

What should I Ask: A Knowledge-driven Approach for Follow-up Questions Generation in Conversational Surveys

Yubin Ge^{*1}, Ziang Xiao^{*1}, Jana Diesner¹, Heng Ji¹, Karrie Karahalios¹, Hari Sundaram¹

¹University of Illinois Urbana-Champaign, USA

{yubinge2, zxiao5, jdiesner, hengji, kkarahal, hs1}@illinois.edu

Abstract

Generating follow-up questions on the fly could significantly improve conversational survey quality and user experiences by enabling a more dynamic and personalized survey structure. In this paper, we proposed a novel task for knowledge-driven follow-up question generation in conversational surveys. We constructed a new human-annotated dataset of human-written follow-up questions with dialogue history and labeled knowledge in the context of conversational surveys. Along with the dataset, we designed and validated a set of reference-free Gricean-inspired (Grice, 1975) evaluation metrics to systematically evaluate the quality of generated follow-up questions. We then propose a two-staged knowledge-driven model for the task, which generates informative and coherent follow-up questions by using knowledge to steer the generation process. The experiments demonstrate that compared to GPT-based baseline models, our two-staged model generates more informative, coherent, and clear follow-up questions.

1 Introduction

A conversational survey collects information from people through an open-ended conversation where an agent asks questions, interprets responses, and probes answers (Xiao et al., 2020b; Gobo, 2011). Compared to structured form-based surveys, a conversational survey enables a more dynamic survey structure and personalized experience through follow-up questions (Xiao et al., 2020a). Although a good follow-up question could probe and prompt more information based on people’s responses, the current practice of delivering follow-up questions in conversational surveys is largely rule-based (Grudin and Jacques, 2019). In this study, we explored an automatic way of generating follow-up questions for conversational surveys.

^{*}denotes equal contribution.

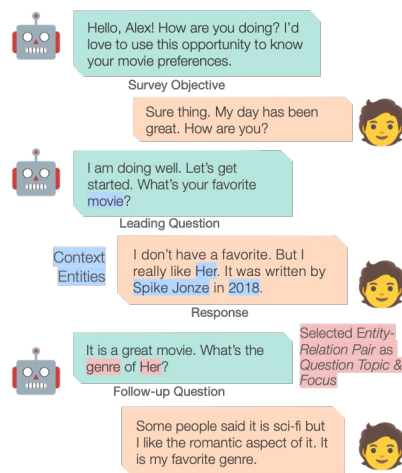


Figure 1: A conversational survey snippet that aims to understand people’s movie preferences where the agent generates a follow-up question based on the dialogue history and a selected entity-relation pair.

Although the conversational survey could be used for various objectives, e.g., personal experiences, public opinion, etc., on different topics, e.g., daily activities, domain knowledge, etc., a single conversational survey often focuses on one survey objective. Therefore, a good follow-up question should be concise and relevant to the survey objective and should be phrased in a clear manner. It requires the system to understand the context deeply, adhere to the survey objective, and generate questions that are easy to understand.

To enable the automatic generation of follow-up questions in an open-domain conversational survey, we aim to tackle three specific challenges. First, there is no dataset for follow-up question generation in the context of open-domain conversational surveys. Related datasets are small or with a very specific focus, such as job interviews (SB et al., 2020) or graduate school admission interviews (Su et al., 2018, 2019). Additionally, no prior dataset considers the background knowledge beyond the dialogue history, which limits the model’s ability to deeply understand the survey objective and context.

Second, existing methods for follow-up question generation are either template filling (Su et al., 2019; Inoue et al., 2020) or seq2seq (Su et al., 2018; Wang et al., 2018; SB et al., 2020). However, both methods lead to unsatisfying results. Template filling limits the diversity of question types and often fails to personalize based on the participant’s response, especially in dynamic and open-ended conversational surveys. As for standard seq2seq methods, they cannot generate questions that adhere to the overall survey objective and are relevant to the context. How to improve the diversity of generated follow-up questions while not limited by templates while utilizing relevant background knowledge remains an unresolved problem.

Third, no established metrics can effectively evaluate the generated follow-up questions for conversational surveys. The same dialogue history can inspire various valid follow-up questions, and the same question can be phrased differently. Hence, common text generation metrics depending on only one ground truth usually underestimate question quality, and human evaluation is hard to scale and compare (Dinan et al., 2018; SB et al., 2020; Xiao et al., 2023).

In this study, we proposed a task of knowledge-driven follow-up question generation in conversational surveys, as shown in Fig. 1. To explore its feasibility and effectiveness, we collected a human-annotated dataset including background knowledge, proposed a two-staged knowledge-driven baseline model, and designed a set of reference-free Gricean-inspired evaluation metrics.

Specifically, we first collected a dataset with knowledge annotation and a human-written follow-up question based on the dialogue history and background knowledge. The follow-up question aimed to collect relevant and valuable information in a coherent and clear manner. To demonstrate feasibility, we proposed a baseline model that leverages knowledge to steer the generative model. We imitated the human question-generation process by using a knowledge selector to identify the question topic and focus (Duan et al., 2008) based on the dialogue history and a background knowledge base. We then combined the dialogue history with the selected question topic and focus as a prompt for a generative model to generate follow-up questions that could lead to valuable and diverse information that contributes to the survey objective. To systematically evaluate the quality of generated follow-

up questions, we designed a new set of reference-free evaluation metrics, *Gricean Scores*, based on Gricean Maxims (Grice, 1975). Gricean Maxims is a collection of communication principles to which both speaker and listener should adhere to engage in effective communication. Our *Gricean Scores* (see Sec. 4) measure the quality of a follow-up question based on the following aspects: *Relevance*, *Informativeness*, *Truthfulness*, *Clarity*, and *Coherence*. *Gricean Scores* align with human evaluation well and can provide more insights than traditional metrics.

Our contributions are as follows:

- A new problem and a dataset for follow-up question generation in conversational surveys, which has background annotation and human-written follow-up questions.
- An effective baseline model that leverages knowledge as a control to generate informative, coherent, and clear follow-up questions.
- A set of reference-free evaluation metrics based on Gricean Maxims, i.e., *Gricean Scores*, that evaluate the quality of open-ended text from various perspectives.

2 Dataset and Annotation

Three principles guide our dataset construction: first, the topic needs to be diverse as conversational surveys can cover many topics. Second, the follow-up question should adhere to the conversation context, and retrieve novel and specific information from the respondent. Third, the knowledge should be explicit, especially salient knowledge that guides the follow-up question generation.

2.1 Question-Answer Pairs Construction

Our dataset is based on OpenDialKG (Moon et al., 2019), which contains open-domain conversations between two human agents about a given topic, such as movies, books, sports, and music. Those topics often appear in conversational surveys to build rapport and collect various responses. Each utterance in the original dataset has been manually annotated with entities from Freebase (Bast et al., 2014). We first extracted question-answer pairs from OpenDialKG and then manually selected those that have a clear information goal, e.g., “how are you doing?” will not be included.

2.2 Follow-up Questions Annotation

We invited workers on Amazon Mechanical Turk to create follow-up questions for each question-answer pair. We first presented a dialogue history, e.g., a question-answer pair, with the overarching survey objective and instructed people to imagine themselves as an interviewer who aims to collect informative and truthful information that contributes to the overall survey objective from their interviewees through a follow-up question. We derive the survey objective based on the topic of the extracted Q & A pair, e.g., [topic] preference. To facilitate the creation of high-quality follow-up questions, we first asked people to select the most interesting and meaningful topic that they wanted to ask about in the follow-up question from mentioned knowledge entities in the dialogue history. Then, they needed to specify the relation of the selected entity as the follow-up question focus. As the last step, we instructed people to write a follow-up question in a clear and coherent manner based on the above-mentioned criteria (Sec. A.2). We manually went through all annotations and made minor edits to ensure quality. We compensated our workers at the rate of \$12 per hour.

3 Models

3.1 Overview

We formalized the task of follow-up question generation as follows: given a dialogue history X_i of the last two turns, consisting of a QA pair (Q_i, A_i) , with a set of context entities $\mathcal{E}_i = \{e_1, e_2, \dots, e_n\}$ (e.g., **Her**, **Spike Jonze** highlighted in Fig. 1) and their relations \mathcal{R} (e.g., the relation **genre** for the entity **Her** highlighted in Fig. 1), a system needs to generate a follow-up question Y_i that is specific to the selected entity-relation pair (e_j, r_k) . Here $e_j \in \mathcal{E}_i$, $r_k \in \mathcal{R}(e_j)$, and we denote $\mathcal{R}(e_j)$ as the set containing all relations connected to the entity e_j .

Inspired by how human experts generate questions in semi-structured interviews (Wilson, 2013), we propose a two-staged framework (as shown in Fig. 2) that (i) selects an entity-relation pair as the question topic and focus that conditioned on the dialogue history (Sec. 3.2), and then (ii) generates a follow-up question that is triggered by the selected entity-relation pair (Sec. 3.3).

3.2 Knowledge Selection

Our knowledge selection model first encodes a dialogue history X_i , a QA pair (Q_i, A_i) , using BERT (Devlin et al., 2019). We process the input text as a concatenation of the question Q_i and its answer A_i , and insert a special tag [CLS] at the beginning and another special tag [SEP] to separate them. The final hidden state that corresponded to [CLS] is used as the aggregated sequence representation of the dialogue history, which is denoted as \mathbf{h}_i^{qa} . Meanwhile, we encode entities and relations with pre-trained knowledge embedding (e.g., TransE (Bordes et al., 2013)), and we represent them as \mathbf{h}_j^e and \mathbf{h}_k^r for the entity e_j and the relation r_k .

We consider two methods to perform the knowledge selection: attention-based and MLP-based. The entity selection works in the same way as the relation selection. Here we only introduce the entity selection for simplicity. In the attention-based version, we regard the score of selecting the entity e_j as the unnormalized dot-product attention score (Vaswani et al., 2017) between \mathbf{h}_j^e and \mathbf{h}_i^{qa} :

$$\mathbf{s}_i = \frac{\mathbf{q}_i \mathbf{K}_i^T}{\sqrt{d_k}}, \quad (1)$$

where \mathbf{q}_i is the query vector such that $\mathbf{q}_i = \mathbf{W}^Q \mathbf{h}_i^{\text{qa}}$, \mathbf{K}_i is the key matrix which is a stack of key vectors $\mathbf{W}^K \mathbf{h}_j^e$ for all $e_j \in \mathcal{E}_i$, and d_k is the dimension of queries and keys. As for the MLP-based version, we concatenate \mathbf{h}_i^{qa} and \mathbf{h}_j^e , then input it into an MLP to predict the score:

$$s_{ij} = \text{MLP}([\mathbf{h}_i^{\text{qa}}; \mathbf{h}_j^e]), \quad (2)$$

where $[\cdot; \cdot]$ denotes the concatenation operation.

Finally, we treat the selection as a binary classification for each candidate entity and set the objective as minimizing the binary cross entropy:

$$\begin{aligned} \mathcal{L}_{\text{ent}} = & -\frac{1}{N} \sum_{i=1}^N \sum_{j: e_j \in \mathcal{E}_i} [y_j^i \cdot \log \sigma(s_{ij}) \\ & + (1 - y_j^i) \cdot \log(1 - \sigma(s_{ij}))], \quad (3) \end{aligned}$$

where $\sigma(\cdot)$ is the sigmoid function, N is the number of samples, y_j^i is the ground truth label for the j^{th} entity in the i^{th} sample, and it equals 1 if the corresponding entity is selected and is 0 otherwise. During the training, we take the sum of \mathcal{L}_{ent} and \mathcal{L}_{rel} as the final objective function.

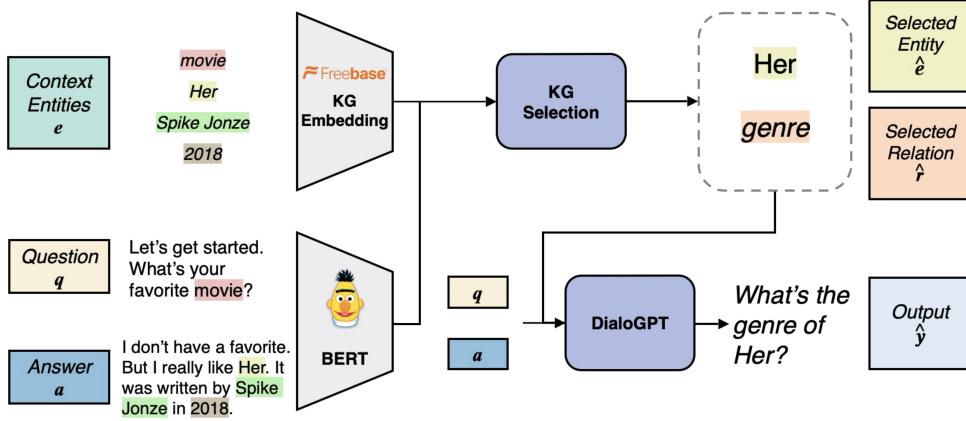


Figure 2: The overall architecture of the proposed framework.

3.3 Follow-up Question Generation

We formulate follow-up question generation as a typical language modeling problem and adopt DialoGPT (Zhang et al., 2020) to solve it.

We serialize the input into one sequence and then input it into DialoGPT to estimate the language model. Specifically, we first concatenate Q_i , A_i and Y_i , and add [EOS] at the end of each utterance such as $[Q_i, \text{EOS}, A_i, \text{EOS}, Y_i, \text{EOS}]$. With the selected entity e_i^* and relation r_i^* , we manually design a discrete template, such as "How to ask about", to combine e_i^* and r_i^* into a prompt P_i . The prompt is treated as an utterance, and is inserted between the dialogue history and the follow-up question to guide the generation. The final input can be represented as one sequence of tokens $[u_1, u_2, \dots, u_m] = [Q_i, \text{EOS}, A_i, \text{EOS}, P_i, \text{EOS}, Y_i, \text{EOS}]$. At each decoding step t , we compute the negative log-likelihood between the predicted distribution and the ground truth from the given response:

$$\mathcal{L}_{\text{NLL}} = - \sum_t \log p(u_t | u_{<t}) \quad (4)$$

Notably, during the inference time, we only input $[Q_i, \text{EOS}, A_i, \text{EOS}, P_i, \text{EOS}]$ into DialoGPT and let it decode the response until reaching [EOS] as the generated follow-up question.

4 Gricean Scores

To systematically assess the quality of generated follow-up questions, we proposed a set of reference-free, customizable and content-based metrics, *Gricean Scores*, grounded on Gricean Maxims (Grice, 1975). The Gricean Maxims is a collection of communication principles, to which both speaker and listener should adhere to engage in effective communication. Researchers have been

using Gricean Maxims to evaluate both human-human conversations (Eskritt et al., 2008; Kleinke, 2010) and human-agent conversations (Xiao et al., 2020b; Langevin et al., 2021). In the context of conducting conversational surveys, a "cooperative" interviewer would obey all the maxims to form questions and probe quality responses. Those maxims include *Quantity*, *Quality*, *Relation*, and *Manner*.

We relied on the Gricean Maxims to define a set of reference-free metrics that quantitatively measure the quality of a follow-up question generated by our models. We measure the quality of a follow-up question and compute *Gricean Scores* from five aspects: *Relevance*, *Informativeness*, *Truthfulness*, *Clarity*, and *Coherence*. Our reference-free metrics allowed us to evaluate a follow-up question without relying on a single ground truth, making it more robust to open-ended text generation and real-world deployment. Also, the *Relevance*, *Informativeness*, and *Truthfulness* are knowledge-driven, which makes our score more contextualized. The survey designers could bring a knowledge graph, customize one, or create one that best fits the survey objective and information value.

Notation. We defined indexed sets \hat{Q}, \mathcal{T} such that a generated follow-up question $\hat{Q}_i \in \hat{Q}$ corresponds to a target follow-up question $\mathcal{T}_i \in \mathcal{T}$. We referred the dialogue history of \hat{Q}_i as its context C_i and further represent \hat{Q}_i as a sequence of tokens $[w_1, w_2, \dots, w_N]$. We marked the recognized entity from \hat{Q}_i as \hat{e}_i , and the set of context entities as \mathcal{E}_i .

Relevance. By the Gricean Maxim of relation, a high-quality follow-up question should be relevant to the prior discussion and adhere to the overall survey objective. Irrelevant follow-up questions elicit useless information and burden the analysis

process. We measured *Relevance* by checking if the recognized entity in the generated follow-up question is from the context entity set:

$$\mathbf{REL}(\hat{Q}_i) = \mathbb{1}[\hat{e}_i \in \mathcal{E}_i] \quad (5)$$

Informativeness. From the Gricean Maxim of Quantity, effective communication should be informative. In the context of a conversational survey, an agent should always ask questions that maximize the information gain in a participant’s response. In a knowledge graph, an entity’s out-degree centrality captures the outreach of other entities which indicates the potential information gain by probing this entity. We measured the *Informativeness* of a generated follow-up question by the out-degree centrality¹ of the recognized entity in the question:

$$\mathbf{INFO}(\hat{Q}_i) = \text{Centrality}(\hat{e}_i) \quad (6)$$

Truthfulness. According to the Gricean Maxim of Quality, a cooperative interlocutor should communicate truthfully. In our case, a high-quality follow-up question should always lead to truthful information. For example, a survey respondent would respond to a question asking for a book’s release year but not a question about a book’s director since the book does not have such an attribute. We measured *Truthfulness* of a follow-up question from an entity-relation perspective where the question should ask a valid relation connected to the recognized entity based on the knowledge base².

Specifically, we first trained a BERT-based relation prediction model (similar to the knowledge selection model introduced in Sec. 3.2) to predict which relation is contained in \hat{Q}_i , and we denoted the predicted relation as \hat{r}_i . Then we measured *Truthfulness* of \hat{Q}_i as checking whether \hat{r}_i is connected to \hat{e}_i in the knowledge base **KG**:

$$\mathbf{TRUTH}(\hat{Q}_i) = \mathbb{1}[(\hat{e}_i, \hat{r}_i) \in \mathbf{KG}] \quad (7)$$

Clarity. The Gricean Maxim of Manner advocates that one should communicate in a clear and orderly manner, and therefore we require a follow-up question to avoid obscurity and ambiguity. We

¹The centrality is normalized by dividing by the maximum possible degree in the graph. For multi-turn conversation, when calculating the out-degree centrality, some edges were removed if mentioned in previous conversation turns.

²To align with the knowledge base we select for our dataset, we choose Freebase here.

regarded *Clarity* of a generated follow-up question as how well it fits in natural language, and we used an external powerful language model as an expert to provide the measurement. Hence, the perplexity of a generated follow-up question is computed under a pre-trained language model and we pick DialoGPT given it’s pre-trained on dialogue data. Specifically, we calculate it as follows:

$$\mathbf{CLA}(\hat{Q}_i) = \exp\left(-\frac{1}{N} \sum_j \log p_\theta(w_j | w_{<j})\right) \quad (8)$$

Coherence. A coherent conversation reflects another aspect of Manner as well. Following previous work (Bommasani and Cardie, 2020), we evaluated the semantic *Coherence* of generated follow-up questions to dialogue histories by predicting the probability of each generated question conditioned on the previous QA pair using the powerful language model, BERT, whose pre-training tasks include this same objective.

$$\mathbf{COH}(\hat{Q}_i) = \mathbb{1}_{\text{BERT}}(\hat{Q}_i | C_i) \quad (9)$$

5 Experiments

We provide the details of experiments in this section and include implementation details in Sec. A.1.

5.1 Knowledge Selection

We first assessed our model’s prediction of the salient knowledge chosen by humans given the dialogue history. The result can inform us of the feasibility of explicitly separating knowledge selection from the full task and choosing the best model to use in our two-staged framework. Meanwhile, separating this task could increase the interpretability of the whole framework by explicitly identifying the knowledge entity and relation that steer the follow-up question generation. We compared the two proposed baseline models, attention-based and MLP-based, and examined how well they interact with different types of knowledge embedding, including TransE (Bordes et al., 2013), TransR (Lin et al., 2015) and TransD (Ji et al., 2015).

Tab. 1 shows the results of recall@k of the selection models. MLP-based models generally perform better than attention-based models for both selection tasks, and the MLP-based model with TransR reached the best performance with 0.654 recall@1 score on entity selection and 0.302 recall@1 score

Model	Entity Selection		Relation Selection	
	R@1	R@1	R@3	R@5
Attention-based				
KG Selector (w/ TransE)	0.642	0.271	0.593	0.761
(w/ TransD)	0.624	0.295	0.616	0.773
(w/ TransR)	0.628	0.288	0.612	0.773
MLP-based				
KG Selector (w/ TransE)	0.647	0.283	0.605	0.768
(w/ TransD)	0.648	0.288	0.613	0.763
(w/ TransR)	0.654	0.302	0.620	0.776

Table 1: Experimental results of various methods on the knowledge selection task.

on relation selection. Based on this, we decided to use the best-performing model, MLP-based with TransR, in the whole framework.

5.2 Follow-up Question Generation With Knowledge

We evaluated our model on the follow-up question generation with selected knowledge in two settings: given the gold-standard knowledge chosen by human annotators, or the knowledge predicted by the best selection model. Additionally, we compared them with the baseline model DialoGPT without adding external knowledge to evaluate the effectiveness of explicitly employing knowledge.

We first used *Gricean Scores* to perform the evaluation and analyze the results using ANOVA tests with Tukey methods (Miller Jr, 1997). As shown in Tab. 2, compared to the baseline model, our proposed knowledge-driven models significantly improved the quality of generated follow-up questions in terms of *Relevance* ($F(2, 3009) = 19.41, p < 0.01^{**}$) and *Truthfulness* ($F(2, 3009) = 85.67, p < 0.01^{**}$). This suggests that integrating selected knowledge is a valid approach to leverage knowledge to guide the generation. For *Clarity* we found the baseline model performed the best ($F(2, 3009) = 15.59, p < 0.01^{**}$). As for *Informativeness* and *Coherence*, three models all achieved good performance with insignificant difference, and we think it's due to the powerful backbone, DialoGPT, we chose. We found the difference between our model with predicted knowledge and gold-standard knowledge for all five dimensions is not statistically significant. This shows that even if the chosen knowledge by the knowledge selection model is different from the ones by human annotators, our generation model can generate appropriate follow-up questions, and demonstrated the applicability of our proposed

method for the open-ended scenario. By examining the outputs, we noticed that including external knowledge through prompts can reduce hallucinations - the fabrication of untruthful information (Maynez et al., 2020). For example, when a dialogue history mentions "Oh, I loved Tom Arnold in *Undercover Blues*", the baseline model that does not involve knowledge generates "Who else starred in *Underpants Blues*", while our model is capable of generating the movie name correctly. This indicates that adding external knowledge through prompts can be a potential solution to entity-level hallucination in text generation.

We further adopt the common metric ROUGE (Lin, 2004) and report the results in Tab. 3. Our experiments again showed that adding knowledge is an effective way to involve knowledge in follow-up question generation, as the two models with knowledge outperform their counterpart, DialoGPT without knowledge. However, we can see there is still a gap between the model with predicted knowledge and gold-standard knowledge, and we think the reason comes from the different knowledge selections, which results in generating totally different follow-up questions. For example, asking about an actor and a movie can be different, but both valid follow-up questions. However, the only one ground truth for one test instance limits the consideration of other possibilities, and this may underestimate models' performance in open-ended text generation. By comparing to the results above, this shows the superiority of our *Gricean Scores* which evaluates multiple aspects of generated text without relying on ground truth text. The reference free nature of the *Gricean Scores* opens its way to the evaluation of other text generation tasks.

Model	REL(%) ↑	INFO ↑	TRUTH(%) ↑	CLA ↓	COH(%) ↑
KG-FQG (w/o knowledge)	60.06	0.47	41.63	2.46	99.07
(w/ predicted knowledge)	72.51	0.55	67.93	2.92	98.55
(w/ gold-standard knowledge)	69.42	0.63	63.35	2.83	98.99

Table 2: Results of various methods on the follow-up question generation task. ↑ indicates the higher score the better, while ↓ means the lower score the better.

Model	R-1	R-2	R-L
KG-FQG (w/o knowledge)	15.46	3.82	15.06
(w/ predicted knowledge)	19.66	7.12	19.32
(w/ gold-standard knowledge)	34.98	16.58	34.00

Table 3: ROUGE scores of various methods on the follow-up question generation task.

5.3 Expert Evaluation

Complementing the result with our objective metrics, we randomly sampled 50 instances from the generated follow-up questions and ground truth texts to perform an expert evaluation. We first discussed a codebook based on four dimensions of Gricean Maxim on a 0 (poor) to 2 (excellent) scale. Then two expert annotators individually rated instances blind to the generation methods. Krippendorff’s alpha ranged from 0.80 to 0.92 for each set of coding. We averaged their scores (Tab. 4) and performed ANOVA tests with Tukey methods to analyze the results. Overall, the results aligned with our proposed *Gricean Scores*, where the knowledge-driven methods outperform the baseline model, DialoGPT, in all dimensions except *Manner*. The Tukey post-hoc tests showed the differences between Ground Truth and our knowledge-driven methods are not statistically significant in all four dimensions, indicating the effectiveness of our approach. In addition, our *Gricean Scores* correlates well with our expert evaluation, Quantity: 0.16, Quality = 0.37*, Relation = 0.24*, and Manner = 0.19³.

6 Error Analysis

Tab. 5 shows several examples that our model with predicted knowledge fails. In the first example, we can observe that our model picks an appropriate mentioned entity, the book *When I’m gone*, to formulate its follow-up question but asks the relation *release_year* which has been covered in the dialogue history. This implies that our current model sometimes may not pay attention to the dialogue history when selecting a relation for asking. We

³* indicates the correlation is significant per Pearson’s test

believe this issue might be solved if we add additional constraints for relation prediction and only allow the knowledge selection model to select unmentioned relations.

Another kind of error can be seen in the second example. The dialogue history mentions *The Runaway Jury*, which is a novel, but our model misunderstands it as a movie even if the phrase "is written by" in the dialogue history indicates it is a book. We notice that such a problem usually happens when different entities share the same name, such as an adapted movie and its original book. Therefore, it is necessary to perform named entity disambiguation during the preprocessing, and exploring better knowledge embedding may be potentially useful in alleviating this problem as well.

7 Related Work

Follow-up Question Generation Previous attempts have explored this task in more specific domains and followed the common pipeline in various text generation tasks (Ge et al., 2021, 2023). For example, several attempts focused on graduate school admission interviews (Su et al., 2018, 2019) and built a small corpus by simulating interviews between participants. They have tried template filling-based and sentence retrieval-based methods for generating follow-up questions. Additionally, job interview is another main domain of the task. To this end, SB et al. (2020) adopted pre-trained GPT-2 and fine-tuned it on their own collected corpus. By contrast, Inoue et al. (2020) pre-defined a set of follow-up questions under different categories and turned the generation task as a question selection problem. Another line of related research lies in question generation in conversational systems. For example, Wang et al. (2018) devised typed decoders to model a type distribution over *interrogatives*, *topic words* and *ordinary words*, and used it to modulate the final generation distribution. By contrast, we focused on the scenario of conversational surveys where the information collection goal is more diverse. We also exploit a

Model	Quantity	Quality	Relation	Manner
DialoGPT	1.41	1.52	1.40	1.62
Ground Truth	1.82**	1.95**	1.82**	1.84
KG-FQG (w/ predicted knowledge)	1.74*	1.86**	1.66	1.84
(w/ gold-standard knowledge)	1.80**	1.74	1.77**	1.75

Table 4: Human Evaluation results of various methods on the follow-up question generation task on Likert Scales from 0-2. The * indicates the difference against DialoGPT is statistically significant. $*p < 0.05$. $**p < 0.01$.

Dialogue History	– Was Two by Two released in 2016? – It was, the same year <u>When I'm Gone</u> was released, which is another similar book.
Predicted Entity:	<i>When I'm Gone</i>
Predicted Relation:	<i>release_year</i>
KG-FQG:	– When was "When I'm gone" released?
Dialogue History	– Do you know <u>The Runaway Jury</u> ? – The Runaway Jury is written by John Grisham, with a genre of Suspense.
Predicted Entity:	<i>The Runaway Jury</i>
Predicted Relation:	<i>subject</i>
KG-FQG:	– What is the subject of the movie?

Table 5: Generated examples from KG-FQG with predicted knowledge for error analysis.

knowledge base to guide the generative model to generate more focused and informative questions.

Dialogue Evaluation Metrics Researchers have shown that standard automatic language evaluation metrics (e.g., BLEU, METEOR) are ineffective for dialogue evaluation (Liu et al., 2016; Deriu et al., 2021; Yeh et al., 2021). Hence, recent research has proposed various automatic metrics specifically for dialogue generation, which can be divided into two categories, reference-required and reference-free, based on whether a reference utterance is required.

Among reference-required metrics, a typical example is RUBER (Tao et al., 2018) which combines a referenced metric and a reference-free metric. The referenced metric computes the cosine similarity of word embeddings between a generated response and a human reference, while a reference-free RNN-based scorer measures the relatedness between a generated response and its context. Based on it, BERT-RUBER (Ghazarian et al., 2019) replaces the RNN in RUBER with BERT (Devlin et al., 2019) to employ contextualized word embeddings as an improvement.

As for reference-free metrics, Lan et al. (2020) proposed PONE based on BERT-RUBER to distinguish positive and negative samples from a training

set. Besides, MAUDE (Sinha et al., 2020) adopts a similar training paradigm but with a different sampling strategy and uses the predicted score of a generated utterance given its context as the metric. Differently, FED (Mehri and Eskenazi, 2020) uses DialoGPT to compute the likelihood of manually designed follow-up utterances to measure multiple qualities of dialog without any supervision.

Our *Gricean Scores* belong to reference-free metrics, and we additionally consider the measurement from the perspective of knowledge base. Another difference is that we utilize different pre-trained language models for different measurements, which are aligned well with their pre-training tasks, so that the gap between downstream inference and pre-training may not be large.

8 Limitations and Future Work

Our work has the following limitations. First, the performance of our knowledge selector is limited by the knowledge embeddings, e.g., TransE. Since the knowledge embeddings mainly focus on the graph structure instead of the semantic meanings of the entities or relations, the selector has a constrained feature space to learn how humans select follow-up question topics and focus. Second, we used a commonsense knowledge graph, Freebase. Although our framework could be adapted to any knowledge graph, the current use of Freebase is not tailored to the survey objective. A conversational survey may have survey objectives that require domain knowledge, such as medical information, a commonsense knowledge graph may not satisfy the need for domain knowledge. Third, the ecological validity of our evaluation is limited. In the future, we plan to integrate our model into a conversational survey and evaluate its effectiveness in information collection with real-world users.

9 Conclusion

We propose a knowledge-driven framework to address three challenges in follow-up question genera-

tion in conversational surveys, diverse and complex responses interpretation, high-quality question construction, and question evaluation. Our framework first selects an entity-relation pair from dialogue histories as question topic and focus, and then uses it to guide a GPT-based model to generate high-quality follow-up questions. To verify the effectiveness of the proposed framework, we collect a new dataset and propose a new set of reference-free evaluation metrics, *Gricean Scores*. Extensive experimental results suggest that our framework outperforms competitive baseline models in both quantitative and qualitative experiments.

References

- Hannah Bast, Florian Bärle, Björn Buchhold, and Elmar Haußmann. 2014. Easy access to the freebase dataset. In *Proceedings of the 23rd International Conference on World Wide Web (WWW)*.
- Rishi Bommasani and Claire Cardie. 2020. Intrinsic evaluation of summarization datasets. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*.
- Antoine Bordes, Nicolas Usunier, Alberto Garcia-Duran, Jason Weston, and Oksana Yakhnenko. 2013. Translating embeddings for modeling multi-relational data. In *Advances in Neural Information Processing Systems (NeurIPS)*.
- Jan Deriu, Alvaro Rodrigo, Arantxa Otegi, Guillermo Echegoyen, Sophie Rosset, Eneko Agirre, and Mark Cieliebak. 2021. Survey on evaluation methods for dialogue systems. *Artificial Intelligence Review*, 54(1):755–810.
- 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 (NAACL-HLT)*.
- Emily Dinan, Stephen Roller, Kurt Shuster, Angela Fan, Michael Auli, and Jason Weston. 2018. Wizard of wikipedia: Knowledge-powered conversational agents. In *International Conference on Learning Representations*.
- Huizhong Duan, Yunbo Cao, Chin-Yew Lin, and Yong Yu. 2008. Searching questions by identifying question topic and question focus. In *Proceedings of Acl-08: HLT*, pages 156–164.
- Michelle Eskritt, Juanita Whalen, and Kang Lee. 2008. Preschoolers can recognize violations of the gricean maxims. *British Journal of Developmental Psychology*, 26(3):435–443.
- Yubin Ge, Ly Dinh, Xiaofeng Liu, Jinsong Su, Ziyao Lu, Ante Wang, and Jana Diesner. 2021. Baco: A background knowledge-and content-based framework for citing sentence generation. In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pages 1466–1478.
- Yubin Ge, Sullam Jeoung, Ly Dinh, and Jana Diesner. 2023. Detection and mitigation of the negative impact of dataset extractivity on abstractive summarization. In *Findings of the Association for Computational Linguistics: ACL 2023*, pages 13963–13976.
- Sarik Ghazarian, Johnny Wei, Aram Galstyan, and Nanyun Peng. 2019. Better automatic evaluation of open-domain dialogue systems with contextualized embeddings. In *Proceedings of the Workshop on Methods for Optimizing and Evaluating Neural Language Generation*, pages 82–89, Minneapolis, Minnesota. Association for Computational Linguistics.
- Giampietro Gobo. 2011. *Back to Likert: Towards a conversational survey*. Sage.
- Herbert P Grice. 1975. Logic and conversation. In *Speech acts*, pages 41–58. Brill.
- Jonathan Grudin and Richard Jacques. 2019. Chatbots, humbots, and the quest for artificial general intelligence. In *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems*, pages 1–11.
- Koji Inoue, Kohei Hara, Divesh Lala, Kenta Yamamoto, Shizuka Nakamura, Katsuya Takanashi, and Tatsuya Kawahara. 2020. Job interviewer android with elaborate follow-up question generation. In *Proceedings of the 2020 International Conference on Multimodal Interaction*.
- Guoliang Ji, Shizhu He, Liheng Xu, Kang Liu, and Jun Zhao. 2015. Knowledge graph embedding via dynamic mapping matrix. In *Proceedings of the 53rd Annual Meeting of the Association for Computational Linguistics and the 7th International Joint Conference on Natural Language Processing (ACL-IJCNLP)*.
- Sonja Kleinke. 2010. Speaker activity and grice’s maxims of conversation at the interface of pragmatics and cognitive linguistics. *Journal of pragmatics*, 42(12):3345–3366.
- Tian Lan, Xian-Ling Mao, Wei Wei, Xiaoyan Gao, and Heyan Huang. 2020. Pone: A novel automatic evaluation metric for open-domain generative dialogue systems. *ACM Transactions on Information Systems (TOIS)*, 39(1):1–37.
- Raina Langevin, Ross J Lordon, Thi Avrahami, Benjamin R Cowan, Tad Hirsch, and Gary Hsieh. 2021. Heuristic evaluation of conversational agents. In *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems*, pages 1–15.

- Chin-Yew Lin. 2004. Rouge: A package for automatic evaluation of summaries. In *Text summarization branches out*, pages 74–81.
- Yankai Lin, Zhiyuan Liu, Maosong Sun, Yang Liu, and Xuan Zhu. 2015. Learning entity and relation embeddings for knowledge graph completion. In *Twenty-ninth AAAI conference on artificial intelligence (AAAI)*.
- Chia-Wei Liu, Ryan Lowe, Iulian Vlad Serban, Mike Noseworthy, Laurent Charlin, and Joelle Pineau. 2016. How not to evaluate your dialogue system: An empirical study of unsupervised evaluation metrics for dialogue response generation. In *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing*, pages 2122–2132.
- Ilya Loshchilov and Frank Hutter. 2018. Decoupled weight decay regularization. In *International Conference on Learning Representations (ICLR)*.
- Joshua Maynez, Shashi Narayan, Bernd Bohnet, and Ryan McDonald. 2020. On faithfulness and factuality in abstractive summarization. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*.
- Shikib Mehri and Maxine Eskenazi. 2020. Unsupervised evaluation of interactive dialog with DialoGPT. In *Proceedings of the 21th Annual Meeting of the Special Interest Group on Discourse and Dialogue*, pages 225–235.
- Rupert G Miller Jr. 1997. *Beyond ANOVA: basics of applied statistics*. CRC press.
- Seungwhan Moon, Pararth Shah, Anuj Kumar, and Rajen Subba. 2019. Opendialkg: Explainable conversational reasoning with attention-based walks over knowledge graphs. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics (ACL)*.
- Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, Alban Desmaison, Andreas Kopf, Edward Yang, Zachary DeVito, Martin Raison, Alykhan Tejani, Sasank Chilamkurthy, Benoit Steiner, Lu Fang, Junjie Bai, and Soumith Chintala. 2019. Pytorch: An imperative style, high-performance deep learning library. In *Advances in Neural Information Processing Systems (NeurIPS)*.
- Pooja Rao SB, Manish Agnihotri, and Dinesh Babu Jayagopi. 2020. Automatic follow-up question generation for asynchronous interviews. In *Proceedings of the Workshop on Intelligent Information Processing and Natural Language Generation*.
- Koustuv Sinha, Prasanna Parthasarathi, Jasmine Wang, Ryan Lowe, William L Hamilton, and Joelle Pineau. 2020. Learning an unreferenced metric for online dialogue evaluation. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 2430–2441.
- Ming-Hsiang Su, Chung-Hsien Wu, and Yi Chang. 2019. Follow-up question generation using neural tensor network-based domain ontology population in an interview coaching system. In *INTERSPEECH*.
- Ming-Hsiang Su, Chung-Hsien Wu, Kun-Yi Huang, Qian-Bei Hong, and Huai-Hung Huang. 2018. Follow-up question generation using pattern-based seq2seq with a small corpus for interview coaching. In *INTERSPEECH*.
- Chongyang Tao, Lili Mou, Dongyan Zhao, and Rui Yan. 2018. Ruber: An unsupervised method for automatic evaluation of open-domain dialog systems. In *Thirty-Second AAAI Conference on Artificial Intelligence*.
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. 2017. Attention is all you need. In *Advances in neural information processing systems (NIPS)*.
- Yansen Wang, Chenyi Liu, Minlie Huang, and Liqiang Nie. 2018. Learning to ask questions in open-domain conversational systems with typed decoders. In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (ACL)*.
- Chauncey Wilson. 2013. *Interview techniques for UX practitioners: A user-centered design method*. Newnes.
- 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 (EMNLP)*.
- Ziang Xiao, Susu Zhang, Vivian Lai, and Q Vera Liao. 2023. Evaluating nlg evaluation metrics: A measurement theory perspective. *arXiv preprint arXiv:2305.14889*.
- Ziang Xiao, Michelle X. Zhou, Wenxi Chen, Huahai Yang, and Changyan Chi. 2020a. *If i hear you correctly: Building and evaluating interview chatbots with active listening skills*. In *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems, CHI '20*, page 1–14, New York, NY, USA. Association for Computing Machinery.
- Ziang Xiao, Michelle X Zhou, Q Vera Liao, Gloria Mark, Changyan Chi, Wenxi Chen, and Huahai Yang. 2020b. Tell me about yourself: Using an ai-powered chatbot to conduct conversational surveys with open-ended questions. *ACM Transactions on Computer-Human Interaction (TOCHI)*, 27(3):1–37.

Yi-Ting Yeh, Maxine Eskenazi, and Shikib Mehri. 2021. A comprehensive assessment of dialog evaluation metrics. In *The First Workshop on Evaluations and Assessments of Neural Conversation Systems*, pages 15–33.

Yizhe Zhang, Siqi Sun, Michel Galley, Yen-Chun Chen, Chris Brockett, Xiang Gao, Jianfeng Gao, Jingjing Liu, and Bill Dolan. 2020. 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 (ACL)*.

A Appendix

A.1 Implementation Details

All the knowledge selection models use BERT_{large} as the backbone encoder which is initialized with the pre-trained uncased weights. We use OpenKE⁴ to pre-train all knowledge embeddings. Models are implemented by Pytorch framework (Paszke et al., 2019) and Huggingface transformers (Wolf et al., 2020). We tune the parameters of each knowledge selection model with the following search space (bold indicate the choice for our final model according to the performance on the validation set): KG embeddings size: {100, 200, 300, **400**, 500}, hidden states: {100, 200, 300, 400, 500, **600**, 700}. The selection models are optimized by AdamW (Loshchilov and Hutter, 2018) with the learning rate of $4e - 5$ and the linear learning scheduler. A default setting trains for 50 epochs, using a batch size of 20. Early stopping is adopted if performance on the validation set doesn't increase for consecutive 10 epochs.

As for the generative model, we initialize it with DialoGPT_{large} and also use AdamW with the linear learning scheduler to fine-tune models. The learning rate is set to $5e - 5$. Similarly, we train the model for 5 epochs with the batch size 1 by default, and use early stopping to stop the training when the performance on the validation set doesn't improve for consecutive 5 epochs. During the generation, we use beam search with beam size 2.

A.2 Dataset Statistics

The final dataset we collect consists of 10040 dialogues, which we divide into 8032 for train, 1004 for validation, and 1004 for the test, and we present a comparison between ours and related datasets in Table 6. A total of 8165 unique entities are

mentioned in the question-answer pairs. On average, each question-answer pair mentioned 2.45 unique entities. And for each mentioned entity, the average number of connected relations is 9.44. More details are shown in Table 7. Regarding the annotated follow-up question for each question-answer pair, there are a total of 3917 'What' questions, 81 'How' questions, 341 'When' questions, 2619 'Who' questions, 12 'Why' questions, 806 'Which' questions, 716 'Where' Questions, and 1548 closed-ended questions.

⁴<http://openke.thunlp.org/>

Datasets	Size	Domain	Knowledge Base
MHMC-IV (Su et al., 2018)	3.4k	School Admission	✗
Interview Coaching (Su et al., 2019)	1.2k	School Admission	✓(ConceptNet)
FQG (SB et al., 2020)	1k	Job Interviews	✗
OURS	10k	General	✓(Freebase)

Table 6: Statistics of our dataset and other previous datasets for follow-up question generation.

	Train	Validation	Test
Number of Dialogues	8032	1004	1004
Number of Utterances	311298	38739	38002
Number of Unique Entities	7429	1910	1920
Avg. Entities per Dialogue	2.44	2.45	2.48

Table 7: Dataset statistics of the Knowledge-Driven Follow-up Question Generation Task.