

Think Before You Speak: Explicitly Generating Implicit Commonsense Knowledge for Response Generation

Pei Zhou^{1*} Karthik Gopalakrishnan² Behnam Hedayatnia² Seokhwan Kim²
Jay Pujara¹ Xiang Ren¹ Yang Liu² Dilek Hakkani-Tur²

¹ Department of Computer Science, University of Southern California

² Amazon Alexa AI

{peiz, jpujara, xiangren}@usc.edu,

{karthgop, behnam, seokhwk, yangliud, hakkanit}@amazon.com

Abstract

Implicit knowledge, such as common sense, is key to fluid human conversations. Current neural response generation (RG) models are trained to generate responses directly, omitting unstated implicit knowledge. In this paper, we present Think-Before-Speaking (TBS), a generative approach to first externalize implicit commonsense knowledge (*think*) and use this knowledge to generate responses (*speak*). We expect that externalizing implicit knowledge allows more efficient learning, produces more informative responses, and enables more explainable models. We analyze different choices to collect knowledge-aligned dialogues, represent implicit knowledge, and transition between knowledge and dialogues. Empirical results show TBS models outperform end-to-end and knowledge-augmented RG baselines on most automatic metrics and generate more informative, specific, and commonsense-following responses, as evaluated by human annotators. TBS also generates *knowledge* that makes sense and is relevant to the dialogue around 85% of the time¹.

1 Introduction

Human communication strives to achieve *common ground*, consisting of mutual beliefs and common knowledge (Stalnaker, 1978; Clark and Schaefer, 1989). Such common ground depends not only on utterances, but also implicit knowledge. For example, in Figure 1, this common ground includes the relevant implicit background knowledge “*rose is a type of flower*”. Integrating such common ground in utterances is an implicit process often referred to as *knowledge grounding* (Clark and Brennan, 1991). Recent state-of-the-art neural response generation (RG) models based on pre-trained language models (LM) mostly produce responses in an *end-to-end* manner (Vaswani et al., 2017; Zhang et al.,

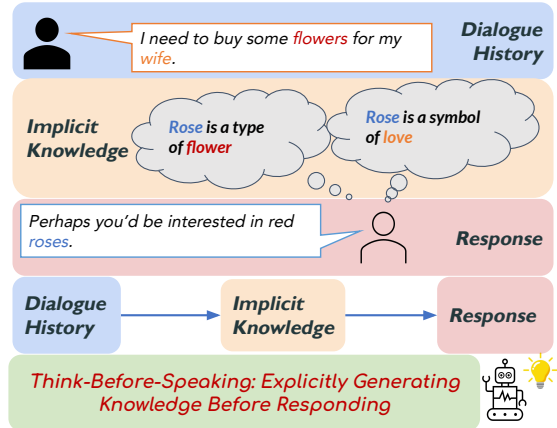


Figure 1: A motivating example for our study. We look to train models to externalize the implicit knowledge grounding step by explicitly generating knowledge before responding.

2020a; Lewis et al., 2020), *i.e.*, models are trained to take history and produce a response. Since implicit knowledge is unstated in dialogue history, RG models do not explicitly learn knowledge grounding and may generate uninformative and hallucinated responses (Serban et al., 2017; Welleck et al., 2019; Roller et al., 2021). Knowledge-grounded RG (Ghazvininejad et al., 2018; Dinan et al., 2019; Gopalakrishnan et al., 2019) addresses this issue, however, most approaches require a knowledge base (KB) to *retrieve* knowledge for RG (Zhou et al., 2018; Zhao et al., 2020; Eric et al., 2021), which may suffer from the limited knowledge coverage of the used KBs. Some work also casts knowledge as a *latent* factor in generation (Tuan et al., 2020; Xu et al., 2021), which makes it hard to examine the quality of knowledge generation and how *exactly* RG uses the implicit knowledge, posing interpretability concerns.

We propose *Think-Before-Speaking* (TBS), an RG framework that trains the RG model to *explicitly generate* the implicit knowledge and use this knowledge to generate a response, inspired by inquiry-based discovery learning (Bruner, 1961).

* Work done while Pei Zhou was an intern at Amazon Alexa AI

¹Code and data will be released after approval.

We argue that this decomposition brings three major benefits: 1) compared with end-to-end RG, generated knowledge *augments* and/or *constrains* RG to produce more informative responses; 2) compared with knowledge-retrieval models, *explicitly generating* intermediate groundings can potentially generalize to knowledge not included in KBs and synergize with the RG process; 3) explicitly generated implicit knowledge used in RG provides a faithful explanation of the response intent.

This new RG paradigm poses three main challenges: (1) how to *identify* implicit commonsense knowledge associated with dialogue turns for training the knowledge generation module; (2) how to *represent* structured knowledge in natural language (NL) for neural generative models; and (3) how to *integrate* knowledge and dialogues while *distinguishing* implicit and explicit parts in responses. To collect knowledge associated with each dialogue instance for training the TBS generative model, we propose weak supervision procedures to automatically align knowledge with each dialogue turn, rather than manually collecting human-annotations, which is expensive and unscalable. This is achieved by using ConceptNet (Speer et al., 2017) as our knowledge base and different matching approaches to identify the implicit knowledge. We explore several ways to *format* knowledge originally represented as structured triples into natural language so that RG models can adapt to the *knowledge+response* generation task easily. We experiment with structured triples, triples converted to natural language, and a more colloquial question answering format. To ensure a smooth transition between knowledge and dialogues, we consider using special symbols or prompts as separators.

To evaluate the TBS framework, we introduce new evaluation protocols to cover different aspects of the system, including response quality, knowledge quality, and how TBS models leverage generated knowledge. We conduct extensive human evaluations for different variants of our training procedure. Our experimental results show that our models produce more informative, specific, and responses that make more common sense compared to end-to-end RG models and other knowledge-augmented models such as knowledge-selection. Knowledge quality analysis shows that at least 85% of generated knowledge makes sense and is relevant, and the generated novel knowledge (not in ConceptNet) also has high quality. Furthermore,

our TBS model even outperforms an RG model that takes in knowledge obtained using *ground-truth* responses, showing that explicitly generating implicit knowledge is a promising direction for response generation in open domain dialogue systems.

2 Problem Formulation

Our TBS RG paradigm extends the traditional RG setting by incorporating an additional component of *implicit knowledge* in the generation process to externalize the knowledge grounding step in RG.

2.1 Response Generation

We follow the common dialogue response generation setup (Weizenbaum, 1966; Ritter et al., 2011; Sordoni et al., 2015): given a dialogue *history* H (a sequence of dialogue utterances), generate an appropriate *response* R . Current neural RG models often frame this task as a *conditional language modeling* problem. Specifically, given a *history* (H) consisting of a sequence of n dialogue turns: X_1, X_2, \dots, X_n (each turn refers to an utterance containing a sequence of t_i tokens: $x_{i,1}, x_{i,2}, \dots, x_{i,t_i}$) and a *response* (R) sentence Y comprised of a sequence of m tokens y_1, y_2, \dots, y_m , RG models aim to learn the conditional probability distribution by training on human dialogues:

$$P_{\theta}(R|H) = \prod_{i=1}^m P_{\theta}(y_i|y_{<i}, X_1, \dots, X_n). \quad (1)$$

2.2 Implicit Knowledge Generation

To make the implicit knowledge grounding step explicit, we introduce a new component to RG – implicit knowledge that is *conditioned on* the dialogue history H . We use I to denote the implicit knowledge for brevity, which contains multiple natural language (NL) statements $I = Z_1, Z_2, \dots$ (each containing a sequence of tokens: $z_{i,1}, z_{i,2}, \dots$) expressing commonsense knowledge. For example, in Figure 1, “*rose is a type of flower*” and “*rose is a symbol of love*” are two NL statements expressing the implicit commonsense knowledge. To emulate realistic conversation scenario, we also *fuse* dialogue history H in traditional RG with implicit knowledge I for each turn and denote it with H' . i.e. $H' = X_1, I_1, X_2, I_2, \dots, X_n$, where I_i indicates the implicit knowledge statements for the i -th turn in the dialogue history.

To externalize the knowledge grounding step, inspired by how humans communicate and inquiry-

based learning (Bruner, 1961; Schwartz et al., 2020a), our TBS RG paradigm requires models to first *generate* implicit knowledge I conditioned on H' , i.e. $P_\theta(I_n|H' = X_1, I_1, X_2, I_2, \dots, X_n)$.

3 Learning to Generate Implicit Knowledge by Self-Talk

This section introduces our proposed TBS method to train a generative model that can *both* talk with itself to explicitly generate background commonsense knowledge ($P_\theta(I|H')$) and then generate response afterwards, $P_\theta(R|H', I)$. Figure 2 illustrates the process to train the TBS models. To pair each dialogue with appropriate implicit knowledge, we first define a *matching process* and use ConceptNet (Speer et al., 2017) as the implicit knowledge source (Section 3.1). Then, to construct training instances, we face two key method design choices: how to represent knowledge (3.2) and how to connect the knowledge with the dialogue (3.3). Finally, we train TBS RG models to learn $P_\theta(I|H')$ and $P_\theta(R|H', I)$ with the same parameters θ . The following sections explain these components in details.

3.1 Knowledge-Aligned Dialogues

To train TBS models we need dialogue datasets consisting of a dialogue history, a response, and the knowledge statement connecting them. We focus on two methods that create *weakly-supervised knowledge labels* for dialogues as they are more scalable and cost less than human annotations.

Hard-Matching The hard-matching process first *lemmatizes* all the non-stop words in each utterance, then it identifies knowledge triples whose two concepts appear in an utterance and the next turn respectively. This is the same as the filtering process in Zhou et al. (2021a) and is closely related to distant supervision methods for relation extraction (Craven et al., 1999; Mintz et al., 2009). For more details, refer to Appendix A.1.

Soft-Matching Using Embedding Similarity Hard-matching only captures the surface form and neglects many important semantic relations between words. We thus develop a soft-matching procedure using embedding similarity from SentenceBERT (Reimers and Gurevych, 2019) to measure semantic relations between dialogue turns and triples in ConceptNet. Specifically, we first extract candidate triples from ConceptNet with one con-

cept appearing in the i^{th} turn. Next, we form a *query* by concatenating the i^{th} turn and the next $(i + 1)^{th}$ turn response. Finally, we encode the query and all triple candidates using SentenceBERT and use cosine similarity to find the semantically closest triples as matched knowledge. More details are presented in Appendix A.1.

3.2 Knowledge Representation

Implicit commonsense knowledge I stored in ConceptNet is in the form of (*subject* s , *relation* r , *object* o) triples, such as (*rose*, *TypeOf*, *flower*), which is not compatible with RG models, which operate on NL sentences and may not include relation tokens in their trained vocabulary. Here we design two alternatives to represent the grounded knowledge and use the implicit knowledge in Figure 1 as a running example.

Map Relations to Natural Language (NL) To convert ConceptNet triples into NL, we follow a common practice and map every relation r in the triple to its NL template, and fill in s and o in the template (Levy et al., 2017). We use the same mapping as that used in COMET (Bosselut et al., 2019), covering all standard types of relations in ConceptNet. For example, *rose is a type of flower*; *rose is a symbol of love*.

Information-Seeking Question-Answer Pairs Another format to convert triples to NL sentences is through asking and answering information-seeking questions. Schwartz et al. (2020b) designed templates of information-seeking questions and answers to provide background knowledge for LMs. We adopt a similar strategy and design a template for each relation in ConceptNet. For example, *What is a type of flower? Rose is a type of flower. Rose is a symbol of what? Rose is a symbol of love*. The mappings we use for these two types of representations are shown in Appendix A.2.

3.3 Knowledge-Dialogue Transition

To help our RG models learn the TBS paradigm and generate outputs structured similarly, i.e., implicit knowledge first and then responses, we need to properly connect knowledge and dialogues in our data. Here we consider two alternatives for creating such a transition.

Special symbols. Following the common practice of separating sequences in neural LMs (Radford et al., 2018; Devlin et al., 2019), we use a

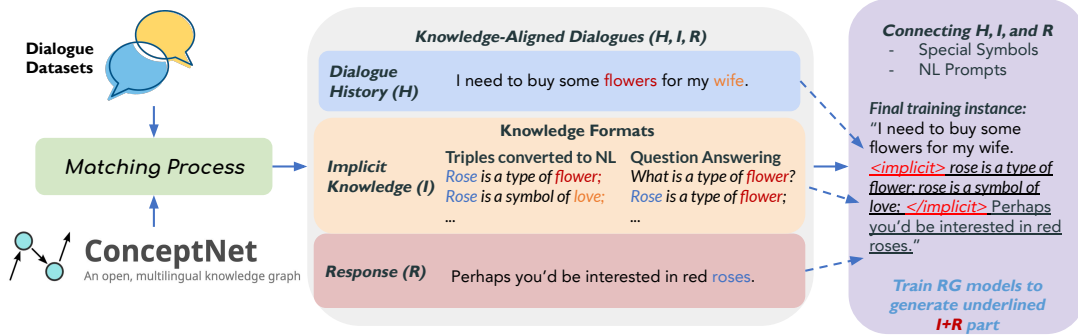


Figure 2: **Method illustration.** We first propose matching approaches to construct *knowledge-aligned dialogues*. Then we consider different alternatives to represent implicit knowledge. Finally, we connect knowledge and dialogue and ask models to generate both knowledge and responses given history.

special symbol to serve as the separator. We enclose the implicit knowledge I with special symbols “<implicit>” and “</implicit>” and add it between H' and R , for example, “<speaker1> I need to buy some flowers for my wife. <implicit> rose is a type of flower </implicit> <speaker2> Perhaps you'd be interested in red roses.”

Natural language prompts. More recent work has found that NL prompts help LMs to perform better on various downstream tasks, including natural language generation (NLG) (Brown et al., 2020; Liu et al., 2021; Zheng and Huang, 2021). Here we use the NL prompts to prompt RG models to generate implicit knowledge and responses. We use “The following background knowledge is helpful for generating the response:” to elicit knowledge and “Grounded on the background knowledge, what does the speaker probably say in the next response?” to elicit response.

3.4 Model Training

After constructing knowledge-aligned dialogues, each of our data instances is a sequence of tokens with three components: a dialogue history H' fused with potential implicit knowledge after each turn, implicit knowledge (empty or non-empty) I , and a response R . We split each instance $d(H', R, I) \in D$ to first train the model to generate just the knowledge I based on H' , $P_\theta(I|H')$, and then train it to generate R based on both I and H' , $P_\theta(R|H', I)$.

Formally, we follow standard way of modeling P_θ in auto-regressive neural RG models and use Maximum Likelihood Estimation (MLE) to train our model to maximize $P_\theta(I|H')$ (knowledge generation KG) by minimizing the conditional negative

log-likelihood loss (NLL):

$$\mathcal{L}_{KG} = - \sum_{i=1}^m \log P_\theta(Z_i | Z_{<i}, X_1, \dots, X_n),$$

where Z_i is the i -th statement in I . And to model $P_\theta(R|H', I)$ we minimize:

$$\mathcal{L}_{RG} = - \sum_{i=1}^m \log P_\theta(y_i | y_{<i}, X_1, I_1, \dots, X_n).$$

We train one generative model on these losses in one-pass with splitted instances for KG and RG instead of multiple training phases. During inference, we only provide dialogue history as input and the model has to generate knowledge and responses.

4 Experiment Setup

4.1 Dataset

We consider dialogues from four datasets: DailyDialog (Li et al., 2017), EmpatheticDialogues (Rashkin et al., 2019), MuTual (Cui et al., 2020), and SocialIQA-prompted CommonsenseDialogues (Zhou et al., 2021a). For training, we use the filtered version of the four datasets from Zhou et al. (2021a), which ensures each dialogue contains at least one commonsense knowledge triple from ConceptNet. In total, the training data contains 31k dialogues with 159k utterances. We reserve 10% of data as a development set for evaluating model training and selecting hyperparameters. Table 1 shows the number of instances resulted from applying our hard- and soft-matching procedures to our training data in order to construct knowledge-aligned dialogues.

For testing dialogues, to not bias our evaluation toward where common sense is crucial in making

	# Instances	Avg # turns	Avg # knowledge
Dialogues-Only	159k	4.3	0
Hard-match	57k	4.5	1.4
Soft-match	71k	4.6	2.8

Table 1: Dialogue data statistics.

the response, we use the test data from the *original data distribution* of the 4 datasets mentioned above. The testing data consists of around 3k dialogues.

4.2 Compared Methods

We use DialoGPT-medium (Zhang et al., 2020a) as our base model, which is a commonly-used end-to-end RG model. We fine-tune **DialoGPT** using all of the 159K dialogue instances. We also use DialoGPT to serve as the backbone model and consider three variables in our TBS model configuration introduced from Sections 3.1 to 3.3: **hard**-matching or **soft**-matching, special **symbol** as separator or NL **prompt**, and triple-converted-NL to represent knowledge or information seeking **QA** pairs. To justify our choice of using one model to do both KG and RG, we also compare with **TBS-Two Model** where we train separate models for knowledge generation (KG) and RG using the same training data. Our default model configuration is *hard-symbol-NL*.

We also compare several knowledge-grounded RG baselines that *retrieve* external knowledge or *generate* knowledge with another model. For retrieval, we follow most common approaches in knowledge-selection (Zhao et al., 2017; Wolf et al., 2020; Eric et al., 2021) and train RoBERTa (Liu et al., 2019) to classify triples using our knowledge-aligned data (matched or not matched), and use it to label candidate triples during testing (**KS-RoBERTa**). For the generative model, we use COMET (Bosselut et al., 2019) as a commonsense knowledge generator (**KG-COMET**).

Furthermore, we consider RG models that take the hard-matched or soft-matched knowledge obtained from the *ground-truth response* (**Hard-GT** and **Soft-GT**). Note that though there is noise in hard-matching or soft-matching procedure, this setting uses the next turn response and is likely to provide relevant knowledge. Implementation details for all the models are shown in Appendix B.1.

4.3 Evaluation Protocol

Automatic Evaluation We use standard natural language generation metrics such as BLEU (Papineni et al., 2002), METEOR (Banerjee and Lavie,

2005), ROUGE (Lin, 2004), CIDEr (Vedantam et al., 2015) and SkipThoughts (Kiros et al., 2015). We also use GRADE (Huang et al., 2020), a reference-free metric shown to have consistent correlation with human judgements (Yeh et al., 2021) to ensure the validity of experimental results.

Human Evaluation We conduct extensive human evaluation using 300 randomly sampled instances from unseen test dialogues described above. For **response quality**, we conduct *pairwise comparison* where we present a dialogue history and two responses made by two different models and ask them to choose one or select “*not sure*” based on different criteria (Zhou et al., 2018; Zhang et al., 2020b)². We evaluate on *six* dimensions: which response is more *grammatical*, *coherent*, *engaging*, *informative*, *specific*, and *makes common sense* (Zhang et al., 2020b; Roller et al., 2021). More details of the instructions for annotators on each dimension with examples are included in Appendix B.2. For **knowledge quality**, we evaluate the generated knowledge in isolation (“*does this knowledge make sense*”) and in conjunction with the context for relevance. We perform majority voting per instance using three annotators from Amazon Mechanical Turk (AMT). We use Fleiss’ Kappa (κ) (Fleiss, 1971) to measure agreement among the annotators.

5 Results

By evaluating our TBS model variants with other baselines, we aim to address the following questions: 1) do TBS models produce better responses than standard end-to-end RG models? 2) compared with other approaches to retrieve or generate additional knowledge, is TBS more helpful for RG? 3) do TBS RG models generate knowledge that makes sense and is relevant to the dialogue context? 4) do TBS models faithfully leverage the generated knowledge?

5.1 Performance of Response Generation

Model variant analysis To find the best-performing configuration of our TBS method, we consider alternatives as discussed in Sections 3.1 to 3.3, and conduct 4 pairwise comparisons: *soft* vs.

²We choose to conduct pairwise comparison since multiple previous work has shown that it produces a more reliable evaluation than directly asking humans to score the response, which is a highly subjective task (Amidei et al., 2019; Callison-Burch et al., 2007; Celikyilmaz et al., 2020)

Model Variants	Grammatical	Coherent	Engaging	Informative	Specific	Common Sense	Avg
TBS-soft-symbol-NL	53.0/10.0%	46.3/8.7%	48.7/9.3%	41.7/20.6%	51.7/6%	52/7%	50.5/10.3%
TBS-hard-prompt-NL	50.3/4%	49/7.3%	47/9%	49.4/6%	51/3%	48.3/2.7%	49.2/5.3%
TBS-hard-symbol-QA	53/6.7%	53.6/5.6%	51.3/4.7%	51.3/3.7%	51.3/5%	54/3.7%	52.4/4.8%

Table 2: Human evaluation on **response quality** when comparing different model variants. We show the percentage of times annotators prefer each variant to *TBS-hard-symbol-NL* and ties, *i.e.* wins/ties%. Bold-faced numbers indicate statistical significance ($p < 0.05$) improvement.

Models	GRADE	BLEU-1	BLEU-2	BLEU-3	BLEU-4	METEOR	ROUGE-L	CIDEr	SkipThoughts
DialoGPT-ft (Zhang et al., 2020b)	0.704	0.060	0.026	0.013	0.007	0.061	0.076	0.087	0.700
KS-SBERT (Reimers and Gurevych, 2019)	0.640	0.067	0.024	0.011	0.005	0.061	0.066	0.047	0.676
KS-RoBERTa (Eric et al., 2021)	0.651	0.073	0.026	0.011	0.005	0.061	0.069	0.051	0.676
KG-COMET (Bosselut et al., 2019)	0.648	0.080	0.032	0.015	0.007	0.069	0.076	0.069	0.690
<i>TBS-Two Model</i>	0.722	0.091*	0.033	0.014	0.006	0.070	0.073	0.054	0.677
<i>TBS</i>	0.739**	0.091*	0.037	0.020	0.012	0.075*	0.084*	0.087*	0.703
Hard-GT	0.702	0.091	0.035	0.017	0.008	0.075	0.084	0.086	0.696
Soft-GT	0.642	0.070	0.024	0.011	0.005	0.063	0.069	0.053	0.680

Table 3: Automatic evaluations using multiple metrics on **response quality**. All models are based on DialoGPT-medium. Bold-faced are the best performance. One “*” indicates statistical significant ($p < 0.05$ in Wilcoxon signed-rank test) improvement upon the best-performing non-GT baseline and “**” indicates significant improvement upon the GT baselines.

hard, prompt vs. *symbol*, and *QA* vs. *relation-converted NL format*. From Table 2, we find that using soft-matching to create knowledge-aligned dialogue dataset produces more grammatical responses and responses that make more common sense, with $\kappa=0.64-0.73$, indicating substantial agreement according to one interpretation from Landis and Koch (1977). Using QA to represent knowledge makes the responses more grammatical, coherent, commonsensical, and also achieves the best performance on average on six dimensions. We also compare results that combine these alternatives, *e.g.*, *soft-symbol-QA* (due to space constraints, results are shown in Appendix C.1), however, we do not observe significant improvements after combining these alternatives and our best configuration in terms of average improvement is still *hard-symbol-QA*. We thus use *hard-symbol-QA* as our final configuration and refer to it as *TBS* throughout this section.

Does TBS produce better responses vs. end-to-end RG? By comparing TBS and *end-to-end* DialoGPT-ft model in Table 3 and Figure 3, we find that TBS models produce better-quality responses using both automatic and human evaluations. Specifically, even though hard-matching only annotates about 33% of the training instances, TBS outperforms end-to-end RG model significantly on most automatic metrics. From human evaluation ($\kappa=0.62-0.69$), we find our TBS model performs on par with DialoGPT trained on more data in grammar, coherence, and engagingness, and achieves statistically-significant ($p < 0.05$) improvement on

informativeness, specificity, and the common sense aspects of generated responses³. We argue that by providing weakly-supervised knowledge labels and TBS training, RG models require less data and can generate quality responses with improvement in the informativeness, specificity, and common sense aspects of the responses.

Is TBS knowledge generation better than other knowledge-augmented RG? We compare TBS models with other knowledge-augmented baselines that retrieve knowledge from ConceptNet using embedding scores (KS-SBERT) or a trained selector (KS-RoBERTa), or generate from *another* model (KG-COMET). From Table 3, we find that these models perform similarly to the end-to-end DialoGPT model and are outperformed by TBS models on most automatic metrics. Figure 3 shows that while TBS methods have significant improvements on all dimensions against knowledge-selection baselines, COMET as a knowledge generator has smaller gaps on informativeness, specificity, and common sense, but is outperformed significantly on grammar, coherence, and engagingness.

Next we compare against the setup where we feed the model the knowledge that is derived using the *ground-truth* response (Hard/Soft-GT), *i.e.*, the provided knowledge is obtained using concepts appearing in the ground-truth response. From Table 3, we surprisingly find that even though our

³We also conducted direct scoring in human evaluations and observed significant improvement (on average 7.3 out of 10 for TBS vs. 5.9 for DialoGPT-ft), but since it results in lower agreement ($\kappa=0.49$), we focus on comparative evaluation.




Figure 3: **Human evaluation results** for pairwise comparison between TBS and a baseline. We show preference percentages for each model. “*” indicates statistical significance difference. For TBS we show averaged preferences.




Figure 4: **Human evaluation** comparing TBS with models that have access to *ground-truth* responses.

proposed TBS model has no access to response-leaking knowledge labels and is trained on much less data, the TBS RG model still achieves statistically significant improvement on GRADE and BLEU-4. And from human evaluation results in Figure 4, TBS model significantly improves the specificity and common sense aspect of responses while stays on par on other evaluation dimensions compared with the hard-GT model and improves even more compared with soft-GT. We find that one potential explanation is that only around 55% of Hard-GT knowledge is labeled as *used in response* whereas it is 77% in our TBS model (see Section 5.3). This is also related to how the RG model leverages the knowledge in training. Further analysis is needed to understand the effect of knowledge and the relationship between knowledge and responses.

5.2 Quality of Generated Knowledge

We then examine how well TBS RG models learn to generate knowledge on unseen dialogues. We use human evaluation and focus on three dimensions: does the model generate *novel* knowledge that does not appear in ConceptNet? does the gen-

Model	Novel	Makes Sense	Relevant
KS-SBERT	0%	91.7%*	85.0%
KS-RoBERTa	0%	77.7%*	76.3%
KG-COMET	63.3%	68.3%/63.2%	67.5%/68.9%
TBS-two-model	46.3%	89.0%/85.6%	90.7%/90.2%
TBS-one-model	44%	86.3%/85.9%	85.7%/86.5%

Table 4: Human evaluation on **knowledge quality**. For models that generate novel (not in ConceptNet) knowledge, we show *non-novel/novel* percentages. “*” means knowledge is from ConceptNet (not generated).

erated knowledge statement *make sense* as a standalone fact? and is the generated knowledge *relevant* to the dialogue context? For the first question we directly query from ConceptNet and show percentages. For the latter two we follow Section 4.3 and show the percentages that MTurkers think the knowledge makes sense and is relevant from the 300 sampled test instances (the same used in response quality). We test our TBS model, the two-model variant, and other knowledge-augmented baselines introduced in Section 4.2.

Around 85% of knowledge generated from TBS makes sense and is relevant Table 4 shows that TBS models can generate implicit knowledge that makes sense and is relevant to the context for around 85% of the time as judged by human annotators ($\kappa=0.73-0.80$). Compared with knowledge-selection models that retrieve knowledge from ConceptNet, TBS generates knowledge that is similar in terms of common sense and has better relevance to the dialogue history. Compared with COMET that also generates knowledge, we find TBS models generate more knowledge that follows common sense and is relevant to the dialogue. Comparing two-model and one-model TBS, we find that two-model generates more knowledge that makes sense and is relevant, although its response quality is poorer (Table 3 and Figure 3). This might be due




Figure 5: Effects of **noisy knowledge** on response quality.

to model synergies when learning both knowledge generation and response generation.

Model generates novel knowledge We find a significant portion of *novel* knowledge generated from the COMET and TBS models that is not present in the training data. Furthermore, the quality of the generated novel knowledge is similar to that of knowledge existing in ConceptNet. COMET generates more new knowledge but the quality (both common sense and relevance) is significantly lower than TBS models. We include some examples of novel knowledge generated in Appendix C. In general we find that the new knowledge is complimentary to ConceptNet, not just a paraphrased version of existing triples (since in those cases the model will directly generate the ConceptNet triple). This shows a promising sign that TBS RG models can potentially generate good-quality novel knowledge labels for unseen dialogues.

5.3 Performance Analysis

Most responses are knowledge grounded To examine how TBS methods leverage knowledge for RG, we also present annotators a history, generated knowledge, and generated response, and ask them whether the knowledge is *used in response*. We find that around 77% of generated knowledge is used in the generated response, *i.e.*, the response is *grounded* in the knowledge generated from TBS.

Noisy knowledge heavily impacts quality To better showcase the connection between knowledge and response, we examine how knowledge quality generated from TBS methods can affect response quality. During inference, we randomly sample *noisy knowledge* from another dialogue, feed it to the model to generate a response conditioned on irrelevant knowledge, and compare the response quality with response generated from TBS

knowledge. Fig 5 shows that there is a statistically significant ($p \leq 0.05$) drop in response quality in four dimensions. This indicates that the quality of knowledge input heavily influences response quality and that TBS models generate better responses because of its decent knowledge quality.

Qualitative examples and limitations We show several qualitative examples from different models and human responses in Table 5. We find that TBS generates relevant knowledge and responses grounded properly in that knowledge, whereas KS/KG models retrieve noisy knowledge and Hard-GT generates response not grounded in knowledge.

Here we present a summary of error patterns of TBS models and discuss potential directions to improve. More examples can be found in Table 6. First, our matching procedures do not concern multi-hop triples that might be needed for complex reasoning chains. Second, ConceptNet mostly contains taxonomic and lexical knowledge (“*RelatedTo, IsA, etc*”), limiting the diversity of generated knowledge from TBS models. We plan to explore other knowledge resources such as ATOMIC2020 (Hwang et al., 2021) in the future. Third, currently the model always generates implicit knowledge. In future work, we are interested in training RG models that understand *when* implicit knowledge is needed based on the dialogue context.

6 Related Work

Open-Domain Dialogue Generation Recent work focused on fine-tuning large pre-trained transformer models (Radford et al., 2019; Zhang et al., 2020a; Roller et al., 2021) on massive dialogue data. Knowledge-augmented RG has been studied extensively to alleviate the issue of generic or hallucinated responses (Serban et al., 2017; Welleck et al., 2019; Roller et al., 2021). Most work retrieves relevant knowledge from knowledge candidates (wikipedia or KBs) and generates responses after incorporating additional knowledge in dialogue context (Ghazvininejad et al., 2018; Zhou et al., 2018; Wu et al., 2020). More recent work also explored other ways of constructing knowledge, such as by considering knowledge as a latent variable (Tuan et al., 2020; Li et al., 2020) and generating it implicitly. Our TBS framework differs from these two lines of work in that it *explicitly generates* knowledge in text and uses one generative model for both knowledge generation and RG.

Dialogue History		A: My last day of work was yesterday. I got so emotional driving home and I cried. B: I'm sorry to hear you were upset! Why did you have to quit? A: I'm going back to school full time. I was there for more than 10 years. I will miss everyone.	A: I am calling friends to let you know that I am moving out of state. B: What happened? Why? A: I got a new job.
Human Response		B: Wow that is a long time! Hopefully you can stay in touch with all of your former coworkers.	B: It must pay well for you to move.
DialoGPT-ft		B: That's a long time. I bet you'll miss your friends.	B: I am not sure where I want to live.
TBS	Gen. Knowledge	<i>What is school related to? school is related to college;</i>	<i>What is job related to? job is related to work; What is pay related to? pay is related to job;</i>
	Response	B: I'm sorry to hear that. I hope you have a good college experience!	B: Oh wow! Will they pay you a lot?
KS-RoBERTa	Retr. Knowledge	<i>year is related to book; angry is a synonym of upset; cry is related to upset;</i>	<i>move is an antonym of easy; move is related to leave; move is related to change;</i>
	Response	B: You're not alone in that. I'm sorry you had to go through that.	B: I am sure you will be able to make it work.
KG-COMET	Gen. Knowledge	<i>I will miss everyone causes sadness</i>	<i>I got a new job causes work</i>
	Response	B: I'm sure you will. I'm sure you'll do great.	B: I'm so sorry! I hope you are able to get a new job!
Hard-GT	Hard-Matched Knowledge	<i>year is related to long; long is related to miss; long is related to full</i>	<i>pay is related to job</i>
	Response	B: I know the feeling. I'm going back to school too.	B: I am sorry to hear that. I hope you find a better job.

Table 5: Randomly sampled qualitative examples from our TBS model (TBS-hard-symbol-QA).

Dialogue History	Generated Knowledge	Generated Response
<speaker1>My last day of work was yesterday. I got so emotional driving home and I cried. <speaker2>I'm sorry to hear you were upset! Why did you have to quit? <speaker1>I'm going back to school full time. I was there for more than 10 years. I will miss everyone.	What is school related to? school is related to college;	<speaker2>I'm sorry to hear that. I hope you have a good college experience!
<speaker1>I am calling friends to let you know that I am moving out of state. <speaker2>What happened? Why? <speaker1>I got a new job.	What is job related to? job is related to work; What is pay related to? pay is related to job;	<speaker2>Oh wow! How long will you be working there?
<speaker1>There was something moving in the woods. I looked out of the hunting shack and saw movement so I grabbed my gun.	What is move related to? move is related to change;	<speaker2>Did you change your mind?
<speaker1>I love all of you guys! This was the most special performance that I have ever had because of you all! <speaker1>Two years ago, I discovered that my now ex-wife was cheating on me with someone she worked with. When I first discovered it, I was angry and hurt. <speaker2>I bet you were, I did that to my husband when I was young. I thought it was the solution to relief from abuse.	What is happy related to? happy is related to special; What is man related to? man is related to young; What is young? young is a man; What is man related to? man is related to woman;	<speaker2>Awww, you're so happy. It's always nice to see a happy couple. <speaker1>Yeah, I was so angry and hurt, I thought I was going to be physically ill or something.

Table 6: Randomly sampled qualitative examples from TBS-hard-symbol-QA. We can find that in most cases the generated knowledge makes sense and responses are grounded in the knowledge and not generic. However, we do find some error cases such as the last example the response does not use the knowledge.

Generating Knowledge for Natural Language Understanding (NLU) Although explicit knowledge generation (KG) for RG has not been explored, similar methods have been proposed for NLU tasks such as question answering (Shwartz et al., 2020b). Previous work has also explicitly generated *rationales* that can be seen as helpful additional knowledge (Rajani et al., 2019). TBS differs from such work in that we consider a generative task and use the same generative model to do both KG and RG.

7 Conclusion

Inspired by how humans contribute to the common ground during communication, We propose to train RG models that explicitly generate implicit knowledge and then respond (TBS). This brings us three main benefits compared with prior end-to-end RG models: 1) more informative and coherent responses by augmenting with knowledge; 2) generated knowledge provides faithful explanations of RG model's inner-workings; 3) models do not rely on external knowledge bases in response gen-

eration time. We first identify implicit knowledge in dialogues, explore different knowledge representation and transition choices, and demonstrate promising results compared with end-to-end and knowledge-grounded RG models from extensive evaluations. We find strong and promising results for TBS RG model compared with end-to-end RG. In particular, TBS can produce good quality and novel knowledge, outperform end-to-end RG models despite training on less data, and even produce better responses than RG models that take ground-truth knowledge. We hope our findings encourage more future studies on making RG models better emulate human communication process and produce better-quality responses.

Ethics and Broader Impact

Our work aims to train RG models that explicitly generate implicit knowledge before responding. Sheng et al. (2021) have found biases in DialoGPT (our base model) responses and Mehrabi et al. (2021) have found representational harms in common sense resources. We acknowledge that the

generated responses from our models might contain biases. All of the dialogue datasets and models are in English, which benefits English speakers more. We have conducted human evaluation using Amazon Mechanical Turks. We pay turkers around \$15 per hour, well above the highest state minimum wage and engage in constructive discussions if they have concerns about the process. We also give each annotation instance enough time so that we do not pressure annotators.

Acknowledgments

We thank anonymous reviewers for providing insightful feedback and members from Amazon Alexa AI team and INK and JAUNTS lab from USC. Pei Zhou, Jay Pujara, and Xiang Ren’s work on this project was funded by the Defense Advanced Research Projects Agency with award N660011924033. The research was also supported by gifts from Google.

References

- Jacopo Amidei, Paul Piwek, and Alistair Willis. 2019. The use of rating and likert scales in natural language generation human evaluation tasks: A review and some recommendations.
- Satanjeev Banerjee and Alon Lavie. 2005. Meteor: An automatic metric for mt evaluation with improved correlation with human judgments. In *Proceedings of the acl workshop on intrinsic and extrinsic evaluation measures for machine translation and/or summarization*, pages 65–72.
- Steven Bird, Ewan Klein, and Edward Loper. 2009. *Natural language processing with Python: analyzing text with the natural language toolkit*. ” O’Reilly Media, Inc.”.
- Antoine Bosselut, Hannah Rashkin, Maarten Sap, Chaitanya Malaviya, Asli Celikyilmaz, and Yejin Choi. 2019. **COMET: Commonsense transformers for automatic knowledge graph construction**. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 4762–4779, Florence, Italy. Association for Computational Linguistics.
- 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*.
- Jerome S Bruner. 1961. The act of discovery. *Harvard educational review*.
- Chris Callison-Burch, Cameron Fordyce, Philipp Koehn, Christof Monz, and Josh Schroeder. 2007. **(meta-) evaluation of machine translation**. In *Proceedings of the Second Workshop on Statistical Machine Translation*, pages 136–158, Prague, Czech Republic. Association for Computational Linguistics.
- Asli Celikyilmaz, Elizabeth Clark, and Jianfeng Gao. 2020. Evaluation of text generation: A survey. *arXiv preprint arXiv:2006.14799*.
- Herbert H Clark and Susan E Brennan. 1991. Grounding in communication.
- Herbert H Clark and Edward F Schaefer. 1989. Contributing to discourse. *Cognitive science*, 13(2):259–294.
- Mark Craven, Johan Kumlien, et al. 1999. Constructing biological knowledge bases by extracting information from text sources. In *ISMB*, volume 1999, pages 77–86.
- Leyang Cui, Yu Wu, Shujie Liu, Yue Zhang, and Ming Zhou. 2020. **MuTual: A dataset for multi-turn dialogue reasoning**. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 1406–1416, Online. Association for Computational Linguistics.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. **BERT: Pre-training of deep bidirectional transformers for language understanding**. In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 4171–4186, Minneapolis, Minnesota. Association for Computational Linguistics.
- Emily Dinan, Stephen Roller, Kurt Shuster, Angela Fan, Michael Auli, and Jason Weston. 2019. **Wizard of wikipedia: Knowledge-powered conversational agents**. In *7th International Conference on Learning Representations, ICLR 2019, New Orleans, LA, USA, May 6-9, 2019*. OpenReview.net.
- Mihail Eric, Nicole Chartier, Behnam Hedayatnia, Karthik Gopalakrishnan, Pankaj Rajan, Yang Liu, and Dilek Hakkani-Tur. 2021. Multi-sentence knowledge selection in open-domain dialogue. In *Proceedings of the 14th International Conference on Natural Language Generation*, pages 76–86.
- Joseph L Fleiss. 1971. Measuring nominal scale agreement among many raters. *Psychological bulletin*, 76(5):378.
- 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 *Proceedings of the Thirty-Second AAAI Conference on Artificial Intelligence*,

- (AAAI-18), *the 30th innovative Applications of Artificial Intelligence (IAAI-18), and the 8th AAAI Symposium on Educational Advances in Artificial Intelligence (EAAI-18)*, New Orleans, Louisiana, USA, February 2-7, 2018, pages 5110–5117. AAAI Press.
- Karthik Gopalakrishnan, Behnam Hedayatnia, Qinglang Chen, Anna Gottardi, Sanjeev Kwatra, Anu Venkatesh, Raefer Gabriel, and Dilek Hakkani-Tür. 2019. Topical-chat: Towards knowledge-grounded open-domain conversations. In *INTERSPEECH*, pages 1891–1895.
- Ari Holtzman, Jan Buys, Li Du, Maxwell Forbes, and Yejin Choi. 2019. The curious case of neural text degeneration. In *International Conference on Learning Representations*.
- Lishan Huang, Zheng Ye, Jinghui Qin, Liang Lin, and Xiaodan Liang. 2020. Grade: Automatic graph-enhanced coherence metric for evaluating open-domain dialogue systems. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 9230–9240.
- Jena D. Hwang, Chandra Bhagavatula, Ronan Le Bras, Jeff Da, Keisuke Sakaguchi, Antoine Bosselut, and Yejin Choi. 2021. Comet-atomic 2020: On symbolic and neural commonsense knowledge graphs. In *AAAI*.
- Ryan Kiros, Yukun Zhu, Russ R Salakhutdinov, Richard Zemel, Raquel Urtasun, Antonio Torralba, and Sanja Fidler. 2015. Skip-thought vectors. In *Advances in neural information processing systems*, pages 3294–3302.
- J Richard Landis and Gary G Koch. 1977. The measurement of observer agreement for categorical data. *biometrics*, pages 159–174.
- Omer Levy, Minjoon Seo, Eunsol Choi, and Luke Zettlemoyer. 2017. Zero-shot relation extraction via reading comprehension. *arXiv preprint arXiv:1706.04115*.
- Mike Lewis, Yinhan Liu, Naman Goyal, Marjan Ghazvininejad, Abdelrahman Mohamed, Omer Levy, Veselin Stoyanov, and Luke Zettlemoyer. 2020. **BART: Denoising sequence-to-sequence pre-training for natural language generation, translation, and comprehension**. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 7871–7880, Online. Association for Computational Linguistics.
- Linxiao Li, Can Xu, Wei Wu, Yufan Zhao, Xueliang Zhao, and Chongyang Tao. 2020. Zero-resource knowledge-grounded dialogue generation. *arXiv preprint arXiv:2008.12918*.
- Yanran Li, Hui Su, Xiaoyu Shen, Wenjie Li, Ziqiang Cao, and Shuzi Niu. 2017. **DailyDialog: A manually labelled multi-turn dialogue dataset**. In *Proceedings of the Eighth International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pages 986–995, Taipei, Taiwan. Asian Federation of Natural Language Processing.
- Chin-Yew Lin. 2004. Rouge: A package for automatic evaluation of summaries. In *Text summarization branches out*, pages 74–81.
- Pengfei Liu, Weizhe Yuan, Jinlan Fu, Zhengbao Jiang, Hiroaki Hayashi, and Graham Neubig. 2021. Pre-train, prompt, and predict: A systematic survey of prompting methods in natural language processing. *arXiv preprint arXiv:2107.13586*.
- Yinhan Liu, Myle Ott, Naman Goyal, Jingfei Du, Mandar Joshi, Danqi Chen, Omer Levy, Mike Lewis, Luke Zettlemoyer, and Veselin Stoyanov. 2019. **Roberta: A robustly optimized bert pretraining approach**. *ArXiv preprint*, abs/1907.11692.
- Ninareh Mehrabi, Pei Zhou, Fred Morstatter, Jay Pujara, Xiang Ren, and Aram Galstyan. 2021. Lawyers are dishonest? quantifying representational harms in commonsense knowledge resources. In *EMNLP*.
- Mike Mintz, Steven Bills, Rion Snow, and Dan Jurafsky. 2009. Distant supervision for relation extraction without labeled data. In *Proceedings of the Joint Conference of the 47th Annual Meeting of the ACL and the 4th International Joint Conference on Natural Language Processing of the AFNLP*, pages 1003–1011.
- Kishore Papineni, Salim Roukos, Todd Ward, and Wei-Jing Zhu. 2002. Bleu: a method for automatic evaluation of machine translation. In *Proceedings of the 40th annual meeting of the Association for Computational Linguistics*, pages 311–318.
- Alec Radford, Karthik Narasimhan, Tim Salimans, and Ilya Sutskever. 2018. Improving language understanding with unsupervised learning.
- Alec Radford, Jeffrey Wu, Rewon Child, David Luan, Dario Amodei, and Ilya Sutskever. 2019. Language models are unsupervised multitask learners. *OpenAI Blog 1.8 (2019)*: 9.
- Nazneen Fatema Rajani, Bryan McCann, Caiming Xiong, and Richard Socher. 2019. Explain yourself! leveraging language models for commonsense reasoning. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 4932–4942.
- Hannah Rashkin, Eric Michael Smith, Margaret Li, and Y-Lan Boureau. 2019. **Towards empathetic open-domain conversation models: A new benchmark and dataset**. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 5370–5381, Florence, Italy. Association for Computational Linguistics.
- 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*. Association for Computational Linguistics.

- Alan Ritter, Colin Cherry, and William B. Dolan. 2011. [Data-driven response generation in social media](#). In *Proceedings of the 2011 Conference on Empirical Methods in Natural Language Processing*, pages 583–593, Edinburgh, Scotland, UK. Association for Computational Linguistics.
- Stephen Roller, Emily Dinan, Naman Goyal, Da Ju, Mary Williamson, Yinhan Liu, Jing Xu, Myle Ott, Eric Michael Smith, Y-Lan Boureau, and Jason Weston. 2021. [Recipes for building an open-domain chatbot](#). In *Proceedings of the 16th Conference of the European Chapter of the Association for Computational Linguistics: Main Volume*, pages 300–325, Online. Association for Computational Linguistics.
- Iulian Serban, Alessandro Sordoni, Ryan Lowe, Laurent Charlin, Joelle Pineau, Aaron Courville, and Yoshua Bengio. 2017. A hierarchical latent variable encoder-decoder model for generating dialogues. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 31.
- Emily Sheng, Kai-Wei Chang, Prem Natarajan, and Nanyun Peng. 2021. [“nice try, kiddo”: Investigating ad hominem in dialogue responses](#). In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 750–767, Online. Association for Computational Linguistics.
- Vered Shwartz, Peter West, Ronan Le Bras, Chandra Bhagavatula, and Yejin Choi. 2020a. [Unsupervised commonsense question answering with self-talk](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 4615–4629, Online. Association for Computational Linguistics.
- Vered Shwartz, Peter West, Ronan Le Bras, Chandra Bhagavatula, and Yejin Choi. 2020b. Unsupervised commonsense question answering with self-talk. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 4615–4629.
- Alessandro Sordoni, Michel Galley, Michael Auli, Chris Brockett, Yangfeng Ji, Margaret Mitchell, Jian-Yun Nie, Jianfeng Gao, and Bill Dolan. 2015. [A neural network approach to context-sensitive generation of conversational responses](#). In *Proceedings of the 2015 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 196–205, Denver, Colorado. Association for Computational Linguistics.
- Robyn Speer, Joshua Chin, and Catherine Havasi. 2017. Conceptnet 5.5: An open multilingual graph of general knowledge. In *Thirty-first AAAI conference on artificial intelligence*.
- Robert C Stalnaker. 1978. Assertion. In *Pragmatics*, pages 315–332. Brill.
- Yi-Lin Tuan, Wei Wei, and William Yang Wang. 2020. Knowledge injection into dialogue generation via language models. *arXiv preprint arXiv:2004.14614*.
- 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 *Advances in Neural Information Processing Systems 30: Annual Conference on Neural Information Processing Systems 2017, December 4-9, 2017, Long Beach, CA, USA*, pages 5998–6008.
- Ramakrishna Vedantam, C Lawrence Zitnick, and Devi Parikh. 2015. Cider: Consensus-based image description evaluation. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 4566–4575.
- Joseph Weizenbaum. 1966. Eliza—a computer program for the study of natural language communication between man and machine. *Communications of the ACM*, 9(1):36–45.
- Sean Welleck, Ilya Kulikov, Stephen Roller, Emily Dinan, Kyunghyun Cho, and Jason Weston. 2019. Neural text generation with unlikelihood training. In *International Conference on Learning Representations*.
- Thomas Wolf, Lysandre Debut, Victor Sanh, Julien Chaumond, Clement Delangue, Anthony Moi, Pierric Cistac, Tim Rault, Remi 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 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.
- Thomas Wolf, Victor Sanh, Julien Chaumond, and Clement Delangue. 2019. Transfertransfo: A transfer learning approach for neural network based conversational agents. *arXiv preprint arXiv:1901.08149*.
- Sixing Wu, Ying Li, Dawei Zhang, Yang Zhou, and Zhonghai Wu. 2020. Diverse and informative dialogue generation with context-specific commonsense knowledge awareness. In *Proceedings of the 58th annual meeting of the association for computational linguistics*, pages 5811–5820.
- Yan Xu, Etsuko Ishii, Zihan Liu, Genta Indra Winata, Dan Su, Andrea Madotto, and Pascale Fung. 2021. Retrieval-free knowledge-grounded dialogue response generation with adapters. *arXiv preprint arXiv:2105.06232*.
- Yi-Ting Yeh, Maxine Eskenazi, and Shikib Mehri. 2021. A comprehensive assessment of dialog evaluation metrics. *arXiv preprint arXiv:2106.03706*.

- Yizhe Zhang, Siqi Sun, Michel Galley, Yen-Chun Chen, Chris Brockett, Xiang Gao, Jianfeng Gao, Jingjing Liu, and Bill Dolan. 2020a. [DIALOGPT : Large-scale generative pre-training for conversational response generation](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics: System Demonstrations*, pages 270–278, Online. Association for Computational Linguistics.
- Yizhe Zhang, Siqi Sun, Michel Galley, Yen-Chun Chen, Chris Brockett, Xiang Gao, Jianfeng Gao, Jingjing Liu, and Bill Dolan. 2020b. [DIALOGPT : Large-scale generative pre-training for conversational response generation](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics: System Demonstrations*, pages 270–278, Online. Association for Computational Linguistics.
- Tiancheng Zhao, Ran Zhao, and Maxine Eskenazi. 2017. Learning discourse-level diversity for neural dialog models using conditional variational autoencoders. In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 654–664.
- 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, Online. Association for Computational Linguistics.
- Chujie Zheng and Minlie Huang. 2021. Exploring prompt-based few-shot learning for grounded dialog generation. *arXiv preprint arXiv:2109.06513*.
- Hao Zhou, Tom Young, Minlie Huang, Haizhou Zhao, Jingfang Xu, and Xiaoyan Zhu. 2018. [Commonsense knowledge aware conversation generation with graph attention](#). In *Proceedings of the Twenty-Seventh International Joint Conference on Artificial Intelligence, IJCAI 2018, July 13-19, 2018, Stockholm, Sweden*, pages 4623–4629. ijcai.org.
- Pei Zhou, Karthik Gopalakrishnan, Behnam Hedayatnia, Seokhwan Kim, Jay Pujara, Xiang Ren, Yang Liu, and Dilek Hakkani-Tur. 2021a. [Commonsense-focused dialogues for response generation: An empirical study](#). In *Proceedings of the 22nd Annual Meeting of the Special Interest Group on Discourse and Dialogue*, pages 121–132, Singapore and Online. Association for Computational Linguistics.
- Pei Zhou, Pegah Jandaghi, Bill Yuchen Lin, Justin Cho, Jay Pujara, and Xiang Ren. 2021b. Probing causal common sense in dialogue response generation. *arXiv preprint arXiv:2104.09574*.

A TBS Framework Details

A.1 Matching Detail

Hard-Matching This process follows that used in Zhou et al. (2021a). We first identify potential candidates for concepts in ConceptNet (Speer et al., 2017). For each utterance, we use a part-of-speech (POS) tagger to find the nouns, verbs, and adjectives that are not stopwords and then construct a set of potential concepts by including the lemmatized version of these words. The POS tagger, lemmatizer, and stopword list are from the Natural Language Toolkit (NLTK) package (Bird et al., 2009). This step results in a set of concept words for *each turn* of a dialogue.

With a set of concepts we extract for every dialogue turn, we then identify a list of candidate triples (e_1, r, e_2) . We use the ConceptNet containing single-word concepts pre-processed by Zhou et al. (2018). For each concept we identified in a turn, we store all triples in ConceptNet that contain this concept, either as subject or object.

After getting a list of commonsense triples (e_1, r, e_2) containing concepts in a particular turn using ConceptNet, we next examine if any of the *other* entity in the triples appears in the concept set of the next turn. If we find such a match, we record this triple to be a commonsense assertion that might be implied in the response.

Soft-Matching We reuse the first several steps of hard-matching to find a set of candidate triples for each dialogue turn, then instead of searching for the exact words in the next turn, we use embedding similarity from SentenceBERT (Reimers and Gurevych, 2019) (specifically the “*all-MiniLM-L6-v2*” variant, which is claimed to be a “All-round model tuned for many use-cases. Trained on a large and diverse dataset of over 1 billion training pairs”)⁴.

To select the final matched knowledge, we choose the top 3 triples from ConceptNet with the highest similarity. After examining the distribution of embedding similarities from SBERT, we also require the similarity to be above 0.4 to be matched to ensure quality matching.

A.2 Mappings

We show complete mappings of relations from ConceptNet for both relation-converted NL and

⁴https://www.sbert.net/docs/usage/semantic_textual_similarity.html

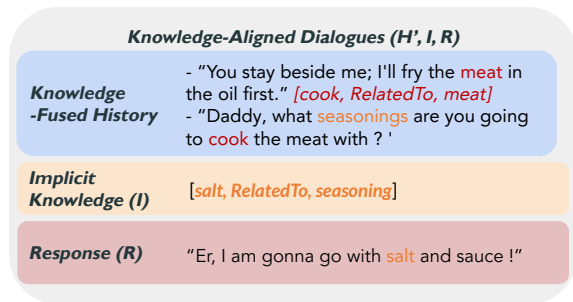


Figure 6: **Data example.** We align implicit knowledge from ConceptNet (Speer et al., 2017) between dialogue turns and form each instance in three components.

information-seeking QA pairs in Table 7.

B Experimental Details

B.1 Implementation Details

We use base models from HuggingFace⁵ and implement TBS based on TransferTransfo (Wolf et al., 2019)⁶. We fine-tune the model for 3 epochs with batch size 4 and set the learning rate to be $6.25e-5$. We perform gradient accumulation for 8 steps and gradient clipping with a max norm of 1.0 and optimize using the Adam optimizer. For decoding, we use top-p nucleus sampling (Holtzman et al., 2019) with temperature T ($p = 0.9$ and $T = 0.7$), and a maximum decoding length of 300 tokens. Note that since we are also generating knowledge, this maximum length is larger than normal RG models. Our TBS models are mostly trained on 4 Quadro RTX 8000 GPUs and take around 5 hours. For automatic metrics, we use the nlg-eval package⁷ and the GRADE repo⁸.

B.2 Evaluation Detail

We present the MTurk interface we use for response quality and knowledge quality evaluation in Figures 7, 8, and 9 including instructions and examples. We require turkers to have at least 500 numbers of HITs approved, with approval rate higher than 95%, and from either Canada, UK, or US since our data is in English.

Relation in ConceptNet	Relation-Converted NL	Information-Seeking QA
DefinedAs	is defined as	What is <concept1>defined as? — <concept1>is defined as <concept2>
DesireOf	desires	What does <concept1>desire of? — <concept1>desires <concept2>
HasA	has a	What does <concept1>have? — <concept1>has <concept2>
HasFirstSubevent	starts with	What does <concept1>start with? — <concept1>starts with <concept2>
HasLastSubevent	ends with	What does <concept1>end with? — <concept1>ends with <concept2>
HasPrerequisite	requires	What does <concept1>require? — <concept1>requires <concept2>
HasProperty	has the property	What property does <concept1>have? — <concept1>is <concept2>
HasSubevent	requires	What subevent does <concept1>have? — <concept1>has subevent of <concept2>
IsA	is a	What is <concept1>? — <concept1>is a <concept2>
MadeOf	is made of	What is <concept1>made of? — <concept1>is made of <concept2>
MotivatedByGoal	is motivated by	What is <concept1>motivated by? — <concept1>is motivated by <concept2>
NotCapableOf	is not capable of	What is <concept1>not capable of? — <concept1>is not capable of <concept2>
NotDesires	does not desire	What does <concept1>not desire? — <concept1>does not desire <concept2>
NotHasA	does not have a	What does <concept1>not have? — <concept1>does not have a <concept2>
NotHasProperty	does not have the property	What property does <concept1>not have? — <concept1>does not have <concept2>
NotIsA	is not a	What <concept1>is not? — <concept1>is not a <concept2>
NotMadeOf	is not made of	What is <concept1>not made of? — <concept1>is not made of <concept2>
PartOf	is part of	What is <concept1>a part of? — <concept1>is a part of <concept2>
RelatedTo	is related to	What is <concept1>related to? — <concept1>is related to <concept2>
SymbolOf	is a symbol of	What is <concept1>a symbol of? — <concept1>is a symbol of <concept2>
UsedFor	is used for	What is <concept1>used for? — <concept1>is used for <concept2>
AtLocation	is located at	Where is <concept1>? — <concept1>is located at <concept2>
CapableOf	is capable of	What is <concept1>capable of? — <concept1>is capable of <concept2>
Causes	causes	What does <concept1>cause? — <concept1>causes <concept2>
CausesDesire	causes the desire to	What desire does <concept1>cause? — <concept1>causes desire of <concept2>
CreatedBy	is created by	What is <concept1>created by? — <concept1>is created by <concept2>
Desires	desires	What does <concept1>desire? — <concept1>desires <concept2>
HasPainCharacter	has pain character of	What pain character does <concept1>have? — <concept1>has pain character of <concept2>
HasPainIntensity	has pain intensity of	What pain intensity does <concept1>have? — <concept1>has pain intensity of <concept2>
InheritsFrom	inherits from	What does <concept1>inherit from? — <concept1>inherits from <concept2>
InstanceOf	is an instance of	What is <concept1>an instance of? — <concept1>is an instance of <concept2>
LocatedNear	is located near	What is <concept1>located near? — <concept1>is located near <concept2>
LocationOfAction	has location of action at	What location of action does <concept1>have? — <concept1>has location of action of <concept2>
ReceivesAction	receives action of	What action does <concept1>receive? — <concept1>received action of <concept2>
Antonym	is an antonym of	What is an antonym of <concept1>? — <concept1>is an antonym of <concept2>
DerivedFrom	is derived from	What is <concept1>derived from? — <concept1>is derived from <concept2>
DistinctFrom	is distinct from	What is <concept1>distinct form? — <concept1>is distinct form <concept2>
EtymologicallyRelatedTo	is etymologically related to	What is <concept1>etymologically related to? — <concept1>is etymologically related to <concept2>
FormOf	is a form of	What is <concept1>a form of? — <concept1>is a form of <concept2>
HasContext	has context of	What context does <concept1>have? — <concept1>has context of <concept2>
SimilarTo	is similar to	What is <concept1>similar to? — <concept1>is similar to <concept2>
Synonym	is a synonym of	What is a synonym of <concept1>? — <concept1>is a synonym of <concept2>
dbpedia:capital	has the capital city	What is the capital city of <concept1>? — <concept1>has capital city of <concept2>
dbpedia:product	has product	What product does <concept1>have? — <concept1>has product of <concept2>

Table 7: Knowledge representation mappings.

C Additional Results

C.1 Models Combining Variants

Table 8 presents the complete results considering all of our models’ variants. We find that the best overall configuration is hard-symbol-QA.

C.2 CEDAR Probing: Do TBS models understand why a response makes sense?

We follow the CEDAR probing framework from Zhou et al. (2021b) that analyzes if RG models assign a higher probability to the response when provided with valid common sense in the form of explanations compared to corrupted explanations. Results comparing to an end-to-end RG model and

a knowledge-selection model are shown in Table 9. We find that by TBS training, RG models become much more sensitive to commonsense explanations against complete corruptions but still fall short against more subtle logical corruptions that require deeper reasoning.

⁵DialoGPT-medium: <https://huggingface.co/microsoft/DialoGPT-medium>

⁶<https://github.com/huggingface/transfer-learning-conv-ai>

⁷<https://github.com/Maluuba/nlg-eval>

⁸<https://github.com/li3cmz/GRADE>

Given a dialogue context and two responses, determine if which of the two responses is more (1) grammatical; (2) coherent; and (3) engaging.

Dialogue History: \${context}

Response 1: : \${baseline_response}

Response 2: : \${model_response}

Which response is more grammatical and fluent?
 Response 1 Response 2 Not sure

Which response is more coherent and has no/small conflicts with the dialogue context?
 Response 1 Response 2 Not sure

Which response is more engaging and makes you want to interact with more?
 Response 1 Response 2 Not sure

Figure 7: **Human evaluation interface** for response quality on dimensions: grammar, coherence, and engagingness.

Given a dialogue context and two responses, determine if which of the two responses is more (1) informative; (2) specific; and (3) follows common sense more.

Dialogue History: \${context}

Response 1: : \${baseline_response}

Response 2: : \${model_response}

Which response is more informative, meaning it carries more content and contributes to the conversations more?
 For example, for a dialogue history: "A: Do you feel anything like that? B: Like what? Bad or miserable?", a response: "A: Like a bad headache or something. It's painful. It feels like I am dying." is MORE informative than "A: Like I am not feeling well"
 Response 1 Response 2 Not sure

Which response is more specific and is less generic?
 For example, for a dialogue history: "A: Can you speak Japanese?", a response: "B: Yep, I learned Japanese for a year in college." is MORE specific and less generic than "B: No I cannot."
 Response 1 Response 2 Not sure

Which response makes more sense and appropriate in terms of common sense?
 For example, for a dialogue history: "A: I want to buy some flowers for my wife", a response: "B: Perhaps you'd be interested in some roses" is makes MORE common sense than "B: Let's sell some flowers."
 Response 1 Response 2 Not sure

Figure 8: **Human evaluation interface** for response quality on dimensions: informativeness, specificity, and common sense.

Given a statement, determine whether the statement makes sense (follows common sense) by itself.
 For examples, "buy is related to money" makes sense but "buy is related to sky" does not.

Statement: : \${statement}

Does the statement make sense (follows common sense)?
 Yes No

Given a statement and a dialogue history, determine whether the statement is relevant to the context of the dialogue.
 For examples, consider the history "A: I just got a raise!", a relevant statement is "raise is related to job" but "raise is related to child" is not relevant in this context, although the latter might make sense by itself.

Dialogue History: : \${context}

Statement: : \${statement}

Is the statement relevant to the context?
 Yes No

Given a statement, a dialogue history, and a next-turn response (not necessarily an appropriate response) continuing the dialogue, determine whether the response makes use of the statement.
 For examples, consider the history "A: I just got a raise!" and response "B: I'm so happy for you", this response makes use of statement "raise has consequence of happy" but the response "B: That's great, you will be richer!" does not.

Dialogue History: : \${context}

Statement: : \${statement}

Response: : \${response}

Does the response make use of the statement?
 Yes No

Figure 9: **Human evaluation interface** for knowledge quality with 3 questions: does the knowledge make sense as a standalone fact, is the knowledge relevant to the context, and does the generated response use the knowledge?

Model Variants	Grammatical	Coherent	Engaging	Informative	Specific	Common Sense	Avg
TBS-soft-symbol-NL	53.0/10.0%	46.3/8.7%	48.7/9.3%	41.7/20.6%	51.7/6%	52/7%	50.5/10.3%
TBS-hard-prompt-NL	50.3/4%	49/7.3%	47/9%	49.4/6%	51/3%	48.3/2.7%	49.2/5.3%
TBS-hard-symbol-QA	53/6.7%	53.6/5.6%	51.3/4.7%	51.3/3.7%	51.3/5%	54/3.7%	52.4/4.8%
TBS-soft-prompt-NL	49.3/6.7%	49.78/8.7%	51.3/4.7%	50.3/2.7%	49.3/8.7%	48.2/6.7%	49.8/5.4
TBS-soft-symbol-QA	51.5/4.2%	52.1/3.5%	51.9/4.9%	49.2/6.7%	49.9/2.7%	45.3/6.9%	51.8/5.6
TBS-hard-prompt-QA	48.3/7.2%	49.9/7.7%	50/5.2%	49.2/5.7%	48.2/6.6%	47.4/2.9%	48.8/6.4
TBS-soft-prompt-QA	50.1/4.7%	50.2/8.7%	49.3/7.9%	48.2/8.7%	48.3/2.7%	49.9/5.7%	49.9/7.2

Table 8: Human evaluation on **response quality** when comparing different model variants with the base model (hard-symbol-NL).

Models	Logical Corruption Average [Accuracy/ Δ NLL]				Complete Corruption Average [Accuracy/ Δ NLL]			
	DD	ED	MuTual	SocialQA	DD	ED	MuTual	SocialQA
<i>Inference Probing</i>								
DialoGPT	0.57/-0.01	0.60/0.03	0.62/0.03	0.64/0.03	0.71/0.15	0.77/0.25	0.79/0.22	0.87/0.40
KS-RoBERTa	0.49/-0.00	0.50/-0.00	0.49/-0.00	0.50/-0.00	0.76/0.23	0.79/0.24	0.78/0.24	0.81/0.27
TBS	0.61/0.15	0.57/0.07	0.57/0.07	0.56/0.05	0.88/1.38	0.86/1.24	0.87/1.14	0.89/1.47
Human	1.0	1.0	0.9	1.0	1.0	1.0	1.0	1.0

Table 9: CEDAR (Zhou et al., 2021b) results where bold-faced numbers indicate statistically significant differences comparing to the second-best model.