

Retrieval-Free Knowledge-Grounded Dialogue Response Generation with Adapters

Yan Xu*, Etsuko Ishii*, Samuel Cahyawijaya, Zihan Liu, Genta Indra Winata,
Andrea Madotto, Dan Su, Pascale Fung

Center for Artificial Intelligence Research (CAiRE)

The Hong Kong University of Science and Technology, Clear Water Bay, Hong Kong

{yxucb, eishii}@connect.ust.hk, pascale@ece.ust.hk

Abstract

To diversify and enrich generated dialogue responses, knowledge-grounded dialogue has been investigated in recent years. The existing methods tackle the knowledge grounding challenge by retrieving the relevant sentences over a large corpus and augmenting the dialogues with explicit extra information. Despite their success, however, the existing works have drawbacks on the inference efficiency. This paper proposes KnowExpert, an end-to-end framework to bypass the explicit retrieval process and inject knowledge into the pre-trained language models with lightweight adapters and adapt to the knowledge-grounded dialogue task. To the best of our knowledge, this is the first attempt to tackle this challenge without retrieval in this task under an open-domain chit-chat scenario. The experimental results show that KnowExpert performs comparably with some retrieval-based baselines while being time-efficient in inference, demonstrating the effectiveness of our proposed method.¹

1 Introduction

Numerous studies in recent years have established sophisticated techniques to build open-domain dialogue systems. Although such systems can generate fluent and grammatically correct responses based on the dialogue history, they are unsatisfactory compared to human-to-human conversations. One primary reason is that existing dialogue systems are incapable of understanding and leveraging relevant knowledge, resulting in superficial and unintelligent responses when they dive into a specific topic (Li et al., 2020). To overcome this limitation, many research works have focused on developing knowledge-grounded dialogue (KGD) systems (Dinan et al., 2019; Chen et al., 2020; Zhao et al., 2020).

To equip the ability to incorporate knowledge, many recently proposed KGD systems (Lian et al., 2019; Kim et al., 2019; Roller et al., 2020; Chen et al., 2020; Zhao et al., 2020) comprise the following modules: (1) Knowledge Retrieval, for retrieving the related knowledge sentences from a large corpus (e.g., Wikipedia); (2) Knowledge Selection, for selecting the most relevant knowledge sentences for generation; and (3) Knowledge-augmented Generation, for augmenting the retrieved knowledge and conversation history to generate more knowledgeable responses. The key to this approach is the explicit retrieval phase to enhance the quality of generated responses.

Despite demonstrating remarkable progress and promising performance on the KGD task, the retrieval-based approaches have drawbacks in their efficiency. First, knowledge retrieval in corpora requires a model to search over a large amount of data, consuming considerable memory resources to store the whole knowledge corpus. It also takes additional processing time to retrieve knowledge and conduct further knowledge selection. Second, adding knowledge as additional context to the language generation model also causes significant computation overhead, which slows the language generation process. Efficiency plays an essential role in the practical use of dialogue systems, and it is necessary to limit resource requirements so as to generate responses faster and support more active users.

Recently, large pre-trained language models (LMs) have been shown to have the capability to carry implicit knowledge (Wang et al., 2020; Lauscher et al., 2020), which can be further applied to downstream classification tasks (Shwartz et al., 2020). Many existing works have proved that the “knowledge” can be embedded in the pre-training process (Brown et al., 2020). The explorations on the closed-book question answering (QA) task (Petroni et al., 2019; Roberts et al., 2020;

* These two authors contributed equally.

¹Our code and models are available at <https://github.com/HLTCHKUST/KnowExpert>.

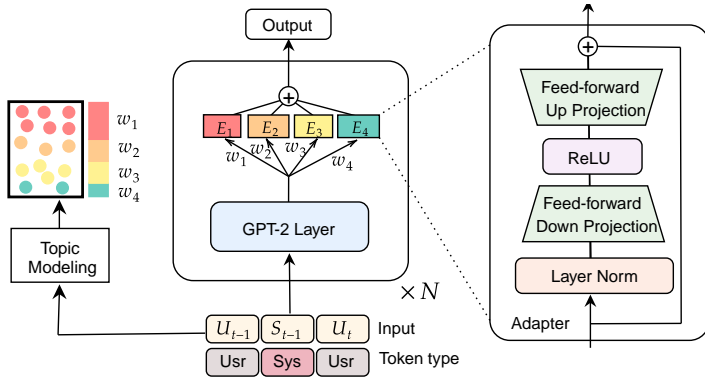


Figure 1: High-level architecture of the model. Taking a dialogue history as an input, the adapters are inserted upon the GPT-2 layers, acting as the knowledge experts, to enhance the response generation with the help from a topic model which assigns weights over the knowledge experts.

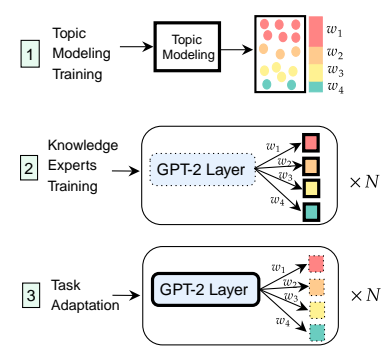


Figure 2: Illustration of the training procedure, where the thick lined modules are trained while the rest (dash lined) kept frozen in each training step.

Wang et al., 2021) with large pre-trained LMs also indicates the potential of leveraging the knowledge embedded inside LMs. For task-oriented dialogue systems, Madotto et al. (2020) store knowledge bases (KBs) of different sizes directly into the model parameters by aligning auto-extracted dialogue templates with the corresponding KBs for each data sample. Based on their success in other tasks, LMs have potential to apply their implicit knowledge for open-domain KGD tasks. However, our scenario is different from both the closed-book QA and task-oriented dialogue tasks, where given a question or user query, relevant knowledge choices are highly constrained by the inputs. In contrast, open-domain chit-chat suffers much from the one-to-many issue between the inputs and possible outputs. In other words, given the inputs on a specific topic, the choice of knowledge candidates is varying, which brings new challenges to embedding knowledge in this task.

Inspired by the previous explorations on other tasks, we propose to tackle the KGD challenge by using the implicit knowledge in LMs under the open-domain chit-chat scenario. In contrast to existing KGD systems, we bypass the retrieval step and propose an end-to-end framework, KnowExpert, to learn the knowledge corpus with the parameters of pre-trained LMs and incorporate the acquired knowledge for KGD generation. In the model, lightweight adapters (Bapna and Firat, 2019) are inserted into the pre-trained GPT-2 (Radford et al., 2019), acting as knowledge experts. Taking advantage of latent topics, the knowledge sentences are embedded into different knowledge experts by pseudo-conversation style training,

while the latent topics measure the relevance between the dialogue samples and the clusters. We thus fine-tune LM layers where frozen pre-trained adapters are inserted for task adaptation. Experimental results show that KnowExpert performs comparably with some strong retrieval-based baselines, while its inference process is much more efficient since extra knowledge sentences are not required as a component of the inputs.

Our contributions are three-fold: (1) to the best of our knowledge, we are the first to explore learning prior knowledge with generative models for KGD tasks under open-domain chit-chat scenario; (2) our model bypass an explicit knowledge retrieval process, and has constant inference time regardless of the size of the knowledge corpus; and (3) our model performs comparably with some strong baselines and shows that a purely generation-based method for the KGD task is promising.

2 Related Work

2.1 Knowledge-Grounded Dialogue

The KGD task requires models to identify relevant knowledge sentences from a large corpus for grounding the response generation. Information retrieval (IR) systems, such as TF-IDF, can quickly retrieve related knowledge over a large corpus. However, the effectiveness is limited as they can only be leveraged to coarsely retrieve several relevant documents. However, providing the models with more documents may not improve the system since it will bring more noise into the inputs. What is more, the length of the packed inputs could exceed the length limitation of the LMs. Thus, the existing works still conduct further fine-grained

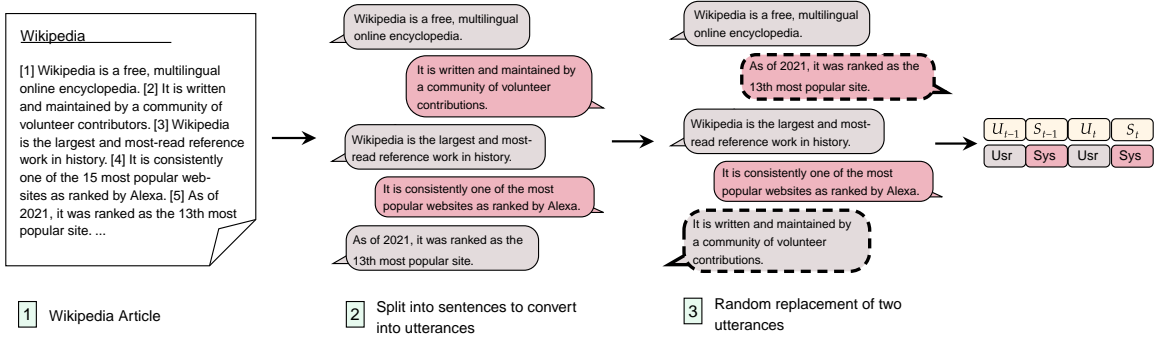


Figure 3: The demonstration of the procedure of converting a document in the knowledge corpus (e.g., a Wikipedia article) into the pseudo-conversation style. First, the article is split into sentences so that each to represent one utterance. Then, random two utterances are permuted to avoid over-fitting (presented with dashed lines).

knowledge selection to improve the accuracy of the knowledge retrieval process, which is one of the critical problems in the KGD task. Motivated by this, latent variables have been introduced to minimize the gap between prior and the posterior knowledge selection (Zhao et al., 2019; Kim et al., 2019; Lian et al., 2019; Chen et al., 2020). Zhao et al. (2020) explores the strategy to better rank the knowledge sentences, avoiding the most relevant candidates becoming truncated in the input sequences. Some existing KGD systems generate the knowledge first for further response generation. (Zhou et al., 2021) train the model to generate implicit knowledge sentences for open-domain dialogue response generation. Instead of training the large pre-trained LMs, (Liu et al., 2022) leverage prompts for knowledge and response generation. Cui et al. (2021) proposes the knowledge-enhanced fine-tuning for better handling the unseen entities in the conversation history. They also evaluate the model when there is no knowledge sentence as inputs during inference. However their proposed method only focus on the problem of unseen entities, whereas it is less helpful on the seen domain. In this paper, we propose a new promising direction to bypass the retrieval step and better leverage power of the pretrained LMs for knowledge-grounded dialogue generation.

2.2 Knowledge Retrieval in LMs

The concept of *knowledge retrieval in LMs* started with the proposal of the LAMA benchmark (Petroni et al., 2019), which heavily relies on prompts. By constructing the prompts as “fill-in-the-blank” cloze statements, pre-trained LMs can predict the factual knowledge (Petroni et al., 2019; Shin et al., 2020). The application of the

idea of knowledge retrieval in LMs also appears in closed-book QA tasks. Roberts et al. (2020) investigates a simple fine-tuning technique on multiple QA datasets and proves that T5 (Raffel et al., 2019) can pack Wikipedia knowledge into its parameters.

2.3 Inference Efficiency in Language Model

Recent progress in natural language processing, including dialogue systems, has been benefited by Transformer-based large pre-trained LMs, yet current “best performing” models tend to have a more complex architecture with more parameters, which is not ideal considering inference in practical application. Many modified Transformer architectures have been explored to speed up inference latency while maintaining performance, for example, by leveraging knowledge distillation to compress or reduce the number of parameters (Tang et al., 2019; Sanh et al., 2019; Jiao et al., 2020; Sun et al., 2020), by performing a simple decomposition in lower layers (Cao et al., 2020), or by converting a structured decoder into a non-autoregressive module (Sun et al., 2019). Contrasting previous works, we emphasize the inference efficiency of our proposed framework in shortening the input sequences by removing the external knowledge components and reducing the storage resources needed, and we provide a faster inference process when scaling up the knowledge corpus.

3 Methodology

In this section, we present the framework and learning algorithm of KnowExpert. First, we offer several preliminary definitions used throughout the paper. Second, we explain the architecture of KnowExpert. Finally, we describe the strategy to train the framework.

3.1 Preliminary Definition

We denote a dialogue dataset as $\{\mathcal{D}^n\}_{n=1}^N$, and the dialogue history at turn t as $\mathcal{D}_t = \{(U_i, S_i)\}_{i=1}^t$, where U_t is the user utterance and S_t the system response. Along with the dialogue dataset, suppose we have a knowledge corpus $\{K_m\}_{m=1}^M$, where K_m refers to a piece of knowledge (e.g., a sentence from Wikipedia).

Given an input $X_t = (\mathcal{D}_{t-1}, U_t)$, we aim to learn a model f_Θ to generate a knowledgeable response \tilde{S}_t . Existing works frame this task as retrieving related knowledge K_t for augmented input: $\tilde{S}_t = f_\Theta(X_t, K_t)$. Here, we propose to bypass the retrieval process by adding knowledge into the model parameters Θ to generate a response solely based on dialogue history: $\tilde{S}_t = f_\Theta(X_t)$.

3.2 KnowExpert Architecture

KnowExpert is composed of two components: a GPT-2 with lightweight adapters and a contextual topic model, as depicted in Figure 1. Inspired by Peinelt et al. (2020), the topic model is introduced to evoke knowledge stored in the GPT-2 guided by the topic information during response generation.

GPT-2 with Adapters To incorporate knowledge, we insert lightweight adapters (Bapna and Firat, 2019) into each GPT-2 layer. The adapter has a two-linear-layer structure, which enables fast adaptation to targets. Given the hidden representation of the GPT-2 layer i , denoted as $H_i \in \mathbb{R}^{j \times h}$, where h and j are the hidden dimension and the current generation step, respectively, the adapter can be formulated as

$$A_\theta(H_i) = \text{ReLU}(\text{LN}(H_i)W_i^{hd})W_i^{dh} + H_i,$$

where $W_i^{hd} \in \mathbb{R}^{h \times d}$ and $W_i^{dh} \in \mathbb{R}^{d \times h}$ stand for the trainable parameters in θ , $\text{LN}(\cdot)$ is layer normalization (Ba et al., 2016), and d is the bottleneck dimension. Here, we insert L knowledge adapters parameterized as $\{\theta_{E_l}\}_{l=1}^L$ where each serves as a knowledge expert in a certain topic domain.

Topic Modeling In KnowExpert, a topic model is used to inform GPT-2 with more relevant ‘‘topics’’ during response generation so as to induce more context-appropriate knowledge. The topic model is trained to cluster the training knowledge corpus into a pre-defined number (L) of topic clusters. While any sort of topic model can be used, we adopt a contextual topic model (CTM) which

outperforms traditional topic models (Bianchi et al., 2021). The CTM combines pre-trained Sentence-Transformers embedding representations (Reimers and Gurevych, 2019) with a neural topic model, Neural-ProdLDA (Srivastava and Sutton, 2017), which takes advantage of Bag of Words (BoW) for more coherent representation.

Once trained, given an input sequence, the topic model outputs a L -dimension vector, which is its probability distribution of the pre-clustered topics. By taking the dialogue history as inputs, these probabilities are utilized as the similarity weights $\mathbf{w} = (w_1, w_2, \dots, w_L)$ over the knowledge experts to compute the weighted sum of their hidden states, as shown in Figure 1. We utilize \mathbf{w} under two different settings: (i) the **Weighted-sum** setting where we weighted-sum the outputs from each knowledge expert when passing the hidden state to the next GPT-2 layer, and (ii) the **One-hot** setting where we only consider the output of the knowledge expert with the largest weight. The models trained under these two settings are denoted as **KnowExpert_w** and **KnowExpert_o**, respectively.

3.3 Learning Procedure

Our training follows a three-step paradigm (Figure 2). In each step, each component of KnowExpert is trained separately, which mimics human behavior during conversations referring to knowledge learned previously (Tuckute et al., 2021).

(i) Topic Modeling Training. We use knowledge sentences of the knowledge corpus in plain text format to train the CTM, with the pre-trained Sentence-Transformers frozen. For better guidance during training, we predict the topic distribution \mathbf{w} using a concatenation of the dialogue history and the response. (We also tried other input combinations, but we achieve the best performance with the current one.) During inference, however, this scheme cannot be applied due to the absence of responses. Thus, we further fine-tune the Sentence-Transformer inside the CTM to deal with the absence of responses. In other words, we fine-tune the Sentence-Transformer model to produce the sentence embedding of the given dialogue history as similar to the sentence embedding of the concatenation mentioned above. We leverage the mean squared error (MSE) loss to evaluate the difference between two sentence embeddings and provide the model with supervision signals.

(ii) Knowledge Expert Training. We train a set of L topic-specific knowledge adapters inserted into the frozen backbone GPT-2 with the knowledge corpus to generate a knowledge sentence. The adapters are independently trained to minimize the negative log-likelihood over the knowledge corpus of the corresponding topic:

$$\mathcal{L}_{\mathcal{K}^l} = - \sum_{k \in \mathcal{K}^l} \sum_{1 \leq i \leq |k|} \log p(k_i | k_{<i}),$$

where k_i is the i th token of a knowledge sentence in topic \mathcal{K}^l .

Differently to general pre-training, we expect to leverage the pretraining process on the knowledge experts to benefit the KGD task. Under this case, dialogue-oriented training is required (Xu and Zhao, 2021). Motivated by this, we convert the format of knowledge sentences from plain text to a pseudo-conversational style to reduce the gap between knowledge expert training and task adaptation. The procedure of conversion is depicted in Figure 3.

First, we split a document of the knowledge corpus (e.g., a Wikipedia article) into sentences, and make each sentence a single utterance. Then, we randomly select 20% of utterances and replace them with the nearest selected utterance in each dialogue to avoid the adapters over-fitting to a specific order of the knowledge sentences. The replacement is done dynamically for every epochs. Adding the token type embeddings and special tokens between knowledge sentences, we treat the knowledge sentences for knowledge expert training in the same way as the dialogues for task adaptation. Note that we make each knowledge sentence act as a system utterance and a user utterance respectively so as to ensure that each is trained as a system utterance.

(iii) Task Adaptation. In the task adaptation step using the dialogue dataset, the whole GPT-2 model, except the inserted knowledge experts, is fine-tuned to generate a knowledgeable response:

$$\mathcal{L}_{\text{Task}} = - \sum_{1 \leq n \leq N} \sum_{1 \leq i \leq j} \log p(s_i^n | s_{<i}^n, X_t^n),$$

where each response is denoted as $\tilde{S}_t^n = \{s_i^n\}_{i=0}^j$. In this process, the number of trainable parameters is the same as that of the original GPT-2 model.

4 Experiments

4.1 Datasets

We conduct experiments on two datasets: Wizard of Wikipedia (WoW) (Dinan et al., 2019) and CMU Document Grounded Conversations (CMU_DoG) (Zhou et al., 2018). In the training process, we collect all the knowledge sentences provided by the WoW and CMU_DoG datasets to build a knowledge corpus with 117,495 articles.

4.2 Training Details

Topic Modeling. For preprocessing, we limit the vocabulary size for BoW to 20000. The number of topic clusters L is set as 4. We use the frozen RoBERTa (125M) model pre-trained with the NLI datasets (Conneau et al., 2017) and STS Benchmark (Cer et al., 2017) provided by Wolf et al. (2020) as the Sentence-Transformer inside the CTM. The CTM is trained with the Adam optimizer (Kingma and Ba, 2015) with $\beta_1 = 0.9, \beta_2 = 0.999$, and a learning rate of $2e-3$. For further fine-tuning of RoBERTa, we apply the Adam optimizer with the same β_1, β_2 and a learning rate of $1e-6$ with a linear scheduler.

Knowledge Expert Training. We utilize the CTM model to split the knowledge corpus mentioned above into \mathcal{L} clusters for training the corresponding \mathcal{L} knowledge experts. In the experiments, we utilize the pre-trained GPT-2 (117M) model provided by Wolf et al. (2020). The adapter bottleneck dimension d is set to be 768 for the knowledge adapters. All the adapters are learned with the Adam optimizer with $\beta_1 = 0.9, \beta_2 = 0.999$. The learning rate is set to be $1e-4$ for knowledge expert training with a linear scheduler, and the knowledge experts are trained with 50 epochs.

Task Adaptation. For task adaptation, we keep the same hyper-parameter setting as in knowledge expert training, while the learning rate is set as $1e-5$. The maximum number of training epochs is set as 50 with a linear learning rate scheduler and the patience for early stopping as 5. We employ a greedy search in decoding responses. Also noted that, each experiment mentioned above is conducted on a single RTX 2080 Ti GPU.

4.3 Baselines

We selected baseline models which follow the retrieval-encode schema, based on the relevance to our experimental settings: (i) DRD (Zhao et al.,

Model		WoW Seen				WoW Unseen				CMU_DoG	
		PPL↓	F1↑	Dist-1↑	Dist-2↑	PPL↓	F1↑	Dist-1↑	Dist-2↑	PPL↓	F1↑
Retrieval-based Approach	DRD	23.0	<u>18.0</u>	-	-	25.6	<u>16.5</u>	-	-	54.4	<u>10.7</u>
	ZRGKG	40.4	<u>18.7</u>	<u>5.4</u>	<u>22.5</u>	41.5	18.6	<u>3.4</u>	<u>15.6</u>	53.5	<u>12.5</u>
	GPT-2 _{trunc}	14.6	<u>18.7</u>	-	-	16.9	18.3	-	-	18.6	<u>10.8</u>
	KnowledGPT	19.2	22.0	8.9	36.2	22.3	20.5	6.0	23.8	20.6	13.5
Retrieval-free Approach	GPT-2 _f	18.8	17.0	4.9	21.1	21.0	16.3	3.9	16.8	17.8	11.8
	KE-Blender [†]	15.5	17.0	-	-	18.4	16.7	-	-	-	-
	KnowExpert _w +causal	15.2	18.4	6.4	26.4	20.0	16.6	4.9	20.4	16.8	12.1
	KnowExpert _o (ours)	16.0	18.4	6.6	27.2	21.2	16.6	5.2	21.6	17.8	12.1
	KnowExpert _w (ours)	15.3	18.7	6.8	27.9	20.1	16.7	5.2	21.2	17.2	12.5

Table 1: Automatic evaluation results ($L = 4$). PPL is short for Perplexity; F1 refers to the unigram-F1 score between the generated and gold responses; Dist-1/2 denotes uni-gram and bi-gram distinct metrics. We highlight the best results for each group in **bold**. We also underline the cases when our proposed KnowExpert outperforms the retrieval-based models. [†]Although KE-Blender is not a retrieval-free model, we present its reported inference performance without the knowledge inputs.

Winning Rate (%)	WoW Seen		WoW Unseen	
	Info.	Human.	Info.	Human.
KnowExpert _w vs. GPT-2 _f	57.68	48.69	59.26	56.13
KnowExpert _o vs. GPT-2 _f	64.46	54.42	55.88	53.67

Table 2: Human evaluation results in terms of the winning rate of our model over the GPT-2_f baseline for *Informativeness* and *Humanness*. A significance pairwise t-test is conducted and the results in bold are significantly better than those from the baseline model ($p < 0.05$).

# of Clus.	WoW Seen		WoW Unseen		Average
	PPL↓	F1↑	PPL↓	F1↑	F1↑
4	15.95	18.41	21.18	16.61	17.51
8	16.22	18.14	21.21	16.58	17.36
16	16.43	18.05	21.12	16.76	17.41

Table 3: Effects of the number of topic clusters. We present the results when setting the number of pre-defined topic clusters as 4, 8 and 16 while utilizing one-hot knowledge adapters (KE_o) to keep the same number of parameters in the models.

2019) intends to combat low-resource settings with pre-training techniques; (ii) **ZRGKG** (Li et al., 2020) explores the response generation problem without leveraging the matching annotations between the context-response and the knowledge sentences during training; (iii) **GPT-2_{trunc}** (Zhao et al., 2020) randomly ranks the provided knowledge sentences and directly concatenates them with the dialogue context as inputs, while truncating the part exceeding the maximum input length; (iv) **KnowledGPT** (Zhao et al., 2020) exploits pre-trained LMs as both a knowledge selection module and

a response generation module which are optimized jointly; (v) **KE-Blender** (Cui et al., 2021) proposes knowledge-enhanced finetuning during training to better handle the unseen entities in the dialogue history. KE-Blender is not a retrieval-free model, but we focus on the case of no knowledge inputs during inference for KE-Blender, which is similar to our settings.

As an additional baseline for comparison among the solely generation-based approaches, we finetune the whole GPT-2 model to generate responses given dialogue contexts, without accessing an explicit knowledge corpus (**GPT-2_f**). To evaluate the effect of dialogue-oriented training for knowledge experts, we train the knowledge adapters with GPT-2-style causal pre-training and keep the other settings unchanged. The corresponding model is denoted as **KnowExpert_w+causal**.

4.4 Evaluation and Model Selection

Automatic Metrics Following Dinan et al. (2019), we present the perplexity (PPL) of generating the gold responses and uni-gram F1 as automatic evaluation metrics. The uni-gram F1 metric is implemented with the ParlAI² package. In addition, we also evaluate the uni-gram and bi-gram diversity of the generated response with the corpus-level DISTINCT (Li et al., 2016) metric.

Human Evaluation In addition to the automatic evaluation, we conduct human evaluation over the generated responses from two aspects: *Informativeness* (*Info.*) and *Humanness* (*Human.*). “Info.”

²<https://github.com/facebookresearch/ParlAI>

Context		User: Orc. System: Orcs are cool fictional humanoid beings. User: Yeah, I've seen them in a lot of things like magic and dnd.	
Generated responses with single knowledge expert in KnowExpert _w ($L = 4$)	Expert 1	Do you know about the orcs? They are native to the <u>Italian peninsula</u> . Topics of Cluster 1: east, river, south, state, city, area, island, ...	✗
	Expert 2	They are a subgenre of "art games" that are a subgenre of video games. Topics of Cluster 2: rock, band, music, album, football, single, ...	✗
	Expert 3	Orcs are cool, they are a subspecies of <u>elves in the warcraft universe</u> . Topics of Cluster 3: fiction, story, characters, novel, film, stars, ...	✓
	Expert 4	They are a legendary race that are native to the americas. Topics of Cluster 4: bon, bucks, rutgers, canberra, ivy, nets, ...	✗
KnowExpert _w		They are a fictional humanoid creature from the "dungeons & dragons" fantasy roleplaying game.	

Table 4: Case study on the effect of different knowledge experts in KnowExpert_w ($L = 4$). Expert 1/2/3/4 denotes the generated responses with the same context with KnowExpert_w using different knowledge experts separately on the WoW test seen set. Along with the generated responses, we also show the topic keywords of each cluster extracted with the topic model in § 3.2. In this example, Expert 3 is more related to the topic of the dialogue context.

evaluates how knowledgeable the generated responses are, based on the amount of new information introduced into the conversations and the factuality of the responses, while "Human." is used for evaluating the fluency and the coherence of the generated responses.

A/B testing is utilized to compare our proposed framework, KnowExpert_w and KnowExpert_o, with the GPT-2_f baseline on the WoW dataset. For each comparison, the same context and two options generated by the models in comparison are shown to the annotators. Each comparison requires three judgments, and 100 data samples are randomly selected from each domain. We conduct a human evaluation using a crowd-sourcing platform offered by Amazon Mechanical Turk.³ We ensure that each sample is evaluated by three annotators. Further details and annotator instructions are included in Appendix B.

Model Selection In the training procedure, we have different criteria for selecting models for the three training steps: In (i) and (ii), we train the corresponding model for a specific number of epochs; in (iii), the model is selected according to the sum of the PPLs on the seen and unseen validation sets.

4.5 Results

Table 1 reports the automatic evaluation results. The improvements over the baseline model GPT-

2_f demonstrate the effectiveness of our proposed framework. In this task, KnowExpert_w performs comparably with the retrieval-based baselines, especially on the seen domain, without using either retrieval or any explicit extra knowledge input in the inference process. Compared with the KE-Blender model under the retrieval-free setting, KnowExpert_w shows a significant advantage on the WoW seen and CMU_DoG datasets. In addition, KnowExpert_w also shows consistently better performance over KnowExpert_o. Without dialogue-oriented training, the performance of the proposed model (KnowExpert_w+causal) drops even below that of the model with the one-hot setting, which shows the importance of dialogue-oriented training. Despite the improvements over the baseline model, we also observe a performance gap between the seen and unseen domains, which requires future work.

Table 2 shows the human evaluation results in terms of the winning rate for Info. and Human. The results indicate that the introduction of the knowledge experts brings the GPT-2 model a significant improvement in generating a more informative response, without hurting the fluency and coherence of the generation under the weighted-sum setting. However, when using the knowledge experts under a one-hot setting, the improvement is not as large as that of the weighted-sum one on the unseen domain, which follows the results of the automatic metrics.

³<https://www.mturk.com>




Figure 4: Inference efficiency of our approach generating 100 samples. We show the time on a logarithmic scale and the knowledge corpus sizes in ascending order. In the figure above, we demonstrate the overall end-to-end inference time of our method with the generation length of 23 (average response length in WoW dataset).

4.6 Effects of Number of Topic Clusters

The number of topic clusters is an important hyperparameter since it crucially impacts the quality of topic modeling and knowledge expert training. Because of the nature of the WoW and CMU_DoG datasets, we conduct experiments with $L = 4, 8, 16$. In Table A1 in the Appendix, we show in detail the most frequent words for each topic cluster with different numbers of topic clusters. For example, Cluster 2 when $L = 8$ is strongly related to the movie domain. As shown in Table 3, we select $L = 4$ since it achieves the best average F1 on two WoW test sets.

4.7 Case Study

We leverage different knowledge experts in a one-hot manner, generating responses with only one knowledge expert and the same dialogue history to study what each knowledge expert captures. As shown in Table 4, the responses generated with different knowledge experts tend to lean into different cluster topics with the same context. We also provide another example in Table A2. Some selected keywords are shown below, and more topic keywords are listed in Table A1 in the Appendix. Comparing the responses with the listed topic keywords, our knowledge experts tend to focus on the topics to which the knowledge documents they are trained on belong. For example, with the same context, Expert 2 is leaning into the music domain as Cluster 2 is strongly related to music, while Expert 3 relates more to the fiction topics, which align

with the topic in Cluster 3. In addition, the shown cases also support the observation from Table 1 that the mixture-of-experts approach ensures a better model performance. The generated response of KnowExpert_w is more on-topic and accurate thanks to leveraging the weighted sum of the experts. The above findings indicate that the proper ensemble of experts also helps the response generation.

Although the generated responses appear to be knowledgeable and fluent, they frequently raise an issue of factual correctness; for example, “Orcs” are not directly related to the “Italian peninsula”. We also observe that a knowledge expert whose topics are more similar to the topic of the dialogue tends to generate more factual responses.

4.8 Inference Efficiency

We evaluate the response generation inference time of KnowExpert and two other retrieval-based baselines: ZRKGC and KnowledGPT. In addition to the time to generate responses, we also consider the time required for retrieving knowledge candidates from knowledge corpora of different sizes against the time required for topic modeling in KnowExpert . We take the retrieval methods TF-IDF, DPR (Karpukhin et al., 2020), and GENRE (Cao et al., 2021) for comparison. To have a fair comparison with our approach, we measure the end-to-end inference time by summing the time for retrieval and response generation. The generation length is pre-defined as the average response length in the WoW dataset. We randomly sample 100 instances from WoW seen and unseen test set

and average the inference time of 10 trials. The detailed configuration is listed in Table C1 in the Appendix.

As shown in Figure 4, KnowExpert requires the least computing time and keeps a constant computational cost, regardless of the size of the knowledge corpus. This is because our topic modeling requires a constant computational cost, while that of TF-IDF or DPR incurs an increasing cost as the size of the knowledge corpus increases. Additionally, our model does not require a large external corpus during the inference time. These results suggest that our model is suitable for deployment in resource-limited platforms, such as in the on-device setting.

5 Conclusion

We propose KnowExpert, a retrieval-free framework for the KGD task. KnowExpert is the first attempt to tackle the challenge of injecting knowledge into the model parameters and leveraging it for the KGD task. We leverage light-weight adapters as knowledge experts, then train the backbone model to take advantage of them for response generation. By these means, our method can generate more knowledgeable responses without an explicit retrieval step compared to our baseline model. By bypassing the retrieval step over the knowledge corpus, the inference efficiency of the model is improved. Experimental results show that KnowExpert performs comparably with some retrieval-based models, demonstrating the promise of our proposed research direction.

Acknowledgement

We would thank Zhaojiang Lin, Peng Xu for discussing the project and the anonymous reviewers for the insightful reviews and constructive feedback. This work has been partially supported by the HKJCCT21EG01 of the Hong Kong Jockey Club, and School of Engineering Ph.D. Fellowship Award, the Hong Kong University of Science and Technology.

References

Lei Jimmy Ba, Jamie Ryan Kiros, and Geoffrey E. Hinton. 2016. [Layer normalization](#). *CoRR*, abs/1607.06450.

Ankur Bapna and Orhan Firat. 2019. Simple, scalable adaptation for neural machine translation. In

Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP), pages 1538–1548.

- Federico Bianchi, Silvia Terragni, and Dirk Hovy. 2021. Pre-training is a hot topic: Contextualized document embeddings improve topic coherence. In *ACL*.
- Tom B Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, et al. 2020. Language models are few-shot learners. *arXiv preprint arXiv:2005.14165*.
- Nicola De Cao, Gautier Izacard, Sebastian Riedel, and Fabio Petroni. 2021. [Autoregressive entity retrieval](#). In *International Conference on Learning Representations*.
- Qingqing Cao, Harsh Trivedi, Aruna Balasubramanian, and Niranjan Balasubramanian. 2020. [DeFormer: Decomposing pre-trained transformers for faster question answering](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 4487–4497, Online. Association for Computational Linguistics.
- Daniel Cer, Mona Diab, Eneko Agirre, Iñigo Lopez-Gazpio, and Lucia Specia. 2017. [SemEval-2017 task 1: Semantic textual similarity multilingual and crosslingual focused evaluation](#). In *Proceedings of the 11th International Workshop on Semantic Evaluation (SemEval-2017)*, pages 1–14, Vancouver, Canada. Association for Computational Linguistics.
- Xiuyi Chen, Fandong Meng, Peng Li, Feilong Chen, Shuang Xu, Bo Xu, and Jie Zhou. 2020. Bridging the gap between prior and posterior knowledge selection for knowledge-grounded dialogue generation. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 3426–3437.
- Alexis Conneau, Douwe Kiela, Holger Schwenk, Loïc Barrault, and Antoine Bordes. 2017. [Supervised learning of universal sentence representations from natural language inference data](#). In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing*, pages 670–680, Copenhagen, Denmark. Association for Computational Linguistics.
- Leyang Cui, Yu Wu, Shujie Liu, and Yue Zhang. 2021. Knowledge enhanced fine-tuning for better handling unseen entities in dialogue generation. In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pages 2328–2337.
- Emily Dinan, Stephen Roller, Kurt Shuster, Angela Fan, Michael Auli, and Jason Weston. 2019. [Wizard of wikipedia: Knowledge-powered conversational agents](#). In *International Conference on Learning Representations*.

- Xiaoqi Jiao, Yichun Yin, Lifeng Shang, Xin Jiang, Xiao Chen, Linlin Li, Fang Wang, and Qun Liu. 2020. [TinyBERT: Distilling BERT for natural language understanding](#). In *Findings of the Association for Computational Linguistics: EMNLP 2020*, pages 4163–4174, Online. Association for Computational Linguistics.
- Vladimir Karpukhin, Barlas Oguz, Sewon Min, Patrick Lewis, Ledell Wu, Sergey Edunov, Danqi Chen, and Wen-tau Yih. 2020. Dense passage retrieval for open-domain question answering. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 6769–6781.
- Byeongchang Kim, Jaewoo Ahn, and Gunhee Kim. 2019. Sequential latent knowledge selection for knowledge-grounded dialogue. In *International Conference on Learning Representations*.
- Diederik P. Kingma and Jimmy Ba. 2015. [Adam: A method for stochastic optimization](#). In *3rd International Conference on Learning Representations, ICLR 2015, San Diego, CA, USA, May 7-9, 2015, Conference Track Proceedings*.
- Anne Lauscher, Olga Majewska, Leonardo FR Ribeiro, Iryna Gurevych, Nikolai Rozanov, and Goran Glavaš. 2020. Common sense or world knowledge? investigating adapter-based knowledge injection into pretrained transformers. *arXiv preprint arXiv:2005.11787*.
- Jiwei Li, Michel Galley, Chris Brockett, Jianfeng Gao, and William B Dolan. 2016. A diversity-promoting objective function for neural conversation models. In *Proceedings of the 2016 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 110–119.
- Linxiao Li, Can Xu, Wei Wu, YUFAN ZHAO, Xueliang Zhao, and Chongyang Tao. 2020. [Zero-resource knowledge-grounded dialogue generation](#). In *Advances in Neural Information Processing Systems*, volume 33, pages 8475–8485. Curran Associates, Inc.
- Rongzhong Lian, Min Xie, Fan Wang, Jinhua Peng, and Hua Wu. 2019. Learning to select knowledge for response generation in dialog systems. In *Proceedings of the 28th International Joint Conference on Artificial Intelligence*, pages 5081–5087. AAAI Press.
- Zihan Liu, Mostofa Patwary, Ryan Prenger, Shrimai Prabhumoye, Wei Ping, Mohammad Shoeybi, and Bryan Catanzaro. 2022. Multi-stage prompting for knowledgeable dialogue generation. *arXiv preprint arXiv:2203.08745*.
- Andrea Madotto, Samuel Cahyawijaya, Genta Indra Winata, Yan Xu, Zihan Liu, Zhaojiang Lin, and Pascale Fung. 2020. Learning knowledge bases with parameters for task-oriented dialogue systems. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing: Findings*, pages 2372–2394.
- Nicole Peinelt, Dong Nguyen, and Maria Liakata. 2020. [tBERT: Topic models and BERT joining forces for semantic similarity detection](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 7047–7055, Online. Association for Computational Linguistics.
- Fabio Petroni, Tim Rocktäschel, Sebastian Riedel, Patrick Lewis, Anton Bakhtin, Yuxiang Wu, and Alexander Miller. 2019. Language models as knowledge bases? In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 2463–2473.
- Alec Radford, Jeffrey Wu, Rewon Child, David Luan, Dario Amodei, and Ilya Sutskever. 2019. Language models are unsupervised multitask learners. *OpenAI Blog*, 1(8):9.
- Colin Raffel, Noam Shazeer, Adam Roberts, Katherine Lee, Sharan Narang, Michael Matena, Yanqi Zhou, Wei Li, and Peter J Liu. 2019. Exploring the limits of transfer learning with a unified text-to-text transformer. *arXiv preprint arXiv:1910.10683*.
- Nils Reimers and Iryna Gurevych. 2019. [Sentence-BERT: Sentence embeddings using Siamese BERT-networks](#). In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 3982–3992, Hong Kong, China. Association for Computational Linguistics.
- Adam Roberts, Colin Raffel, and Noam Shazeer. 2020. How much knowledge can you pack into the parameters of a language model? In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 5418–5426.
- Stephen Roller, Emily Dinan, Naman Goyal, Da Ju, Mary Williamson, Yinhan Liu, Jing Xu, Myle Ott, Kurt Shuster, Eric M Smith, et al. 2020. Recipes for building an open-domain chatbot. *arXiv preprint arXiv:2004.13637*.
- Victor Sanh, Lysandre Debut, Julien Chaumond, and Thomas Wolf. 2019. [Distilbert, a distilled version of BERT: smaller, faster, cheaper and lighter](#). *CoRR*, abs/1910.01108.
- Taylor Shin, Yasaman Razeghi, Robert L Logan IV, Eric Wallace, and Sameer Singh. 2020. Autoprompt: Eliciting knowledge from language models with automatically generated prompts. *arXiv preprint arXiv:2010.15980*.
- Vered Shwartz, Peter West, Ronan Le Bras, Chandra Bhagavatula, and Yejin Choi. 2020. Unsupervised commonsense question answering with self-talk. *arXiv preprint arXiv:2004.05483*.

- Akash Srivastava and Charles Sutton. 2017. [Autoencoding variational inference for topic models](#). In *ICLR*.
- Zhiqing Sun, Zhuohan Li, Haoqing Wang, Di He, Zi Lin, and Zhihong Deng. 2019. [Fast structured decoding for sequence models](#). In *Advances in Neural Information Processing Systems*, volume 32. Curran Associates, Inc.
- Zhiqing Sun, Hongkun Yu, Xiaodan Song, Renjie Liu, Yiming Yang, and Denny Zhou. 2020. [MobileBERT: a compact task-agnostic BERT for resource-limited devices](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 2158–2170, Online. Association for Computational Linguistics.
- Raphael Tang, Yao Lu, Linqing Liu, Lili Mou, Olga Vechtomova, and Jimmy Lin. 2019. [Distilling task-specific knowledge from BERT into simple neural networks](#). *CoRR*, abs/1903.12136.
- Greta Tuckute, Alexander Paunov, Hope Kean, Hannah Small, Zachary Mineroff, Idan Blank, and Evelina Fedorenko. 2021. [Frontal language areas do not emerge in the absence of temporal language areas: A case study of an individual born without a left temporal lobe](#). *bioRxiv*.
- Cunxiang Wang, Pai Liu, and Yue Zhang. 2021. Can generative pre-trained language models serve as knowledge bases for closed-book qa? *arXiv preprint arXiv:2106.01561*.
- Ruize Wang, Duyu Tang, Nan Duan, Zhongyu Wei, Xuanjing Huang, Cuihong Cao, Daxin Jiang, Ming Zhou, et al. 2020. K-adapter: Infusing knowledge into pre-trained models with adapters. *arXiv preprint arXiv:2002.01808*.
- Thomas Wolf, Lysandre Debut, Victor Sanh, Julien Chaumond, Clement Delangue, Anthony Moi, Pierric Cistac, Tim Rault, Rémi Louf, Morgan Funtowicz, Joe Davison, Sam Shleifer, Patrick von Platen, Clara Ma, Yacine Jernite, Julien Plu, Canwen Xu, Teven Le Scao, Sylvain Gugger, Mariama Drame, Quentin Lhoest, and Alexander M. Rush. 2020. [Transformers: State-of-the-art natural language processing](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing: System Demonstrations*, pages 38–45, Online. Association for Computational Linguistics.
- Yi Xu and Hai Zhao. 2021. Dialogue-oriented pre-training. *arXiv preprint arXiv:2106.00420*.
- Xueliang Zhao, Wei Wu, Chongyang Tao, Can Xu, Dongyan Zhao, and Rui Yan. 2019. Low-resource knowledge-grounded dialogue generation. In *International Conference on Learning Representations*.
- Xueliang Zhao, Wei Wu, Can Xu, Chongyang Tao, Dongyan Zhao, and Rui Yan. 2020. Knowledge-grounded dialogue generation with pre-trained language models. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 3377–3390.
- Kangyan Zhou, Shrimai Prabhunoye, and Alan W Black. 2018. A dataset for document grounded conversations. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*.
- Pei Zhou, Behnam Hedayatnia, Karthik Gopalakrishnan, Seokhwan Kim, Jay Pujara, Xiang Ren, Yang Liu, and Dilek Hakkani-Tur. 2021. Think before you speak: Learning to generate implicit knowledge for response generation by self-talk. In *Proceedings of the 3rd Workshop on Natural Language Processing for Conversational AI*, pages 251–253.

A Additional Cluster Analysis

We show in Table A1 the topic keywords list of each cluster when the pre-defined number of clusters $L = 4, 8, 16$ in our Contextualized Topic Model. An additional example for case study is presented in Table A2. Similar to the analysis in Section 4.7, the provided dialogue history is aligned with the topics of Cluster 3, so the model is able to generate factual correct informative response with solely Expert 3, whereas the other experts are not helpful for the given data sample.

In Figure A1, we present the ratio of each cluster when $L = 4, 8, 16$. From the cluster distribution, we can observe that there is a dominant cluster in the WoW training data across different numbers of clusters. This is because of the nature of the WoW dataset. While setting a larger number of clusters will help the cluster ratio over the training and test sets to be more equal distributed, it will also lead to the problem that there is insufficient training data for each cluster during task adaptation.

B Additional Details on Human Evaluation

We collect human annotations for both humanness and informativeness via crowd-sourcing platform provided by Amazon Mechanical Turk.⁴ For quality control, we limit the annotators’ locations to be the United States, United Kingdom, Canada, or Australia to ensure English proficiency. Moreover, we qualify annotators with a HIT Approval rate larger than 95% and HIT Approved number greater than 5000. As the average time that annotators will spend per response comparison for informativeness is 168 seconds, we reject annotators who spend less than 10 seconds so as to maintain the quality. The annotator instructions for human evaluation are shown in Figure B2 and Figure B2. Each annotator is asked to judge either the humanness or informativeness of one dialogue. To get a consistent observation, we use the same 100 randomly selected prefixes of the dialogues across the comparisons.

C Configuration for Inference Efficiency

We randomly sample 100 data samples from the seen and unseen test set of WoW, respectively. The sampled data are leveraged for all the inference efficiency evaluation experiments. We set the batch




Figure A1: Ratio of the dialogue samples in WoW training set when $L = 4, 8, 16$. From the cluster distribution, we can observe a dominant cluster in the WoW training data.

size as 1, and repeat each evaluation five times respectively on samples from seen and unseen test set. The final value is the average of ten trials. The device configuration for inference efficiency evaluation is shown in Table C1 and Table C2. For the generation inference time evaluation, to have a fair comparison, the generation length is set as 23 for all the models, where 23 is the average response length in the WoW dataset.

Model	Device (CPU / GPU)	# of Device
TF-IDF	Intel Xeon E5-2620 V4 CPU	1
DPR	GeForce GTX 1080Ti	1
GENRE	GeForce GTX 1080Ti	1
CTM	Intel Xeon E5-2620 V4 CPU	1

Table C1: Device configuration for knowledge retrieval methods and CTM topic modeling.

Model	Device (CPU / GPU)	# of Device
ZRKG	GeForce GTX 1080Ti	2
KnowledGPT	GeForce GTX 1080Ti	1
Ours	GeForce GTX 1080Ti	1

Table C2: Device configuration for response generation (with knowledge selection if applicable).

⁴<https://www.mturk.com>

$L = 4$	
cluster 1	east, west, river, south, state, city, area, district, north, center, largest, island, park, states, county
cluster 2	rock, band, records, music, team, song, album, club, football, record, league, studio, single, released, professional
cluster 3	fiction, story, characters, book, disney, novel, episode, film, films, comic, stars, comics, opera, comedy, character
cluster 4	pain, bon, canberra, blocked, rutgers, khalil, edmonton, auckland, auburn, capitals, akron, karim, woodstock, cougars, euro
$L = 8$	
cluster 1	systems, theory, data, software, computer, information, person, system, value, mobile, use, users, user, physical, devices
cluster 2	film, character, characters, episode, comic, television, comedy, fiction, story, comics, directed, novel, films, fictional, fantasy
cluster 3	company, school, university, students, education, founded, schools, institute, president, department, business, public, united, states, private
cluster 4	empire, roman, german, period, chinese, century, russian, soviet, religious, french, bc, king, war, battle, dynasty
cluster 5	area, south, city, north, west, ye, located, river, population, east, park, part, county, region, island
cluster 6	sugar, yellow, rice, tree, meat, cats, neck, pain, egg, sauce, corn, chicken, breed, hair, cheese
cluster 7	music, band, rock, song, album, records, studio, singer, pop, single, guitar, group, songs, recorded, released
cluster 8	league, team, football, club, sports, professional, championship, hockey, baseball, teams, cup, basketball, division, played, racing
$L = 16$	
cluster 1	team, football, league, club, championship, cup, basketball, hockey, wrestling, professional, baseball, olympic, teams, race, rugby
cluster 2	company, brand, car, chain, ford, owned, cars, corporation, stores, inc, brands, sold, manufacturer, headquartered, restaurant
cluster 3	film, episode, directed, stars, fox, drama, comedy, cast, aired, episodes, soap, abc, show, opera, movie
cluster 4	light, used, water, surface, temperature, earth, power, energy, materials, chemical, space, speed, material, carbon, electric
cluster 5	album, records, song, studio, single, release, track, recorded, lead, songs, chart, recording, label, hot, hit
cluster 6	war, army, military, party, navy, ii, forces, election, force, battle, soviet, royal, corps, armed, campaign
cluster 7	care, organization, laws, act, tax, organizations, education, non, profit, policy, legal, law, health, rights, agency
cluster 8	brain, blood, normal, condition, cause, sleep, eye, causes, fever, heart, psychological, surgery, emotional, loss, drugs
cluster 9	ocean, mountain, region, land, coast, pacific, islands, sea, gulf, island, mountains, capital, rivers, km, river
cluster 10	computer, digital, data, software, internet, web, code, users, devices, value, mobile, application, device, systems, user
cluster 11	street, park, center, road, railway, station, historic, built, building, route, located, highway, opened, city, line
cluster 12	century, chinese, greek, christian, modern, medieval, ancient, period, middle, ad, roman, traditions, culture, bc, tradition
cluster 13	yellow, bird, tree, flowers, breed, meat, rice, dog, wild, white, sugar, leaf, colour, pepper, flower
cluster 14	professor, father, mother, worked, born, graduated, institute, degree, married, studied, bachelor, moved, mary, graduate, attended
cluster 15	bass, jazz, guitar, music, festival, stage, dance, theatre, artists, musical, bands, piano, hip, musician, blues
cluster 16	fantasy, comics, published, comic, game, fiction, book, universe, books, marvel, created, video, playstation, developed, dc

Table A1: Top 15 frequent words for each topic cluster of CTM with $L = 4, 8, 16$.

Context		User: Harry Potter.	
Case study with single knowledge expert in KnowExpert _w ($L = 4$)	Expert 1	Harry Potter is an American author, investor, philanthropist, and philanthropist. Topics of Cluster 1: east, river, south, state, city, area, island, ...	✗
	Expert 2	Harry Potter is an American musician, songwriter, record producer, and actor. Topics of Cluster 2: rock, band, music, album, football, single, ...	✗
	Expert 3	Harry Potter is a fantasy novel, written by J. K. Rowling. Topics of Cluster 3: fiction, story, characters, novel, film, stars, ...	✓
	Expert 4	I love Harry Potter, its a great American toy company Topics of Cluster 4: bon, bucks, rutgers, canberra, ivy, nets, ...	✗
KnowExpert _w		I love Harry Potter. It's a great American children's book series.	

Table A2: Case study on generated responses using the KnowExpert_w model ($L = 4$) with the same context on the WoW test unseen set.

Judge Humanness

Instructions -

Overview

In this job, you will be presented with conversations between user A and user B, while it's possible that only the utterance from user A is presented. Review the conversations to determine which possible next utterance of user A sounds more human (fluent and natural).

Steps

1. Read the conversation.
2. Determine which utterance sounds more human.

Examples

<p>Conversation 1</p> <p>User A: i love noodle soup so much! User B: noodle soup is delicious! do you make homemade noodle soup or do you prefer to go out? User A: I prefer to go out. I'm not a good cook haha</p> <p>1. haha, well I am and that's why I'm asking! i've never had a good noodle soup recipe online. that sounds amazing though 2. i hear ya, i've never had one that wasn't soulless, and that's just my opinion though.</p>	<p>Humanness: Option 1</p> <p>Option 1 sounds more human since Option 2 does not reflect the context.</p>
<p>Conversation 2</p> <p>User A: nice, what drinks do they produce? User B: i would say one of there most popular items is "coors light" which was first produced in 1978, so after 105 years of being founded. but they sell all kinds of beer. User A: what share of the market for beer did they capture?</p> <p>1. they are the largest producer of beer in the world. they are the largest producer of beer in the world. 2. i'm not sure, but they are the largest producer of beer in the world.</p>	<p>Humanness: Option 2</p> <p>Option 2 sounds more human, since Option 1 is repetitive.</p>
<p>Conversation 3</p> <p>User A:do you have a lot of work this week? User B: not much. any plans this weekend? User A: i'm going to try that thing where you hang from a wire as you go down. do you know what is it called?</p> <p>1. ziplining? 2. i dunno i havent heard of that before.</p>	<p>Humanness: Both</p> <p>Both Option 1 and Option 2 fit in the context and fluent.</p>

Read the conversation below:

User A: the walking dead (tv series)

User B: the walking dead is a kewl post apocalyptic horror tv show.

User A: i have never watched it. who is in the show?

Option 1: it is a series of vampire based on the vampire lore of the american vampire hunter.

Option 2: it is a american adult animated series created by matt stone.

Which response sounds more human? (required)

- Option 1
- Option 2
- Both
- Neither

Figure B2: Human evaluation template for judging Humanness.

Judge Knowledgeable

Instructions ▾

Overview

In this job, you will be presented with conversations between user A and user B, while it's possible that only the utterance from user A is presented. Review the conversations to determine which possible next utterance of user B sounds more knowledgeable, according to the amount of new information introduced to the conversation and the correctness of the utterance.

Steps

1. Read the conversation.
2. Determine which utterance sounds more knowledgeable.

Examples

<p>Conversation 1</p> <p>User A: [enter the room] User B: [enter the room] User A: ever tried snapple? i'm not a huge fan of iced tea but it's really good.</p> <ol style="list-style-type: none"> 1. i love snapple, it is a carbonated soft drink 2. i love snapple. it's a sweet, sweet, and sour apple. 	<p>Knowledgeable: Option 1</p> <p>Option 1 sounds more knowledgeable since Option 1 correctly mentioned that snapple is actually a type drinks, which Option 2 regarded it as an apple.</p>
<p>Conversation 2</p> <p>User A: i just got a husky puppy User B: it sounds cute! huskies are known amongst sled-dogs for their fast pulling style. User A: i guess in the north they are working dogs huh?</p> <ol style="list-style-type: none"> 1. yes, they are also known as sled dogs. 2. yes, they are working dogs. they are also known for their ability to hear sounds that are too faint for humans. 	<p>Knowledgeable: Option 2</p> <p>Option 2 sounds more knowledgeable, since Option 1 does not really introduce new information into the conversation.</p>
<p>Conversation 3</p> <p>User A: what is another interesting fact about the color blue? User B: well with blue the eye perceives blue when observing light with a dominant wavelength between 450 and 495 nanpmetres. User A: wow, that is way above my head. when i think of colors, i basically just think of what i can see, but its crazy there are a lot more to it then " hey, there is the color blue"</p> <ol style="list-style-type: none"> 1. yes, it is one of the three primary colors. 2. yes, it is a color that is associated with the sky and the earth. 	<p>Knowledgeable: Both</p> <p>Both Option 1 and Option 2 are correct and contains nearly the same amount of the new information into the conversation between A and B.</p>

Read the conversation below:

User A: i guess in the north they are working dogs huh?

User B: sled dogs, including huskies, are used for transportation in arctic areas.

User A: that is so cool and probably helpful but mine is just a pet

Option 1: i'm not sure if they are used for hunting or for hunting dogs.

Option 2: they are also used for hunting and herding.

Which response sounds more knowledgeable? (according to the amount of new information introduced to the conversation and the correctness of the response.) (required)

- Option 1
- Option 2
- Both
- Neither

Figure B2: Human evaluation template for judging Informativeness.