge representation $\mathbf{h}_p^{K_{\tau-1,j}}$, the knowledge shifter predicts the shifting knowledge distribution $P(K_\tau|K_{\tau-1,j}, shi)$ over knowledge pool K_τ , which is estimated as follows:

$$\begin{aligned} P(K_\tau|K_{\tau-1,j}, shi) &= \text{softmax}(\mathbf{Q}_{shi} \mathbf{K}_{shi}^T) \in \mathbb{R}^{1 \times |K_\tau|} \\ \mathbf{Q}_{shi} &= \text{mlp}([\mathbf{h}_p^{X_\tau}; \mathbf{h}_p^{K_{\tau-1,j}}]) \in \mathbb{R}^{1 \times d} \\ \mathbf{K}_{shi} &= \text{mlp}(\mathbf{H}_p^{K_\tau}) \in \mathbb{R}^{|K_\tau| \times d}. \end{aligned} \quad (4)$$

3.4.3 Posterior knowledge tracker. Given context representations $\mathbf{h}_p^{X_{\tau-1}}$ and $\mathbf{h}_p^{Y_{\tau-1}}$ and shifted knowledge representation $\mathbf{h}_p^{K_{\tau,j}}$, the posterior knowledge tracker predicts the posterior knowledge tracking distribution $P(K_{\tau-1}|K_{\tau,j}, pos)$ over knowledge pool $K_{\tau-1}$, which

is estimated as follows:

$$\begin{aligned} P(K_{\tau-1}|K_{\tau,j}, pos) &= \text{softmax}(\mathbf{Q}_{pos}\mathbf{K}_{pos}^\top) \in \mathbb{R}^{1 \times |K_{\tau-1}|} \\ \mathbf{Q}_{pos} &= \text{mlp}([\mathbf{h}_p^{X_{\tau-1}}; \mathbf{h}_p^{Y_{\tau-1}}; \mathbf{h}_p^{K_{\tau,j}}]) \in \mathbb{R}^{1 \times d} \\ \mathbf{K}_{pos} &= \text{mlp}(\mathbf{H}_p^{K_{\tau-1}}) \in \mathbb{R}^{|K_{\tau-1}| \times d}. \end{aligned} \quad (5)$$

3.5 Decoding layer

We feed $[\mathbf{H}^{X_\tau}; \mathbf{H}^{K_{\tau,j}}]$ into a transformer decoder [42] equipped with a copying mechanism [10, 26, 33] to generate Y_τ token by token, where $\mathbf{H}^{X_\tau} \in \mathbb{R}^{|X_\tau| \times d}$ is the context representation before pooling (see Eq. 2) and $\mathbf{H}^{K_{\tau,j}} \in \mathbb{R}^{|K_{\tau,j}| \times d}$ is the shifted knowledge representation before pooling based on the prediction from the knowledge shifter. During the training process, we always give the representations of the ground truth shifted knowledge $\mathbf{H}^{K_{\tau,l}}$ to the decoder, where l refer to the ground truth label. Specifically, the probability of generating $y_{\tau,t}$ at t is modeled as:

$$P(y_{\tau,t}) = P(g)P(y_{\tau,t}|g) + P(c_c)P(y_{\tau,t}|c_c) + P(c_k)P(y_{\tau,t}|c_k), \quad (6)$$

where $P(y_{\tau,t}|g)$ is the probability of generating a token from the predefined vocabulary V :

$$P(y_{\tau,t}|g) = \text{mlp}(\mathbf{h}_{\tau,t}) \in \mathbb{R}^{1 \times |V|}, \quad (7)$$

where $\mathbf{h}_{\tau,t} = \text{TransformerDecoder}(\text{emb}(y_{\tau,<t}), [\mathbf{H}^{X_\tau}; \mathbf{H}^{K_{\tau,j}}]) \in \mathbb{R}^d$; $\text{TransformerDecoder}$ is a stack of Transformer decoder blocks [42]; $\text{emb}(y_{\tau,<t})$ denotes the embedding of $y_{\tau,<t}$.

$P(y_{\tau,t}|c_c)$ is the probability copying a token from the context X_τ :

$$P(y_{\tau,t}|c_c) = \sum_{i: x_{\tau,i}=y_{\tau,t}} \alpha_{\tau,t,i}^c \in \mathbb{R}^{1 \times |X_\tau|}, \quad (8)$$

where $x_{\tau,i}$ is the i -th token in X_τ and $\alpha_{\tau,t,i}^c$ is the attention distribution on X_τ with $\mathbf{h}_{\tau,t}$ attentively reading $\mathbf{H}^{X_\tau}$ (see Eq.10). $P(y_{\tau,t}|c_k)$ is the probability copying a token from the knowledge $K_{\tau,j}$, which is calculated in a similar way.

$P(g)$, $P(c_c)$, and $P(c_k)$ are the coordination probabilities among the above three modes: g , c_c and c_k , which are estimated as follows:

$$[P(g), P(c_c), P(c_k)] = \text{softmax}(\text{mlp}([\mathbf{h}_{\tau,t}; \mathbf{c}_{\tau,t}^g; \mathbf{c}_{\tau,t}^k])) \in \mathbb{R}^{1 \times 3}, \quad (9)$$

where $\mathbf{c}_{\tau,t}^g$ and $\mathbf{c}_{\tau,t}^k$ are attention vectors derived from $\mathbf{h}_{\tau,t}$ attending to $\mathbf{H}^{X_\tau}$ and $\mathbf{H}^{K_{\tau,j}}$, respectively. Finally, $\mathbf{c}_{\tau,t}^g$ is calculated as follows:

$$\begin{aligned} \mathbf{c}_{\tau,t}^g &= \alpha_{\tau,t}^g \mathbf{H}^{X_\tau} \in \mathbb{R}^{1 \times d} \\ \alpha_{\tau,t}^g &= \text{softmax}(\mathbf{Q}_g \mathbf{K}_g^\top) \in \mathbb{R}^{1 \times |X_\tau|} \\ \mathbf{Q}_g &= \text{mlp}(\mathbf{h}_{\tau,t}) \in \mathbb{R}^{1 \times d}, \mathbf{K}_g = \text{mlp}(\mathbf{H}^{X_\tau}) \in \mathbb{R}^{|X_\tau| \times d}. \end{aligned} \quad (10)$$

And $\mathbf{c}_{\tau,t}^k$ is calculated in a similar way.

3.6 Dual knowledge interaction learning

We devise a DukeL scheme to learn DukeNet, which can be divided into two phases: *warm-up training phase* and *dual interaction training phase*, as shown in Algorithm 1.

3.6.1 Warm-up training phase. We first employ the commonly used MLE loss to maximize the likelihood of the demonstrated examples in the training set [14]:

$$\begin{aligned} \mathcal{L}_{pri}(\theta) &= -\log P(K_{\tau-1,l} | pri) \\ \mathcal{L}_{pos}(\theta) &= -\log P(K_{\tau-1,l} | K_{\tau,l}, pos) \\ \mathcal{L}_{shi}(\theta) &= -\log P(K_{\tau,l} | K_{\tau-1,l}, shi) \\ \mathcal{L}_g(\theta) &= -\sum_{t=1}^{|Y_\tau|} \log P(y_{\tau,t} | y_{\tau,<t}, X_\tau, K_{\tau,l}), \end{aligned} \quad (11)$$

where θ are all the parameters of DukeNet and l refer to the ground truth label. $\mathcal{L}_{pri}(\theta)$ is the prior tracking loss; $\mathcal{L}_{pos}(\theta)$ is the posterior tracking loss; $\mathcal{L}_{shi}(\theta)$ is the shifting loss; and $\mathcal{L}_g(\theta)$ is the generation loss. The *final loss* is the sum of the four functions just defined:

$$\mathcal{L}(\theta) = \mathcal{L}_{pri}(\theta) + \mathcal{L}_{pos}(\theta) + \mathcal{L}_{shi}(\theta) + \mathcal{L}_g(\theta). \quad (12)$$

3.6.2 Dual interaction training phase. For each iteration, given an example sampled from the training set, we first optimize the knowledge shifter. We feed the representation of the ground truth tracked knowledge $\mathbf{h}_p^{K_{\tau-1,l}}$ to the knowledge shifter to get the knowledge shifting distribution $P(K_\tau | K_{\tau-1,l}, shi)$ (Line 5–6 in Algorithm 1). Then we sample knowledge $K_{\tau,s}$ from $P(K_\tau | K_{\tau-1,l}, shi)$ and feed its representation $\mathbf{h}_p^{K_{\tau,s}}$ to the posterior knowledge tracker to get the probability of recovering the ground truth tracked knowledge $K_{\tau-1,l}$ (Line 7–8 in Algorithm 1). We regard this recovering probability as a reward:

$$\begin{aligned} \mathbb{E}[R] &= \mathbb{E}[R \log P(K_{\tau,s} | K_{\tau-1,l}, shi)] \\ R &= \log[P(K_{\tau-1,l} | K_{\tau,s}, pos)], \end{aligned} \quad (13)$$

where R is the reward of the sampled knowledge $K_{\tau,s}$. After that, we use policy gradients [39] to maximize $\mathbb{E}[R]$ and compute the gradient for the parameters θ_{shi} (Line 9 in Algorithm 1):

$$\nabla_{\theta_{shi}} \mathbb{E}[R] = \mathbb{E}[R \nabla_{\theta_{shi}} \log P(K_{\tau,s} | K_{\tau-1,l}, shi)]. \quad (14)$$

We then start to optimize the posterior knowledge tracker (the parameters θ_{pos}), which is done in a similar way (Line 11–15 in Algorithm 1). After above dual process, we feed the representation of the ground truth shifted knowledge $\mathbf{h}_p^{K_{\tau,l}}$ to the optimized posterior knowledge tracker to again predict the posterior knowledge tracking distribution $P(K_{\tau-1} | K_{\tau,l}, pos)$. We then distill the gains from $P(K_{\tau-1} | K_{\tau,l}, pos)$ to $P(K_{\tau-1} | pri)$ via KLDivLoss (Line 17 in Algorithm 1) $\mathcal{L}_{kl}(\theta) =$

$$\sum_{j=1}^{|K_{\tau-1}|} P(K_{\tau-1,j} | K_{\tau,l}, pos) \log \frac{P(K_{\tau-1,j} | K_{\tau,l}, pos)}{P(K_{\tau-1,j} | pri)}. \quad (15)$$

To reduce the impact of inaccurate reward estimation [23], we combine $\mathcal{L}_{kl}(\theta)$ with the MLE losses $[\mathcal{L}_{pos}(\theta), \mathcal{L}_{shi}(\theta), \mathcal{L}_g(\theta)]$ linearly and jointly train them:

$$\mathcal{L}(\theta) = \mathcal{L}_{kl}(\theta) + \lambda[\mathcal{L}_{pos}(\theta) + \mathcal{L}_{shi}(\theta) + \mathcal{L}_g(\theta)], \quad (16)$$

where λ is a hyper-parameter to control the effect of MLE. We repeat the above process until convergence.

4 EXPERIMENTS

4.1 Research questions

We aim to answer the following research questions:

- (RQ1) What is the performance of DukeNet? Does DukeNet outperform state-of-the-art methods? (See §5.1 and §5.2)
- (RQ2) Where does the improvements of DukeNet come from? How do the different components contribute to its performance? (See §6.1)
- (RQ3) Does the dual knowledge interaction improve *knowledge tracking* and *knowledge shifting* jointly? (See §6.2)
- (RQ4) Is DukeNet able to generate better responses? Are there any failures? (See §6.3)

4.2 Dataset

Following Kim et al. [14], we evaluate our model on two public KGC benchmark datasets, Wizard of Wikipedia [7] and Holl-E [25], both of which contain the ground truth labels for KS. We split the data into training, validation and test as per the original papers.

Wizard of Wikipedia is the largest unstructured KGC dataset that is based on sentences to date. The conversations are conducted between two speakers about some given open-domain topics. The one speaker acts as the wizard (expert), who can use a retrieval system to acquire knowledge sentences from Wikipedia and chooses any to form a response. The other speaker acts as the apprentice (learner), who is eager to talk with the wizard about a topic but does not have access to external knowledge. It contains 18,430, 1,948 and 1,933 conversations for training, validation and test, respectively. The test set is further split into two subsets, Test Seen and Test Unseen. The former contains conversations on topics overlapping with topics in the training set, and the latter contains conversations on topics never seen in the training and validation set. The average number of sentences in a knowledge pool is about 67.

Holl-E is a document based dataset, i.e., a single document is given as knowledge per conversation. Kim et al. [14] have changed it to a sentence based one and modified the ground truth labels for KS, so we use the version released by them. It contains 7,228, 930 and 913 conversations for training, validation and test, respectively. There are two versions of the test set: one with a single reference and the other with multiple references (more than one ground truth knowledge sentences and corresponding responses for each given conversation context). The average number of sentences in a knowledge pool is roughly 60.

4.3 Baselines

We compare DukeNet with state-of-the-art KGC methods in the sentence based setting.

- **Seq2Seq** [38] maps the conversation context into the response with an encoder-decoder framework, which does not use any knowledge information.
- **Transformer** [42] implements an encoder-decoder framework by solely relying on multi-head attention mechanism and dispensing with recurrence, which does not use any knowledge information either.
- **MemNet** [9] combines a Seq2Seq model with an external memory network to store knowledge.

- **TMemNet** [7] combines a transformer model with an external memory network in an end-to-end manner, which further introduces an auxiliary loss to better supervise knowledge selection.
- **PostKS** [19] uses response and conversation context to jointly form a posterior knowledge distribution and regards it as pseudo-labels to supervise KS.
- **SKT** [14] sequentially models the history of KS at previous turns via a sequential latent variable model [35]. In addition, SKT uses BERT to encode conversation context and knowledge and incorporates a copying mechanism [10, 33] to promote response generation.

For a fair comparison, we also report the results of PostKS+BERT and TMemNet+BERT, which also use BERT as encoders.

4.4 Evaluation metrics

We conduct both automatic and human evaluations. For automatic evaluation we follow previous KGC studies [24, 51], and use BLEU-4 [27], METEOR [5], ROUGE-1, ROUGE-2, and ROUGE-L [20] for evaluating response generation. In addition, we report Hit@1 (the top 1 accuracy) to evaluate knowledge selection [14, 22] at each turn. For the evaluation of multiple references in the Holl-E dataset, we follow Moghe et al. [25] and Kim et al. [14]. For evaluating RG, we take the max score between responses generated by models and multiple ground truth responses. For evaluating KS, we regard the knowledge selected by model as correct if it matches any of the ground truth knowledge.

For human evaluation we randomly sample 300 examples from each test set. For each example, we ask three workers to conduct a pairwise comparison between the knowledge sentences/responses selected/generated by DukeNet and the ones selected/generated by a baseline on Amazon Mechanical Turk¹. Specifically, given the conversation context, the knowledge pool² used at the current turn, the selected knowledge, as well as the generated responses, each worker needs to give a preference (ties are allowed) in terms of three aspects: (1) *Appropriateness*, i.e., which selected knowledge is more appropriate/relevant w.r.t. the given conversation context; (2) *Informativeness*, i.e., which response looks more informative [14, 22, 52]; (3) *Engagingness*, i.e., which response is better in general [14]. Model names were masked out during evaluation.

4.5 Implementation details

For models without BERT encoder, the word embedding size and hidden size are all set to 256. For models with BERT encoder, we use BERT-Base pre-trained weights and the hidden size is 768. We use the Adam optimizer [15] to train all models. In particular, we train DukeNet model for 10 epochs for the warm-up training phase (learning rate 0.00002) and 5 epochs for the dual interaction training phase (learning rate 0.00001). λ in Eq. 16 is set to 0.5. We use gradient clipping with a maximum gradient norm of 0.4. We use the BERT vocabulary³ (the size is 30,522) for all models. We select the best models based on performance on the validation set. The code is available online⁴.

¹<https://www.mturk.com/>

²To reduce the burden of workers, we only show them a reduced pool including at most 10 knowledge sentences, where the ones selected by the models are in it.

³<https://github.com/huggingface/transformers>

⁴<https://github.com/ChuanMeng/DukeNet/>

Table 1: Automatic evaluation results on the Wizard of Wikipedia dataset. Bold face indicates the best result in terms of the corresponding metric. Significant improvements over the best baseline results are marked with * (t-test, $p < 0.05$).

Methods	Test Seen (%)						Test Unseen (%)					
	BLEU-4	METEOR	ROUGE-1	ROUGE-2	ROUGE-L	Hit@1	BLEU-4	METEOR	ROUGE-1	ROUGE-2	ROUGE-L	Hit@1
Seq2Seq	0.46	12.22	20.32	3.03	14.46	–	0.34	10.21	19.01	2.16	13.55	–
Transformer	0.39	12.82	20.50	3.27	15.00	–	0.39	11.36	18.81	2.17	14.33	–
MemNet	0.41	12.35	21.51	3.17	15.33	4.27	0.32	11.75	20.06	2.51	14.68	4.13
PostKS	0.57	13.50	21.61	3.66	16.07	4.70	0.36	12.60	20.79	2.52	14.88	4.45
TMemNet	1.35	14.52	22.84	4.31	16.77	21.57	0.43	12.82	21.33	3.05	15.39	12.10
PostKS + BERT	0.77	14.16	22.68	4.27	16.59	4.83	0.39	12.59	20.82	2.73	15.25	4.39
TMemNet + BERT	1.61	15.47	24.12	4.98	17.00	23.86	0.60	13.05	21.74	3.63	15.60	16.33
SKT	1.76	16.04	24.61	5.24	17.61	25.36	1.05	13.74	22.84	4.40	16.05	18.19
DukeNet	2.43*	17.09*	25.17	6.81*	18.52*	26.38	1.68*	15.06*	23.34	5.29*	17.06*	19.57

Table 2: Automatic evaluation results on the Holl-E dataset.

Methods	Single golden reference (%)						Multiple golden references (%)					
	BLEU-4	METEOR	ROUGE-1	ROUGE-2	ROUGE-L	Hit@1	BLEU-4	METEOR	ROUGE-1	ROUGE-2	ROUGE-L	Hit@1
Seq2Seq	4.84	17.12	26.25	8.41	20.08	–	7.41	21.47	30.68	12.01	24.58	–
Transformer	5.09	16.39	25.96	8.62	19.64	–	7.58	21.01	30.43	12.25	24.60	–
MemNet	5.49	17.70	26.88	9.51	21.15	3.39	7.75	21.60	31.63	12.21	24.94	5.32
PostKS	5.85	18.53	27.52	9.21	21.23	3.60	8.01	22.23	31.57	12.55	25.15	5.95
TMemNet	6.77	20.67	28.25	9.97	22.37	24.15	8.98	25.29	32.46	13.05	26.37	33.95
PostKS + BERT	6.54	19.30	28.94	9.89	22.15	3.95	8.49	23.97	32.85	13.10	26.17	6.40
TMemNet + BERT	8.99	24.48	31.65	13.24	25.90	28.44	12.36	28.61	35.29	16.14	29.51	37.30
SKT	17.81	29.41	35.28	21.74	30.06	28.99	24.69	35.78	41.68	28.30	36.24	39.05
DukeNet	19.15*	30.93*	36.53	23.02*	31.46*	30.03	26.83*	37.73*	43.18*	30.13*	38.03*	40.33

5 EXPERIMENTAL RESULTS

5.1 Automatic evaluation (RQ1)

We list the results of all methods on both the Wizard of Wikipedia and Holl-E datasets in Tables 1 and 2, respectively. Generally, DukeNet significantly outperforms all baselines on both datasets. From the results, we have three main observations.

First, DukeNet outperforms the strongest baseline SKT by around 1–1.5% in terms of Hit@1 on both datasets. In particular, the improvement of DukeNet over SKT on the test unseen is 1.38%, while on test seen it is 1.02%, which means that DukeNet can better handle unseen cases. The reason is that SKT only models the unidirectional interaction from knowledge tracking to shifting and merely uses demonstrated examples in the training set to optimize model via MLE. In contrast, DukeNet benefits from DukeL, which regards knowledge tracking and shifting as dual tasks to let them boost each other. The learning process of DukeL explores extra knowledge besides the demonstrated ground truth in the training set, which improves the generalization ability on unseen cases.

Second, DukeNet substantially outperforms the other baselines that do not explicitly model knowledge tracking and shifting in terms of Hit@1, e.g., MemNet, PostKS and TMemNet. Especially, DukeNet outperforms the most competitive baseline TMemNet + BERT by around 1.5–3.5% on both datasets. The gains show that explicitly modeling knowledge tracking and shifting can better capture the interaction between the knowledge at adjacent turns.

In addition, the tracked knowledge can provide extra evidence and clues to infer the shifted knowledge compared to use context only, which narrows the search space for KS. DukeNet also outperforms all baselines, including SKT, in terms of Response Generation (RG), i.e., the BLEU-4, METEOR, ROUGE-1, ROUGE-2, and ROUGE-L scores are significantly improved. Note that nothing special is proposed for the decoder in DukeNet, so the higher scores on the generation metrics indicate that better KS performance of DukeNet also improves the quality of RG.

Third, interestingly, we found that the improvement from BERT is limited compared to that on QA tasks [6], e.g., BERT based models have been shown to be much more effective and already outperformed humans on the SQuAD dataset [30]. We think the reason is that BERT is pretrained on language modeling tasks, which will help improve the language modeling performance mostly. However, on the KGC task, the main bottleneck now is the KS, which BERT can only make a limited contribution.

5.2 Human evaluation (RQ1)

Although automatic evaluation metrics have been shown to be reliable on the KGC task [7, 24], we still conduct human evaluations to confirm the improvement of DukeNet.

We compare DukeNet with the three most competitive baselines TMemNet + BERT, PostKS + BERT, and SKT on the more challenging Wizard of Wikipedia dataset. The results are shown in Table 3.

Table 3: Human evaluation on the Wizard of Wikipedia dataset.

Methods	Test Seen (%)									Test Unseen (%)								
	Appropriateness			Informativeness			Engagingness			Appropriateness			Informativeness			Engagingness		
	Win	Tie	Lose	Win	Tie	Lose	Win	Tie	Lose	Win	Tie	Lose	Win	Tie	Lose	Win	Tie	Lose
DukeNet vs PostKS + BERT	51	46	3	74	23	3	59	36	5	74	25	1	70	26	4	78	21	1
DukeNet vs TMemNet + BERT	37	55	8	19	79	2	48	42	10	45	50	5	23	74	3	38	57	5
DukeNet vs SKT	15	79	6	15	80	5	16	77	7	20	76	4	15	81	4	21	68	11

Table 4: Ablation study on the Wizard of Wikipedia dataset. -DukeL denotes removing dual interaction training phase. -Pri and -Pos denote removing prior and posterior knowledge tracker, respectively. -KL denotes removing KLDivLoss in Eq. 15.

Methods	Test Seen (%)						Test Unseen (%)					
	BLEU-4	METEOR	ROUGE-1	ROUGE-2	ROUGE-L	Hit@1	BLEU-4	METEOR	ROUGE-1	ROUGE-2	ROUGE-L	Hit@1
Full model	2.43	17.09	25.17	6.81	18.52	26.38	1.68	15.06	23.34	5.29	17.06	19.57
-Pri,-KL	2.09	16.74	25.10	6.36	18.08	25.93	1.39	14.05	22.52	4.68	15.99	18.85
-DukeL	1.99	16.34	24.54	6.19	17.83	25.57	1.28	13.75	22.20	4.67	15.76	18.56
-DukeL,-Pri,-Pos,-KL	1.74	15.99	24.34	5.88	17.66	23.82	0.99	13.39	22.69	4.35	15.90	16.25

Generally, DukeNet achieves the best performance in terms of all metrics on both datasets. In particular, we find that the performance gaps between DukeNet and the other three baselines are more obvious on the test unseen subset. For instance, the wins of DukeNet over SKT is 15% and 20% on the test seen and test unseen in terms of Appropriateness, respectively, which is consistent with the automatic evaluation results and again indicates that exploring data besides the ground truth in DukeL can indeed promote the generalization ability. DukeNet is even better than SKT in terms of Informativeness, despite that the fact that they both use a copying mechanism [10] to make use of knowledge during decoding. We think that this is because, in many cases, SKT selects inappropriate knowledge for a given conversation context; the response generated based on this is more likely to be less relevant and will be considered as less informative by workers. DukeNet gets the best score in terms of Engagingness, which shows that the workers prefer the responses from DukeNet in general, mostly because DukeNet selects more appropriate knowledge which will result in more relevant and natural responses.

6 ANALYSIS

6.1 Ablation study (RQ2)

To analyze where the improvements of DukeNet come from, we conduct an ablation study on the Wizard of Wikipedia dataset; see Table 4. Here, we consider three settings. (1) No prior knowledge tracker (i.e., -Pri,-KL in Table 4), i.e. we directly use posterior knowledge tracker during inference, and feed zero vector to replace the shifted knowledge. (2) No dual interaction training phase in §3.6.2 (i.e., -DukeL in Table 4). (3) No knowledge tracking (i.e., -DukeL,-Pri,-Pos,-KL in Table 4).

The results show that all parts are helpful to DukeNet because removing any of them will decrease the results. Without knowledge tracking, the performance of DukeNet drops sharply in terms of all metrics, almost degenerating to TMemNet + BERT. Specifically, it drops around 2.5–3.5% in terms of Hit@1, which means that knowledge tracking is essential for KS, and modeling KS only by

Table 5: The results of knowledge tracking and shifting with or without dual interaction training phase on the Wizard of Wikipedia dataset. Tra-Hit@1 denotes the top 1 accuracy for knowledge tracking.

Methods	Test Seen (%)		Test Unseen (%)	
	Tra-Hit@1	Hit@1	Tra-Hit@1	Hit@1
-DukeL	81.66	25.57	76.04	18.56
Full model	82.77	26.38	76.82	19.57

modeling conversation context is far from enough. Without DukeL, the performance of DukeNet also drops a lot in terms of all metrics, almost degenerating to SKT. It drops by 0.81% and 1.01% in terms of Hit@1 in the test seen and test unseen conditions, respectively. This indicates that DukeL can not only model the dual interaction between knowledge tracking and shifting to improve them jointly but also alleviate the many-to-many mapping phenomenon in KGC via exploring the knowledge that is not limited to ground truth. Without the prior knowledge tracker, we find that the gain from DukeL is very limited, though it still slightly outperforms SKT. It only improves around 0.3% in terms of Hit@1 compared to the case without DukeL. The posterior knowledge tracker without shifted knowledge as input cannot perform well during inference due to the incompatible dual processes between training and inference, which restricts the effect of the knowledge shifter.

6.2 Dual knowledge interaction (RQ3)

To analyze whether dual knowledge interaction improves both knowledge tracking and knowledge shifting after the dual interaction training phase (see §3.6.2), we report the Tra-Hit@1 and Hit@1 with or without this phase, respectively. The results on the Wizard of Wikipedia dataset are shown in Table 5.

We see that the performances improve on both the test seen and test unseen for knowledge tracking and shifting, which indicates that they two indeed teach each other during this process. Thus, during inference, the improved knowledge tracking can ground tracked

Table 6: Case study. Due to the space limitations, we only show one merged pool ($K_{\tau-1}+K_{\tau}$) and three knowledge sentences.

	Example 1 (Test seen)	Example 2 (Test unseen)
Context	$X_{\tau-1}$: pizza delivery $Y_{\tau-1}$: for dinner i had pizza delivered to my house by a pizzeria X_{τ} : love a good pizzeria or restaurant that specializes in pissa	$X_{\tau-1}$: what can you tell me about the story of harry potter? $Y_{\tau-1}$: harry potter is a 7 book fantasy series written by jk rowling X_{τ} : 7 books ! wow
Knowledge pool	K_1 : pizza delivery is a service in which a pizzeria or pizza chain delivers a pizza to a customer . K_2 : pizzas may be delivered in pizza boxes or delivery bags , and deliveries are made with either an automobile . K_3 : an order is typically made either by telephone or over the internet to the pizza chain, in which the customer can request pizza type , motorized scooter	K_1 : harry potter is a series of fantasy novels written by british author j. k. rowling . K_2 : the main story arc concerns harry ' s struggle against lord voldemort , a dark wizard who intends to become immortal . K_3 : the books have sold more than 500 million copies worldwide , making them the best-selling book series in history
Tracking to shifting	PostKS + BERT: $\emptyset \rightarrow K_1$ ✗ TMemeNet + BERT: $\emptyset \rightarrow K_1$ ✗ SKT: $K_2 \rightarrow K_2$ ✗ DukeNet: $K_1 \rightarrow K_3$ ✓	PostKS + BERT: $\emptyset \rightarrow K_1$ ✗ TMemeNet + BERT: $\emptyset \rightarrow K_1$ ✗ SKT: $K_1 \rightarrow K_3$ ✗ DukeNet: $K_1 \rightarrow K_2$ ✓
Response	PostKS + BERT: i love pizza delivery . i love pizza delivery . TMemeNet + BERT: i can do it with my pizza delivery , but i can pay my pizza at home . SKT: i love frozen pizza ! i love the delivery bagels with either an automobile , or a motorized scooter . DukeNet: i love pizza delivery and i could pay online or online ordering .	PostKS + BERT: i know that harry potter is a series of fantasy novels . TMemeNet + BERT: i 've never read the book, but i know that harry potter is a series of fantasy novels written by british author j. k. rowling . SKT: they have sold more than 500 million copies worldwide . DukeNet: the main story is about harry ' s struggle against lord voldemort .

knowledge more accurately, and the more accurately tracked knowledge can be fed to the improved knowledge shifting to select more appropriate knowledge, which eventually boosts the quality of RG.

6.3 Case study (RQ4)

We randomly select examples from the Wizard of Wikipedia test sets to compare the performance of DukeNet, SKT, TMemNet+BERT and PostKS+BERT in Table 6. We can see that DukeNet conducts knowledge tracking and shifting more precisely, and hence results in more engaging responses. For instance, in Example 1, DukeNet captures the knowledge interaction from the general definition of “Pizza delivery” to its online ordering. In Example 2, DukeNet shifts the topic from “harry potter novel” to its main story content properly. In contrast, in Example 1, SKT messes up the knowledge K_1 and K_2 , and makes a mistake by considering K_2 as the tracked knowledge. The possible reason is that SKT does not explicitly distinguish knowledge tracking and shifting, and uses the unidirectional interaction from knowledge tracking to shifting, which cannot utilize the dual interaction between them to further improve both. TMemNet+BERT and PostKS+BERT get the worst performances in the two examples, which again shows that only using context to model KS is far from enough. We also observe a few bad cases of DukeNet. Although DukeNet alleviates the problem of the many-to-many mapping phenomenon to a certain extent, we found that DukeNet is still more likely to make mistakes when given a tracked knowledge sentence that can be mapped to multiple reasonable shifted knowledge ones. This is because the knowledge interaction between two adjacent turns can only provide limited hints, which fails to model the long-term dual interactions between knowledge tracking and shifting.

7 CONCLUSION AND FUTURE WORK

In this paper, we propose DukeNet, which explicitly models knowledge tracking and knowledge shifting as dual tasks to improve performance on the Knowledge Selection (KS) task. We also devise an unsupervised DukeL to explore knowledge beyond the demonstrated ones in the dataset during training for KS. Extensive experiments on two benchmark datasets show that DukeNet achieves new state-of-the-art performance, indicating that DukeNet enhanced by DukeL can select more appropriate knowledge and hence generate more informative and engaging responses.

A limitation of DukeNet is that it only considers the dual knowledge interaction between two adjacent conversation turns. In future work, we plan to extend the dual knowledge interaction to complete conversations, such that we can leverage long-term dual interactions between knowledge tracking and shifting to further improve Knowledge Selection (KS).

ACKNOWLEDGMENTS

We want to thank our anonymous reviewers for their feedback.

This work was partially supported by the Natural Science Foundation of China (61972234, 61902219, 61672324, 61672322), the Key Scientific and Technological Innovation Program of Shandong Province (2019JZZY010129), the Tencent AI Lab Rhino-Bird Focused Research Program (JR201932), the Fundamental Research Funds of Shandong University, and the Innovation Center for Artificial Intelligence (ICAI).

All content represents the opinion of the authors, which is not necessarily shared or endorsed by their respective employers and/or sponsors.

REFERENCES

- [1] Mohammad Aliannejadi, Hamed Zamani, Fabio Crestani, and W Bruce Croft. 2019. Asking Clarifying Questions in Open-domain Information-seeking Conversations. In *SIGIR*. 475–484.
- [2] Daniel Cer, Yinfei Yang, Sheng-yi Kong, Nan Hua, Nicole Limtiaco, Rhomni St John, Noah Constant, Mario Guajardo-Cespedes, Steve Yuan, Chris Tar, et al. 2018. Universal Sentence Encoder. *arXiv preprint arXiv:1803.11175* (2018).
- [3] Hongshen Chen, Zhaochun Ren, Jiliang Tang, Yihong Eric Zhao, and Dawei Yin. 2018. Hierarchical variational memory network for dialogue generation. In *WWW*. 1653–1662.
- [4] Shaobo Cui, Rongzhong Lian, Di Jiang, Yuanfeng Song, Siqi Bao, and Yong Jiang. 2019. DAL: Dual Adversarial Learning for Dialogue Generation. In *Proceedings of the Workshop on Methods for Optimizing and Evaluating Neural Language Generation*. 11–20.
- [5] Michael Denkowski and Alon Lavie. 2014. Meteor Universal: Language Specific Translation Evaluation for Any Target Language. In *Proceedings of the ninth workshop on statistical machine translation*. 376–380.
- [6] Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding. In *NAACL*. 4171–4186.
- [7] Emily Dinan, Stephen Roller, Kurt Shuster, Angela Fan, Michael Auli, and Jason Weston. 2019. Wizard of Wikipedia: Knowledge-Powered Conversational Agents. In *ICLR*.
- [8] Jianfeng Gao, Michel Galley, Lihong Li, et al. 2019. Neural Approaches to Conversational AI. *Foundations and Trends in Information Retrieval* 13, 2-3 (2019), 127–298.
- [9] Marjan Ghazvininejad, Chris Brockett, Ming-Wei Chang, Bill Dolan, Jianfeng Gao, Wen-tau Yih, and Michel Galley. 2018. A knowledge-grounded neural conversation model. In *AAAI*. 5110–5117.
- [10] Jiatao Gu, Zhengdong Lu, Hang Li, and Victor OK Li. 2016. Incorporating Copying Mechanism in Sequence-to-Sequence Learning. In *ACL*. 1631–1640.
- [11] Di He, Yingce Xia, Tao Qin, Liwei Wang, Nenghai Yu, Tie-Yan Liu, and Wei-Ying Ma. 2016. Dual Learning for Machine Translation. In *NeurIPS*. 820–828.
- [12] Minlie Huang, Xiaoyan Zhu, and Jianfeng Gao. 2020. Challenges in Building Intelligent Open-domain Dialog Systems. *TOIS* (2020).
- [13] Xisen Jin, Wenqiang Lei, Zhaochun Ren, Hongshen Chen, Shangsong Liang, Yihong Zhao, and Dawei Yin. 2018. Explicit State Tracking with Semi-Supervision for Neural Dialogue Generation. In *CIKM*. 1403–1412.
- [14] Byeonchang Kim, Jaewoo Ahn, and Gunhee Kim. 2020. Sequential Latent Knowledge Selection for Knowledge-Grounded Dialogue. In *ICLR*.
- [15] Diederik P Kingma and Jimmy Ba. 2014. Adam: A Method for Stochastic Optimization. In *ICLR*.
- [16] Wenqiang Lei, Xisen Jin, Min-Yen Kan, Zhaochun Ren, Xiangnan He, and Dawei Yin. 2018. Sequicity: Simplifying task-oriented dialogue systems with single sequence-to-sequence architectures. In *ACL*. 1437–1447.
- [17] Qintong Li, Hongshen Chen, Zhaochun Ren, Zhumin Chen, Zhaopeng Tu, and Jun Ma. 2019. EmpGAN: Multi-resolution Interactive Empathetic Dialogue Generation. *arXiv preprint arXiv:1911.08698* (2019).
- [18] Zekang Li, Cheng Niu, Fandong Meng, Yang Feng, Qian Li, and Jie Zhou. 2019. Incremental Transformer with Deliberation Decoder for Document Grounded Conversations. In *ACL*. 12–21.
- [19] Rongzhong Lian, Min Xie, Fan Wang, Jinhua Peng, and Hua Wu. 2019. Learning to Select Knowledge for Response Generation in Dialog Systems. In *IJCAI*.
- [20] Chin-Yew Lin. 2004. Rouge: A Package for Automatic Evaluation of Summaries. In *Text Summarization Branches Out*. 74–81.
- [21] Shuman Liu, Hongshen Chen, Zhaochun Ren, Yang Feng, Qun Liu, and Dawei Yin. 2018. Knowledge Diffusion for Neural Dialogue Generation. In *ACL*. 1489–1498.
- [22] Zhibin Liu, Zheng-Yu Niu, Hua Wu, and Haifeng Wang. 2019. Knowledge Aware Conversation Generation with Reasoning on Augmented Graph. In *EMNLP*. 1782–1792.
- [23] Fuli Luo, Peng Li, Jie Zhou, Pengcheng Yang, Baobao Chang, Zhifang Sui, and Xu Sun. 2019. A Dual Reinforcement Learning Framework for Unsupervised Text Style Transfer. In *IJCAI*.
- [24] Chuan Meng, Pengjie Ren, Zhumin Chen, Christof Monz, Jun Ma, and Maarten de Rijke. 2020. RefNet: A Reference-aware Network for Background Based Conversation. In *AAAI*.
- [25] Nikita Moghe, Siddhartha Arora, Suman Banerjee, and Mitesh M Khapra. 2018. Towards Exploiting Background Knowledge for Building Conversation Systems. In *EMNLP*. 2322–2332.
- [26] Kyosuke Nishida, Itsumi Saito, Kosuke Nishida, Kazutoshi Shinoda, Atsushi Otsuka, Hisako Asano, and Junji Tomita. 2019. Multi-style Generative Reading Comprehension. In *ACL*. Florence, Italy, 2273–2284.
- [27] Kishore Papineni, Salim Roukos, Todd Ward, and Wei-Jing Zhu. 2002. BLEU: A Method for Automatic Evaluation of Machine Translation. In *ACL*. 311–318.
- [28] Prasanna Parthasarathi and Joelle Pineau. 2018. Extending Neural Generative Conversational Model using External Knowledge Sources. In *EMNLP*. 690–695.
- [29] Lianhui Qin, Michel Galley, Chris Brockett, and Xiaodon Liu. 2019. Conversing by Reading: Contentful Neural Conversation with On-demand Machine Reading. In *ACL*. 5427–5436.
- [30] Pranav Rajpurkar, Jian Zhang, Konstantin Lopyrev, and Percy Liang. 2016. SQuAD: 100,000+ Questions for Machine Comprehension of Text. In *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing (EMNLP)*. 2383–2392.
- [31] Siva Reddy, Danqi Chen, and Christopher D. Manning. 2019. CoQA: A Conversational Question Answering Challenge. *Transactions of the Association for Computational Linguistics* 7 (2019), 249–266.
- [32] Pengjie Ren, Zhumin Chen, Christof Monz, Jun Ma, and Maarten de Rijke. 2020. Thinking Globally, Acting Locally: Distantly Supervised Global-to-Local Knowledge Selection for Background Based Conversation. In *AAAI*.
- [33] Abigail See, Peter J Liu, and Christopher D Manning. 2017. Get To The Point: Summarization with Pointer-Generator Networks. In *ACL*. 1073–1083.
- [34] Lifeng Shang, Zhengdong Lu, and Hang Li. 2015. Neural Responding Machine for Short-text Conversation. In *ACL*. 1577–1586.
- [35] Shiv Shankar and Sunita Sarawagi. 2019. Posterior Attention Models for Sequence to Sequence Learning. In *ICLR*.
- [36] Sainbayar Sukhbaatar, Jason Weston, Rob Fergus, et al. 2015. End-to-end Memory Networks. In *NeurIPS*. 2440–2448.
- [37] Yibo Sun, Duyu Tang, Nan Duan, Shujie Liu, Zhao Yan, Ming Zhou, Yuanhua Lv, Wenpeng Yin, Xiaocheng Feng, Bing Qin, et al. 2019. Joint Learning of Question Answering and Question Generation. *TKDE* (Early access) (2019).
- [38] Ilya Sutskever, Oriol Vinyals, and Quoc V Le. 2014. Sequence to Sequence Learning with Neural Networks. In *NeurIPS*. 3104–3112.
- [39] Richard S Sutton, David A McAllester, Satinder P Singh, and Yishay Mansour. 2000. Policy Gradient Methods for Reinforcement Learning with Function Approximation. In *NeurIPS*. 1057–1063.
- [40] Duyu Tang, Nan Duan, Tao Qin, Zhao Yan, and Ming Zhou. 2017. Question Answering and Question Generation as Dual Tasks. *arXiv preprint arXiv:1706.02027* (2017).
- [41] Yi-Lin Tuan, Yun-Nung Chen, and Hung-yi Lee. 2019. DyKgChat: Benchmarking Dialogue Generation Grounding on Dynamic Knowledge Graphs. In *EMNLP*. 1855–1865.
- [42] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Lukasz Kaiser, and Illia Polosukhin. 2017. Attention Is All You Need. In *NeurIPS*. 5998–6008.
- [43] Wenjie Wang, Minlie Huang, Xin-Shun Xu, Fumin Shen, and Liqiang Nie. 2018. Chat More: Deepening and Widening the Chatting Topic via A Deep Model. In *SIGIR*. 255–264.
- [44] Yiren Wang, Yingce Xia, Tianyu He, Fei Tian, Tao Qin, ChengXiang Zhai, and Tie-Yan Liu. 2019. Multi-Agent Dual Learning. In *ICLR*.
- [45] Yijun Wang, Yingce Xia, Li Zhao, Jiang Bian, Tao Qin, Guiquan Liu, and Tie-Yan Liu. 2018. Dual Transfer Learning for Neural Machine Translation with Marginal Distribution Regularization. In *AAAI*.
- [46] Yingce Xia, Jiang Bian, Tao Qin, Nenghai Yu, and Tie-Yan Liu. 2017. Dual Inference for Machine Learning. In *IJCAI*. 3112–3118.
- [47] Yingce Xia, Tao Qin, Wei Chen, Jiang Bian, Nenghai Yu, and Tie-Yan Liu. 2017. Dual Supervised Learning. In *ICML*. 3789–3798.
- [48] Yingce Xia, Xu Tan, Fei Tian, Tao Qin, Nenghai Yu, and Tie-Yan Liu. 2018. Model-level Dual Learning. In *ICML*. 5383–5392.
- [49] Zili Yi, Hao Zhang, Ping Tan, and Minglun Gong. 2017. Dualgan: Unsupervised Dual Learning for Image-to-image Translation. In *ICCV*. 2849–2857.
- [50] Hainan Zhang, Yanyan Lan, Jiafeng Guo, Jun Xu, and Xueqi Cheng. 2018. Reinforcing Coherence for Sequence to Sequence Model in Dialogue Generation. In *IJCAI*. 4567–4573.
- [51] Wen Zheng and Ke Zhou. 2019. Enhancing Conversational Dialogue Models with Grounded Knowledge. In *CIKM*. 709–718.
- [52] Hao Zhou, Tom Young, Minlie Huang, Haizhou Zhao, Jingfang Xu, and Xiaoyan Zhu. 2018. Commonsense Knowledge Aware Conversation Generation with Graph Attention. In *IJCAI*. 4623–4629.
- [53] Kangyan Zhou, Shrimai Prabhumoye, and Alan W Black. 2018. A Dataset for Document Grounded Conversations. In *EMNLP*. 708–713.