

Stay Hungry, Stay Focused: Generating Informative and Specific Questions in Information-Seeking Conversations

Peng Qi Yuhao Zhang Christopher D. Manning

Stanford University

Stanford, CA 94305

{pengqi, yuhaozhang, manning}@stanford.edu

Abstract

We investigate the problem of generating informative questions in information-asymmetric conversations. Unlike previous work on question generation which largely assumes knowledge of what the answer might be, we are interested in the scenario where the questioner is not given the context from which answers are drawn, but must reason pragmatically about how to acquire new information, given the shared conversation history. We identify two core challenges: (1) formally defining the informativeness of potential questions, and (2) exploring the prohibitively large space of potential questions to find the good candidates. To generate pragmatic questions, we use reinforcement learning to optimize an informativeness metric we propose, combined with a reward function designed to promote more specific questions. We demonstrate that the resulting pragmatic questioner substantially improves the informativeness and specificity of questions generated over a baseline model, as evaluated by our metrics as well as humans.

1 Introduction

Conversations are a primary means to seek and communicate information between humans, where asking the right question is an important prerequisite for effective exchange of knowledge. Learning to ask questions in conversations can help computer systems not only acquire new knowledge, but also engage human interlocutors by making them feel heard (Huang et al., 2017).

Previous work on question generation often falls into three classes: generating questions according to a discrete schema or end goal (Bordes et al., 2017; Zhang et al., 2018b), transforming the answer statement into a question (Mitkov et al., 2003; Rus et al., 2010; Heilman and Smith, 2010), or generating questions with data-driven systems by conditioning on the context where the answer comes

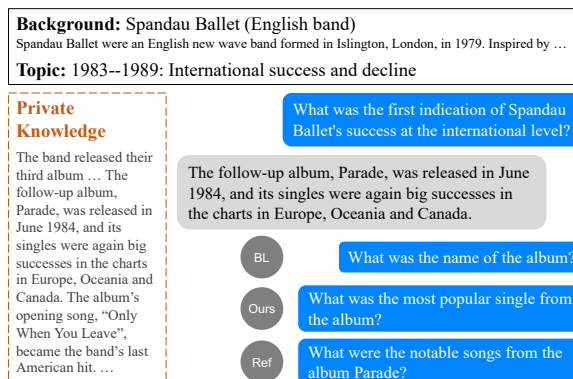


Figure 1: Asking questions in a conversation to acquire information. In this communication setting, the question asker has access to the background and topic, but no access to the private textual knowledge that contains the answer. In this example, the baseline non-pragmatic question generator (BL) generates an uninformative question (one that has already been answered), while our pragmatic system (Ours) and humans (Ref) actively seek new information.

from (Du et al., 2017; Zhou et al., 2017). Despite their successful adaptation to conversations to predict the question that elicits the observed answer (Gao et al., 2019; Pan et al., 2019; Nakanishi et al., 2019), they are not suitable for modeling communication of knowledge in open-domain conversations, because the crucial problem of *what to communicate* has already been assumed to be addressed by conditioning on the schema of information need or the context that contains the answer.

We instead study the problem of question generation in a more realistic setting, *i.e.*, in *open-domain information-seeking conversations* where the question asker cannot access the answering context. This is an important step towards practical natural language processing (NLP) systems that can reason about the state of mind of agents they interact with purely through natural language interactions, so that they can generate more help-

ful responses. In this paper, we build a question generator that reasons pragmatically about what information the answerer can provide, and generates questions to gather new information in a conversation (see Figure 1 for an example).

We identify several key challenges in this task: (1) generating informative questions without access to potential answers; (2) evaluating generated questions beyond comparing them to the reference question, because multiple questions can reveal unseen information despite being very different to each other; (3) navigating a large search space of potential questions to improve informativeness by reasoning about the other agent’s knowledge, which is more complex than limited reference games in previous work on computational pragmatics.

To address these issues, we first develop a baseline question generation model that generates questions in a conversation without conditioning on the unseen knowledge. We then propose automatic metrics to quantify how much new information questions reveal, as well as how specific they are to the conversation. Next, we use reinforcement learning to optimize our question generator on these metrics. In our experiments on the QuAC dataset, we show that the proposed method substantially improves the specificity and informativeness of the generated questions as evaluated by our automatic metrics. These results are corroborated by blinded human evaluation, where questions generated by our system are also of higher overall quality than those by the baseline system as judged by humans. To recap, our main contributions are:¹

- To the best of our knowledge, our work represents the first attempt at studying question generation to seek information in open-domain communication, which involves challenging NLP problems, *e.g.*, evaluation of open-ended language generation and pragmatic reasoning;
- To address these problems, we propose automatic metrics to quantify the informativeness and specificity of questions, which are essential for efficient iterative system development;
- We show that optimizing the proposed metrics via reinforcement learning leads to a system that behaves pragmatically and has improved communication efficiency, as also verified by human evaluation. This represents a practical

¹We release our code and models at <https://github.com/qipeng/stay-hungry-stay-focused>.

method for pragmatic reasoning in an open-domain communication setting.

2 Related Work

Question Generation. Question generation has long been studied in the education and psychology communities as a means to assess and promote reading comprehension in humans (Davey and McBride, 1986). In natural language processing, question generation has been explored to improve the systems in various natural language processing tasks, *e.g.*, the quality of question answering systems (Duan et al., 2017) as well as information retrieval in an open-domain question answering system (Nogueira et al., 2019).

Some of the first question generation systems are rule-based (Mitkov et al., 2003; Rus et al., 2010; Heilman and Smith, 2010), while large-scale question answering datasets, *e.g.*, SQuAD (Rajpurkar et al., 2016, 2018), have recently kindled research interest in data-driven approaches. Du et al. (2017) and Zhou et al. (2017) apply sequence-to-sequence (seq2seq) models to generate SQuAD questions from Wikipedia sentences containing the answers.

The release of large conversational question answering datasets such as QuAC (Choi et al., 2018) and CoQA (Reddy et al., 2019) enabled Gao et al. (2019), Pan et al. (2019) and Nakanishi et al. (2019) to extend previous neural seq2seq question generators by conditioning them on the conversation history and the context that contains the answer, while Scialom and Staiano (2019) remove answers to the reference question to generate curiosity-driven questions from the rest of the context.

Despite their success, most existing approaches to question generation are limited to either *reading comprehension* settings where potential answers are known *a priori*, or *goal-oriented settings* where the schema of knowledge is limited (Bordes et al., 2017; Zhang et al., 2018b). This prevents them from being applied to an *open-domain communication* setting, where the purpose of questions is to acquire information that is unknown ahead of time.

Evaluating System-generated Questions. Automatic evaluation of system-generated text has long been an important topic in NLP. Traditional *n*-gram overlap-based approaches (Papineni et al., 2002; Lin, 2004) are computationally efficient, but have been shown to correlate poorly with human judgement of quality (Novikova et al., 2017). More recently, Zhang et al. (2020) leverage large

pretrained language models (BERT, Devlin et al., 2019) to relax the limitation of exact n -gram overlap. Hashimoto et al. (2019) combine human judgement with system-reported likelihood of generated text to make population-level estimates of quality and diversity. However, most existing metrics either evaluate generated text against very few references, or provide only relative ranking for multiple systems at a population level rather than reliable feedback for each example. This renders them inapplicable to generating informative questions in a conversation, where multiple questions can be equally informative and relevant in a given scenario, and per-example feedback is necessary.

Pragmatic Reasoning for Informativeness.

Pragmatic reasoning is tightly related to informativeness and efficiency in communication. Starting from the cooperative maxims for conversational pragmatic reasoning (Grice, 1975), Frank and Goodman (2012) developed a computational framework that has been applied to reference games with images (Andreas and Klein, 2016) and colors (Monroe et al., 2017), as well as generating descriptions for images (Cohn-Gordon et al., 2019). Decision-theoretic principles (van Rooy, 2003) have also been applied to quantify the informativeness of community questions (Rao and Daumé III, 2018). These approaches usually assume that either the list of *referents* (images, colors, or answers) or the space of utterances (descriptions or questions) is enumerable or can be directly sampled from, or both. More crucially, the speaker agent usually has complete access to this information to readily gauge the effect of different utterances. We instead study a more realistic information-seeking setting, where the questioner cannot access the answers, let alone aggregate them for pragmatic reasoning, and where these simplifying assumptions will not hold.

3 Method

In this section, we outline the setup for the communication problem we set out to address, present a baseline system, and lay out our approach to extending it to reason pragmatically to acquire information more efficiently.

3.1 Problem Setup

We consider a communication game between two agents, a teacher and a student (see Figure 1 for an example). The two agents share a common topic

of discussion \mathcal{T} (*Background* and *Topic* in the figure), as well as a common goal for the student to acquire some knowledge \mathcal{K} on this topic that only the teacher has direct access to (*Private Knowledge* in the figure). We consider the scenario where the agents can only communicate to each other by engaging in a conversation, where the conversation history \mathcal{H} is shared between the agents. We further constrain the conversation to one where the student asks questions about the shared topic, and the teacher provides answers based on \mathcal{K} . Note that this setup is very similar to that of the “Game of Interrogation” by (Groenendijk, 1999), except we relax the definition, using natural language instead of focusing on predicate logic, as we will detail in the sections that follow.

In this paper, we are interested in building a model of the student (question asker) in this scenario. Specifically, we investigate how to enable the student to reason pragmatically about which questions to ask to efficiently acquire knowledge, given only the topic \mathcal{T} and the conversation history \mathcal{H} . This setting of *information-seeking conversations* involves many interesting and challenging problems in natural language processing:

- **Quantifying textual information.** We need to be able to quantify how much knowledge the student has acquired from \mathcal{K} .
- **Evaluating language generation when a single reference is insufficient.** At any state in the conversation, there is usually more than one valid question, some more effective and more appropriate than others. To address this problem, we need to come up with evaluation metrics and objective functions accordingly, rather than relying on the similarity between generated questions and the single reference that is available in existing datasets.
- **Pragmatic reasoning with partial information and a large search space.** In order to train computational agents capable of pragmatic reasoning, previous work typically takes the approach of either limiting the space of referents, or the space of possible utterances, or both. However, the former is infeasible in a communication setting as the student does not have access to \mathcal{K} beyond what is already revealed in the conversation, and the latter is also impractical for natural conversations that cover a diverse set of topics.

We address these challenges by proposing two automatic reward functions that evaluate the informativeness and specificity of questions, and optimizing them with reinforcement learning.

3.2 Generating Questions in Conversations

Before we delve into the proposed approaches for training a question generator model to be pragmatic, an introduction of the model itself is due.

For the purposes of this paper, we assume that the shared topic \mathcal{T} , the shared conversation history \mathcal{H} , and the teacher’s knowledge \mathcal{K} (which the student has no access to) are all made available to agents in natural language. Since we consider information-seeking conversations only, the conversation history is grouped into pairs of questions and answers: $\mathcal{H} = [(q_1, a_1), (q_2, a_2), \dots, (q_{|\mathcal{H}|}, a_{|\mathcal{H}|})]$.

To generate conversational questions, we build a sequence-to-sequence model that encodes the information available to the student and decodes it into the next question in the conversation (see Figure 2(a)). Specifically, we first model the shared topic \mathcal{T} with a bi-directional LSTM (BiLSTM) (Hochreiter and Schmidhuber, 1997), and use the resulting topic representation $h_{\mathcal{T}}$ in the conversation encoder. Then we obtain a representation of the conversation with hierarchical LSTM encoders: we first encode each pair of question and answer with $h_{\mathcal{T}}$ using a BiLSTM, then feed these pair representations into a *unidirectional* LSTM in the direction that the conversation unfolds. To generate the question, we apply an LSTM decoder with attention both on the topic and the conversation history (Bahdanau et al., 2015). This allows us to efficiently batch computation for each conversation by sharing these representations across different turns. We include detailed description of the model in Appendix A.

As a baseline, we train this model to minimize the negative log likelihood (NLL) of questions observed in the training set:

$$\ell_{\text{NLL}} = -\frac{1}{N_{\text{p}}} \sum_{i=1}^N \sum_{j=1}^{|\mathcal{H}^{(i)}|} \log P_{\theta}(q_j^{(i)} | \mathcal{H}_{<j}^{(i)}, \mathcal{T}), \quad (1)$$

where θ stands for model parameters, N is the number total conversations in the training dataset, $\mathcal{H}^{(i)}$ the conversation history of the i -th conversation in the dataset, and $N_{\text{p}} = \sum_{i=1}^N |\mathcal{H}^{(i)}|$ is the total number of question-answer pairs in the training dataset. Intuitively, this trains the model to mimic the observed questions in the dataset, but does not provide guarantees or assessment of how

well generated questions are actually able to acquire information from the teacher agent.

3.3 Evaluating Informativeness through Question Answering

In order to train the question generation model to generate pragmatically apt questions that reveal new information from \mathcal{K} , we need to be able to quantify informativeness in communication first. However, informativeness is difficult to quantify in an open-domain dialogue, and sometimes even subjective. In this paper, we focus on providing an objective metric for how much new information is revealed by a question. Since questions do not reveal information directly, but rather rely on the answers to them to introduce new facts into the conversation, we begin by defining the *informativeness* of an answer a once it is provided. Specifically, we are interested in characterizing how much new information an answer a reveals about \mathcal{K} beyond what is already provided in the conversation history $\mathcal{H}_{<j}$ up until this point in the conversation. Theoretical quantities like mutual information might seem appealing in this context given their strong grounding in information theory. However, applying them would potentially require us to fully specify the state space the world can be in for an open-domain conversation, as well as estimating the probability distribution over potential configurations, neither of which is trivial, if feasible. Therefore, we turn to more practical quantities in defining the informativeness of an answer a given the conversation history $\mathcal{H}_{<j}$ by leveraging the observation that, the more new information an answer reveals about \mathcal{K} , the more likely it involves words that have not already been mentioned in $\mathcal{H}_{<j}$. Therefore, making use of the unigram precision function $\text{Prec}(a, a')$ between the predicted answer a and an answer a' that is already provided in the conversation history $\mathcal{H}_{<j}$, we define the informativeness of the predicted answer as follows

$$I_{\text{ans}}(a; \mathcal{H}_{<j}) := 1 - \max_{1 \leq k < j} \text{Prec}(a, a_k), \quad (2)$$

Intuitively, the more a overlaps with *any* of the previously revealed answers, the less new information it contains. This metric of informativeness has the advantages of objectivity and ease of automatic evaluation. Also note that the choice of unigram precision is here not one of necessity but simplicity and practicality. It is in principle interchangeable

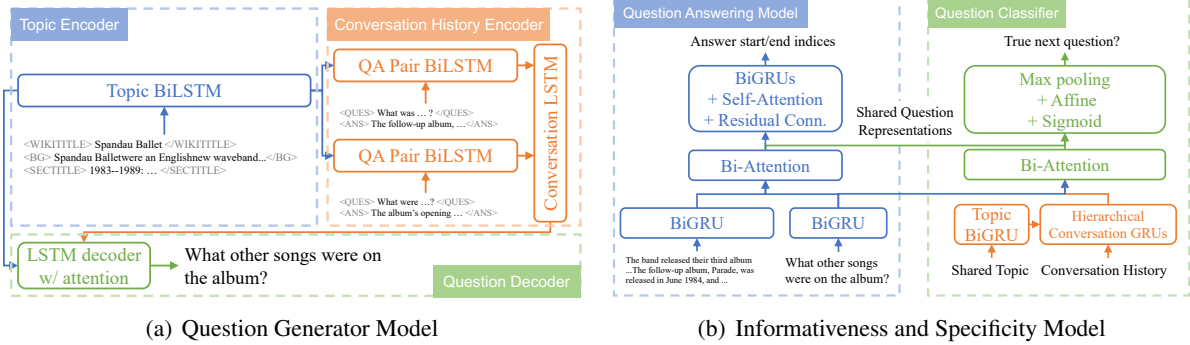


Figure 2: Model architectures of (a) our question generation model, which takes only the shared topic and conversation history to generate the next question in the conversation; and (b) the model to evaluate how informative and specific generated questions are.

with more sophisticated models of fuzzy text overlap (e.g., BERTScore (Zhang et al., 2020)).

We use this definition of answer informativeness to define the utility of potential questions. Specifically, we define the *informativeness* of a question as the amount of new information it can immediately reveal through its answer

$$I(q; \mathcal{C}_{<j}) := I_{\text{ans}}(\text{QA}(q, \mathcal{C}_{<j}), \mathcal{H}_{<j}), \quad (3)$$

where $\mathcal{C}_{<j} = (\mathcal{H}_{<j}, \mathcal{T}, \mathcal{K})$ is the complete context available to the teacher up until the question is raised, $\text{QA}(q, \mathcal{C}_{<j})$ is a pretrained conversational question answering (QA) model that answers the question q from the knowledge source \mathcal{K} given this context. This is equivalent to using a point estimate for $P(a|q, \mathcal{C}_{<j})$ to evaluate q 's expected utility, which is practical for pragmatic reasoning at scale by avoiding the need for aggregating over a large set of candidate answers for each question. In contrast, previous work on pragmatics often require probabilistic normalization in the space of speaker utterances (questions) and listener actions (answers), which is intractable in our setting.

This definition of informativeness is also *explainable*: it is easy for a human to inspect the answer provided by the QA model and compare it to previous ones to understand how much new information has been revealed. Note that this definition itself also does not rely on any specific QA model, although more accurate QA models could result in more accurate estimates of informativeness. For simplicity, we use a bidirectional attention flow model (Seo et al., 2017) with self-attention (Clark and Gardner, 2018) as adapted for conversational QA by Choi et al. (2018) (see Figure 2(b)).

3.4 Evaluating Question Specificity

Now that we have a metric to evaluate informativeness, can we maximize it and obtain a good model for generating pragmatic conversational questions? It turns out that there are two issues with naively optimizing this value: generated questions could be overly generic or disruptive of the conversation flow while still acquiring new information. For instance, questions like *What else?* almost always reveal new information. On the other hand, in our example in Figure 1, *Did they go on tour for their 1983 album?* seems more disruptive (topic-changing) as the next question in the conversation than the candidate questions in the figure.

To address this, we take a similar approach to previous work by selecting negative examples to target these issues and training a classifier to distinguish them from questions that were actually part of the conversation (Lowe et al., 2017; Rao and Daumé III, 2018). Once this classifier is trained, we can make use of the score it assigns different candidate questions to evaluate how specific each is to the current conversation history. Specifically, we select two kinds of negative questions: frequent questions from the training set (frequency > 1) and random questions other than the observed one from the same conversation. We train a model (with shared parameters with the QA model, see Figure 2(b)) to assign a probability that a question is the true next question (positive) given the conversation history, and define this quantity as the *specificity* of the question

$$S(q; \mathcal{H}_{<j}, \mathcal{T}) := P_{\xi}(q \text{ is positive} | \mathcal{H}_{<j}, \mathcal{T}), \quad (4)$$

where ξ is the parameters of the classifier optimized with binary cross entropy loss. Once this classifier

is trained jointly with the QA model, we can use this specificity reward to bias the model towards generating questions that are not only informative, but also specific to the given conversation history.

Conceptually, our specificity idea is related to a few separate but connected concepts in NLP, namely discourse coherence, relevance, and reducing genericness in natural language generation.

The coherence and relevance of a piece of text in a discourse is highly correlated with the perceived quality of the generated text. Previous work has approached generating coherent utterances in conversations through encouraging the model to learn similar distributed representations throughout the conversation (Baheti et al., 2018; Xu et al., 2018; Zhang et al., 2018a). In contrast, we achieve the same goal with a discriminative classifier, which is trained to contrast the true follow-up question (relevant and coherent) against randomly sampled questions (irrelevant) from other conversations and out-of-order questions (uncoherent). The idea of discerning discourse consistency has also been applied to large pretrained language models (Devlin et al., 2019; Iyer et al., 2020), which is demonstrated to sometimes yield performance gains when they are finetuned on downstream tasks.

On the other hand, since we sample frequent questions in the training set as negative examples for the classifier, it also discourages the model from generating overly generic questions. Previous work has attacked the problem of genericness in conversational natural language generation by proposing auxiliary training objectives, *e.g.*, ones that maximize the utility of the generated utterance estimated with adversarial networks (Rao and Daumé III, 2019), specificity estimates that are estimated from data (Ko et al., 2019a,b), or the mutual information between the generated turn and previous ones (Li et al., 2016). Our proposed method can be viewed as a generalization of these approaches, where the objective to be optimized at the time of generation is implicitly specified via a parameterized model by choosing negative examples for contrast.

3.5 Generating Informative and Specific Questions

Given the informativeness metric and specificity reward, we can improve upon these by maximizing the following reward function that blends the two

in a weighted sum

$$R(q; \mathcal{C}_{<j}) = \lambda_1 I(q; \mathcal{C}_{<j}) + (1 - \lambda_1) S(q; \mathcal{H}_{<j}, \mathcal{T}). \quad (5)$$

Since this quantity can only be evaluated once a complete question has been generated, the non-differentiability of the decoding process prevents us from directly optimizing it with respect to θ using gradient-based optimization. However, we can still estimate the gradient of the expected reward of generated questions, $\mathbb{E}_{q \sim P_\theta} [R(q)]$ using REINFORCE (Williams, 1992), a reinforcement learning technique. For an example q , the gradient estimate is the gradient of the following loss function

$$\ell_R(q) = -[R(\hat{q}) - b(q)] \log P_\theta(\hat{q}) \quad (6)$$

where \hat{q} is a sample from P_θ and we dropped the dependency on $\mathcal{C}_{<j}$ for notational clarity. $b(q)$ is called the *baseline function*, which, if chosen carefully, reduces the variance of this gradient estimate and results in faster convergence. We apply a technique called self-critical sequence training (Rennie et al., 2017), which selects $b(q) = R(q_G)$, the reward obtained by the greedily decoded sequence, q_G , from the question generator.

To ensure that the generator maximizes the desired reward function without losing fluency in generated questions, we combine ℓ_R with negative log likelihood during model finetuning (Paulus et al., 2018). We finetune a pretrained question generator (with ℓ_{NLL}) using the following objective

$$\ell = \lambda_2 \ell_R + (1 - \lambda_2) \ell_{\text{NLL}}. \quad (7)$$

Here, $\ell_R = \frac{1}{N_p} \sum_{i=1}^N \sum_{j=1}^{M_i} \ell_R(q_j^{(i)})$. We choose $\lambda_1 = 0.5$ and $\lambda_2 = 0.98$ in our experiments, which were chosen by tuning the model on the dev set.

4 Experiments

Data. For our experiments, we use the QuAC dataset presented by Choi et al. (2018). Although other similar datasets share some common characteristics, some crucial differences render them inapplicable for our experiments. For instance, CoQA (Reddy et al., 2019) gives both agents access to the context, while Wizard of Wikipedia (Dinan et al., 2019) does not assign the student agent clear goals of acquiring new information.

Since QuAC’s test set is held private for fair evaluation, for this work we repurpose the original dev set as our test set. We randomly split the training set

Split	Orig.	# Entities	# Dialogues	# QA pairs
Train	Train	2727	10572	76338
Dev	Train	264	995	7230
Test	Dev	721	1000	7354

Table 1: Data split of the QuAC dataset used in our experiments.

into training and development partitions, ensuring that the Wikipedia entities discussed in conversations do not overlap between these partitions. The goal of the split is to obtain a development set that is roughly as large as the repurposed test set. The statistics of our data split can be found in Table 1.

Training. We follow the recipe available in AllenNLP (Gardner et al., 2018) to train the QA model on QuAC, and make sure that it obtains performance on par with that reported by Choi et al. (2018) on the official dev set (with multiple answer references).² We use the Adam optimizer (Kingma and Ba, 2015) with default hyperparameters to train and finetune our question generator, and anneal the learning rate by 0.5 whenever dev performance does not improve for more than 3 consecutive epochs (patience=3). When training finishes, the specificity classifier achieves approximately 75% F₁ on the dev set when the true next question, sampled frequent questions and random questions from the same conversation have a balanced ratio of 1:1:1. For unanswerable questions in QuAC, we revise Equation (3) and set informativeness to zero if the predicted answer is CANNOTANSWER, as the answer does not reveal new information about the hidden knowledge \mathcal{K} .

5 Results

5.1 Metric-based Evaluation

For the baseline model and our model finetuned for informativeness and specificity, we generate predictions with greedy decoding for simplicity. We evaluate them on conventionally used metrics such as perplexity (PPLX) of the reference question and the F₁ score of the ROUGE-L metric (ROUGE-L) (Lin, 2004) between the predicted questions and the reference. The former helps verify the overall quality of our model, while the latter helps us compare single-reference metrics to our proposed ones. We also report the informativeness metric (INFO)

²However, in practice, we remove the ELMo component, which greatly speeds up computation at the cost of only losing 2–3 F₁ in answer prediction.

and specificity reward (SPEC) for these models, and compare them to the reference questions on these measures on both the dev and test sets.

As shown in Table 2, the baseline model and our pragmatically finetuned model achieve comparable performance when evaluated against the reference question using n -gram overlap metrics (ROUGE-L), and the perplexity of the reference question is only slightly worse. As expected, these metrics tell us nothing about how well the model is going to fare in actual communication, because perplexity does not evaluate the usefulness of generated questions, and ROUGE-L can barely tell these systems apart.

We can also see in Table 2 that the finetuned model improves upon the baseline model on both informativeness and specificity. Further, we notice that despite their high specificity, the reference questions are only about as informative as our baseline questions on average, which is a bit surprising at first sight. Further analysis reveals that about 12.6% of dev questions and 15.7% test ones are considered unanswerable by crowd workers, which is a byproduct of the information-asymmetric setting adopted when the data was collected. As a result, many reference questions could be considered uninformative by our definition, since they might cause the QA model to abstain from answering.

5.2 Human Evaluation

Although the results in Table 2 show that our model sees substantial improvements on the proposed informativeness and specificity metrics, it remains unclear whether these improvements correlate well with human judgement of quality, which is critical in the application of the resulting system. To study this, we conduct a comparative human evaluation.

We randomly selected 200 turns from the test set, and asked two NLP PhD students to evaluate the reference questions, as well as those generated by the baseline model and our model. These questions are evaluated on their overall quality, informativeness, and specificity, where the annotators are asked to rank the candidate questions on each metric with ties allowed. System identity is hidden from the annotators, and the order of the systems is shuffled for each comparison. Prior to annotation, both annotators were educated to follow the same guidelines to encourage high agreement (see Appendix D for details).

As shown in Table 3, human annotators favor our system over the baseline on informativeness

System	dev				test			
	PPLX	ROUGE-L	INFO	SPEC	PPLX	ROUGE-L	INFO	SPEC
Reference	–	–	0.639	0.825	–	–	0.635	0.821
Baseline	8.46	25.5	0.653	0.761	8.08	24.3	0.664	0.779
Our Model	9.24	25.6	0.740	0.834	8.79	24.9	0.752	0.835

Table 2: Evaluation of the baseline system, our pragmatically finetuned system, and the reference questions on conventional metrics as well as our proposed metric and reward functions.

	Win	Tie	Lose	<i>p</i> -value
Ours vs. Baseline				
Overall	29.00%	52.50%	18.50%	0.006
INFO	19.25%	74.50%	6.25%	3×10^{-5}
SPEC	25.50%	56.50%	18.00%	0.029
Ours vs. Human Reference				
Overall	17.75%	26.25%	56.00%	$< 2 \times 10^{-6}$
INFO	8.75%	74.50%	19.75%	3×10^{-4}
SPEC	15.25%	29.00%	55.75%	$< 2 \times 10^{-6}$

Table 3: Human evaluation comparing questions our system generated to those from the baseline, as well as the original reference questions in QuAC. We perform a bootstrap test with 10^6 samples for the difference between pairs of systems and report the *p*-values here.

(93.75% of our questions are considered equally or more informative), and to a lesser extent, overall quality (81.5%) and specificity (82%). We find that 26.1% of questions generated by these systems are identical, which inflates the number of ties in human evaluation. We expect a starker contrast if a sampling-based decoding strategy were applied for generation diversity, which we leave to future work. We also attribute this difference in human-perceived quality on these three aspects partly to the inherent nature of these annotation tasks: while our annotators agree on 80.3% of the pair-wise judgments regarding informativeness, agreement decreases to 70.7% for overall quality and 69.2% for specificity since they are more subjective. It is encouraging, however, that our system is also considered equally or more informative than the human reference 80.25% of the time. What negatively affects human’s perception of the overall quality of questions our system generates is largely attributable to the over-genericness of these questions compared to the references, and a sometimes blatant lack of common sense (e.g., questions like “What did he do after his death?”).

Background: Spandau Ballet (English band)
Topic: 1983–1989: International success and decline
Candidate Questions
BL₁: What happened in 1983?
Ours₁: What happened in 1983?
Ref₁: What was the first indication of Spandau Ballet’s success at the international level?
Ans₁: The follow-up album, <i>Parade</i> , was released in June 1984, and its singles were again big successes in the charts in Europe, Oceania and Canada.
BL₂: What was the name of the album?
Ours₂: What was the most popular single from the album?
Ref₂: What were the notable songs from the album <i>Parade</i> ?
Ans₂: The album’s opening song, “ <i>Only When You Leave</i> ”.
BL₃: What was the name of the album that was released?
Ours₃: What other songs were on the album?
Ref₃: How did the opening song do on the charts?
Ans₃: Became the band’s last American hit.
BL₄: What was the last album that they released?
Ours₄: What other songs were on the album?
Ref₄: Are there any other interesting aspects about this article?

Figure 3: A success example where our automatic metrics align well with human judgement of informativeness and specificity, when comparing questions generated by the baseline (BL), our system (Ours), and the original human-written reference (Ref). All answers in the figure are from the original human-human conversations in QuAC in answer to reference questions.

6 Analysis

We further analyze concrete examples of generated questions in conversations to understand the behavior of our informativeness and specificity metrics.

Case Study. To sanity check whether our informativeness metric and specificity reward match human intuition, we manually inspect a few examples from the test set. Figure 3 represents a case where all the questions our system generated are considered equal to or more informative than the reference and baseline generated questions by our metric. As shown in the example, the baseline system is prone to generating topical but uninformative questions (BL₂ and BL₃). Our system finetuned on our reward function is more pragmatic and asks about relevant questions that can likely be answered from the unseen paragraph \mathcal{K} . Our informativeness

Background: Alison Moyet (English singer)
Topic: 1990s: Further recordings and hiatus
Candidate Questions
BL₁: What was the first album released?
Ours₁: What was her first recording?
Ref₁: What did she do in 1990?
INFO: $Ref = BL = Ours$ SPEC: $Ref \approx Ours > BL$
Ans₁: CANNOTANSWER
BL₂: What was her first album?
Ours₂: What was her first album?
Ref₂: What else did she record in the 1990's?
INFO: $Ref = BL = Ours$ SPEC: $BL=Ours>Ref$
Ans₂: Hoodoo.
BL₃: What was the name of the album?
Ours₃: What was her first album?
Ref₃: Did he reach bill board charts?
INFO: $BL = Ref > Ours$ SPEC: $Ours>BL>Ref$
Ans₃: Yes, and Moyet was nominated for a Grammy for the single
BL₄: What was the song called?
Ours₄: What other songs did she release?
Ref₄: Did she receive an award?
INFO: $BL = Ours > Ref$ SPEC: $BL \approx Ours \approx Ref$

Figure 4: A failure case of our informativeness metric and specificity reward, when comparing reference questions (Ref), baseline generated questions (BL), our those generated by our model (Ours). All answers in the figure are from the original human-human conversations in QuAC in answer to reference questions.

metric also correctly identifies that both Ours₃ and Ref₃ are good questions that reveal new information about \mathcal{K} , although there is very little overlap between the two. On the other hand, the specificity reward successfully identifies that BL₃ and Ref₄ are the least specific questions of their respective turn, where the former is disconnected from the most recent topic under discussion (the song), the latter is phrased in an overly generic way.

We also demonstrate some clear failure cases. In Figure 4, we see that our informativeness and specificity measures make judgements a human will unlikely make, as the topic implies \mathcal{K} is unlikely to contain information about Moyet’s first album/recording. In fact, the QA model fails to recognize that these questions (BL_{1,2}, Ours_{1,2,3}, Ref₁) are unanswerable, and instead assigns them high informativeness. The specificity model, on the other hand, fails to recognize near paraphrases (BL₁ vs Ours₁) and a question that was likely just answered (BL₃). A positive finding in this example is that the informativeness metric is well-aligned with pragmatic behavior in the fourth turn—had Moyet won the Grammy, the previous answer (A₃) would have mentioned it instead of just her nomination.

We include the answering contexts for these examples in Appendix B for the reader’s reference.

Explainable Informativeness. As stated in Section 3.3, our definition of informativeness is explainable to humans—we demonstrate this with concrete examples. For instance, in the example in Figure 3, although the question *What happened in 1983?* is phrased rather vaguely, the QA model is able to identify its correct answer from the paragraph *The band released their third album, True, in March 1983*, which offers new information (note the answer in the figure only reflects the actual human-human conversation, not this hypothetical one). Similarly, the QA model correctly identifies that the question our model generated on the second turn (Ours₂) has the same answer as the human reference (Ref₂), which introduces a new entity into the conversation. BL₂ and BL₃ are deemed uninformative in this case since the QA model offered the same answer about the album *True* again. Although this answer is about an incorrect entity in this context (the album *True* instead *Parade*, which is the focus of discussion), the large amount of overlap between this answer and Ans₁ is still sufficient to regard these questions as less informative.

We note that this informativeness metric does have an exploitable flaw—it does not prevent the questioner from asking vague, open-ended questions (*e.g.*, *What else do you know?*) to acquire knowledge. In fact, we find this strategy is also adopted by QuAC’s crowd workers. However, our specificity reward penalizes genericness, and therefore alleviates this issue in the questions our system generates. We show that our system repeats n -grams from previous questions less frequently, and refer the reader to Appendix C for details.

7 Conclusion

In this paper, we presented a question generation system in information-seeking conversations. By optimizing our proposed automatic metrics for informativeness and specificity, the model is able to generate pragmatically relevant and specific questions to acquire new information about an unseen source of textual knowledge. Our proposed method presents a practical if shallow implementation of pragmatics in an open-domain communication setting beyond simple reference games. We hope that our work brings the community’s attention to this important problem of natural language communication under information asymmetry.

Acknowledgments

The authors would like to thank Christopher Potts, Kawin Ethayarajh, Alex Tamkin, Allen Nie, and Emrah Budur, among other members of the Stanford NLP Group, for their feedback on earlier versions of this paper. We are also grateful to the anonymous reviewers for their insightful comments and constructive discussions. This research is funded in part by Samsung Electronics Co., Ltd. and in part by the SAIL-JD Research Initiative.

References

- Jacob Andreas and Dan Klein. 2016. Reasoning about pragmatics with neural listeners and speakers. In *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing*, pages 1173–1182.
- Dzmitry Bahdanau, Kyunghyun Cho, and Yoshua Bengio. 2015. Neural machine translation by jointly learning to align and translate. In *International Conference on Learning Representations*.
- Ashutosh Baheti, Alan Ritter, Jiwei Li, and Bill Dolan. 2018. Generating more interesting responses in neural conversation models with distributional constraints. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*.
- Antoine Bordes, Y-Lan Boureau, and Jason Weston. 2017. Learning end-to-end goal-oriented dialog. In *International Conference on Learning Representations*.
- Kyunghyun Cho, Bart van Merriënboer, Caglar Gulcehre, Dzmitry Bahdanau, Fethi Bougares, Holger Schwenk, and Yoshua Bengio. 2014. Learning phrase representations using RNN encoder–decoder for statistical machine translation. In *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing (EMNLP)*.
- Eunsol Choi, He He, Mohit Iyyer, Mark Yatskar, Wentaoh Yih, Yejin Choi, Percy Liang, and Luke Zettlemoyer. 2018. QuAC: Question answering in context. In *Conference on Empirical Methods in Natural Language Processing (EMNLP)*.
- Christopher Clark and Matt Gardner. 2018. Simple and effective multi-paragraph reading comprehension. In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 845–855.
- Reuben Cohn-Gordon, Noah Goodman, and Christopher Potts. 2019. An incremental iterated response model of pragmatics. In *Proceedings of the Society for Computation in Linguistics (SCiL) 2019*, pages 81–90.
- Beth Davey and Susan McBride. 1986. Effects of question-generation training on reading comprehension. *Journal of Educational Psychology*, 78(4):256.
- 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.
- Emily Dinan, Stephen Roller, Kurt Shuster, Angela Fan, Michael Auli, and Jason Weston. 2019. Wizard of Wikipedia: Knowledge-powered conversational agents. In *International Conference on Learning Representations (ICLR)*.
- Xinya Du, Junru Shao, and Claire Cardie. 2017. Learning to ask: Neural question generation for reading comprehension. In *Association of Computational Linguistics*.
- Nan Duan, Duyu Tang, Peng Chen, and Ming Zhou. 2017. Question generation for question answering. In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing*, pages 866–874.
- Michael C Frank and Noah D Goodman. 2012. Predicting pragmatic reasoning in language games. *Science*, 336(6084):998–998.
- Yifan Gao, Piji Li, Irwin King, and Michael R. Lyu. 2019. Interconnected question generation with coreference alignment and conversation flow modeling. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 4853–4862.
- Matt Gardner, Joel Grus, Mark Neumann, Oyvind Tafjord, Pradeep Dasigi, Nelson F. Liu, Matthew Peters, Michael Schmitz, and Luke Zettlemoyer. 2018. AllenNLP: A deep semantic natural language processing platform. In *Proceedings of Workshop for NLP Open Source Software (NLP-OSS)*, pages 1–6.
- Herbert P Grice. 1975. Logic and conversation. In *Syntax and Semantics*, volume 3, pages 41–58. Academic Press.
- Jeroen Groenendijk. 1999. The logic of interrogation. In *Semantic and Linguistics Theory*.
- Tatsunori Hashimoto, Hugh Zhang, and Percy Liang. 2019. Unifying human and statistical evaluation for natural language generation. 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 1689–1701.

- Michael Heilman and Noah A Smith. 2010. Good question! Statistical ranking for question generation. In *Human Language Technologies: The 2010 Annual Conference of the North American Chapter of the Association for Computational Linguistics*, pages 609–617.
- Sepp Hochreiter and Jürgen Schmidhuber. 1997. Long short-term memory. *Neural computation*, 9(8):1735–1780.
- Karen Huang, Michael Yeomans, Alison Wood Brooks, Julia Minson, and Francesca Gino. 2017. It doesn’t hurt to ask: Question-asking increases liking. *Journal of personality and social psychology*, 113(3):430.
- Dan Iter, Kelvin Guu, Larry Lansing, and Dan Jurafsky. 2020. Pretraining with contrastive sentence objectives improves discourse performance of language models. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*.
- Diederik P Kingma and Jimmy Ba. 2015. Adam: A method for stochastic optimization. In *International Conference on Learning Representations*.
- Wei-Jen Ko, Greg Durrett, and Junyi Jessy Li. 2019a. Domain agnostic real-valued specificity prediction. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 33, pages 6610–6617.
- Wei-Jen Ko, Greg Durrett, and Junyi Jessy Li. 2019b. Linguistically-informed specificity and semantic plausibility for dialogue generation. 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 3456–3466.
- Jiwei Li, Will Monroe, Alan Ritter, Dan Jurafsky, Michel Galley, and Jianfeng Gao. 2016. Deep reinforcement learning for dialogue generation. In *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing*.
- Chin-Yew Lin. 2004. ROUGE: A package for automatic evaluation of summaries. In *Text Summarization Branches Out*, pages 74–81.
- Ryan Lowe, Michael Noseworthy, Iulian Vlad Serban, Nicolas Angelard-Gontier, Yoshua Bengio, and Joelle Pineau. 2017. Towards an automatic Turing test: Learning to evaluate dialogue responses. In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1116–1126.
- Ruslan Mitkov et al. 2003. Computer-aided generation of multiple-choice tests. In *Proceedings of the HLT-NAACL 03 workshop on Building educational applications using natural language processing*, pages 17–22.
- Will Monroe, Robert X. D. Hawkins, Noah D. Goodman, and Christopher Potts. 2017. Colors in context: A pragmatic neural model for grounded language understanding. *Transactions of the Association for Computational Linguistics*, 5:325–338.
- Mao Nakanishi, Tetsunori Kobayashi, and Yoshihiko Hayashi. 2019. Towards answer-unaware conversational question generation. In *Proceedings of the 2nd Workshop on Machine Reading for Question Answering*, pages 63–71.
- Rodrigo Nogueira, Wei Yang, Jimmy Lin, and Kyunghyun Cho. 2019. Document expansion by query prediction. *arXiv preprint arXiv:1904.08375*.
- Jekaterina Novikova, Ondřej Dušek, Amanda Cercas Curry, and Verena Rieser. 2017. Why we need new evaluation metrics for NLG. In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing*, pages 2241–2252.
- Boyuan Pan, Hao Li, Ziyu Yao, Deng Cai, and Huan Sun. 2019. Reinforced dynamic reasoning for conversational question generation. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 2114–2124.
- 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.
- Romain Paulus, Caiming Xiong, and Richard Socher. 2018. A deep reinforced model for abstractive summarization. In *International Conference on Learning Representations*.
- Jeffrey Pennington, Richard Socher, and Christopher D Manning. 2014. GloVe: Global vectors for word representation. In *Proceedings of the 2014 conference on empirical methods in natural language processing (EMNLP)*, pages 1532–1543.
- Ofir Press and Lior Wolf. 2017. Using the output embedding to improve language models. In *Proceedings of the 15th Conference of the European Chapter of the Association for Computational Linguistics: Volume 2, Short Papers*, pages 157–163.
- Pranav Rajpurkar, Robin Jia, and Percy Liang. 2018. Know what you don’t know: Unanswerable questions for SQuAD. In *Association of Computational Linguistics (ACL)*.
- Pranav Rajpurkar, Jian Zhang, Konstantin Lopyrev, and Percy Liang. 2016. SQuAD: 100,000+ questions for machine comprehension of text. In *Conference on Empirical Methods in Natural Language Processing (EMNLP)*.
- Sudha Rao and Hal Daumé III. 2018. Learning to ask good questions: Ranking clarification questions using neural expected value of perfect information. In

- Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 2737–2746.
- Sudha Rao and Hal Daumé III. 2019. Answer-based Adversarial Training for Generating Clarification Questions. 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)*.
- Siva Reddy, Danqi Chen, and Christopher D. Manning. 2019. [CoQA: A conversational question answering challenge](#). *Transactions of the Association for Computational Linguistics*, 7:249–266.
- Steven J Rennie, Etienne Marcheret, Youssef Mroueh, Jerret Ross, and Vaibhava Goel. 2017. Self-critical sequence training for image captioning. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 7008–7024.
- Robert van Rooy. 2003. Questioning to resolve decision problems. *Linguistics and Philosophy*, 26(6):727–763.
- Vasile Rus, Brendan Wyse, Paul Piwek, Mihai Lintean, Svetlana Stoyanchev, and Cristian Moldovan. 2010. The first question generation shared task evaluation challenge. In *Proceedings of the 6th International Natural Language Generation Conference*.
- Thomas Scialom and Jacopo Staiano. 2019. Ask to learn: A study on curiosity-driven question generation. *arXiv preprint arXiv:1911.03350*.
- Minjoon Seo, Aniruddha Kembhavi, Ali Farhadi, and Hannaneh Hajishirzi. 2017. Bidirectional attention flow for machine comprehension. In *International Conference on Learning Representations (ICLR)*.
- Ronald J Williams. 1992. Simple statistical gradient-following algorithms for connectionist reinforcement learning. *Machine learning*, 8(3-4):229–256.
- Xinnuo Xu, Ondřej Dušek, Ioannis Konstas, and Verena Rieser. 2018. Better conversations by modeling, filtering, and optimizing for coherence and diversity. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*.
- Hainan Zhang, Yanyan Lan, Jiafeng Guo, Jun Xu, and Xueqi Cheng. 2018a. Reinforcing coherence for sequence to sequence model in dialogue generation. In *Proceedings of the Twenty-Seventh International Joint Conference on Artificial Intelligence, IJCAI-18*, pages 4567–4573.
- Junjie Zhang, Qi Wu, Chunhua Shen, Jian Zhang, Jianfeng Lu, and Anton Van Den Hengel. 2018b. Goal-oriented visual question generation via intermediate rewards. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 186–201.
- Tianyi Zhang, Varsha Kishore, Felix Wu, Kilian Q. Weinberger, and Yoav Artzi. 2020. [BERTScore: Evaluating text generation with BERT](#). In *International Conference on Learning Representations*.
- Qingyu Zhou, Nan Yang, Furu Wei, Chuanqi Tan, Hangbo Bao, and Ming Zhou. 2017. Neural question generation from text: A preliminary study. In *National CCF Conference on Natural Language Processing and Chinese Computing*, pages 662–671.

A Model Details

In this section, we include the details of the question generation model and the informativeness/specificity model we used in our experiments.

A.1 Question Generator

For the input to the encoder and decoder models in our question generator, we tokenize them with the spaCy toolkit,³ and initialize word representations with 100-dimensional GloVe vectors (Pennington et al., 2014). As shown in Figure 2, we also introduce special XML-like symbols to delimit different parts of the input to various models. The representations of these special symbols are randomly initialized, and finetuned with those of the top 1000 most frequent words in the training set during training.

For the topic containing the title of the Wikipedia page and the background on the entity after concatenating them with special symbols, we feed them into a topic BiLSTM model and obtain the topic representation with a multi-layer perceptron (MLP) attention mechanism, using the concatenated final state from each direction of the BiLSTM as the key

$$h_{\mathcal{T}} = \text{BiLSTM}_{\mathcal{T}}(x_{\mathcal{T}}), \quad (8)$$

$$h_{\mathcal{T}}^{\text{attn}} = \text{MLP}_{\text{SelfAttn}}(h_{\mathcal{T}}). \quad (9)$$

We use this representation to initialize the BiLSTM we use to encode each pair of turns in the conversation that contains a question and its corresponding answer

$$h_{\mathcal{H}_j}^0 = \text{BiLSTM}_{\mathcal{H}, \text{pair}}(x_{\mathcal{H}_j}, h_{\mathcal{T}}^{\text{attn}}), \quad (10)$$

which we in turn use as the input to our unidirectional LSTM model to obtain the representation of the entire conversation up until a certain turn

$$h_{\mathcal{H}} = \text{BiLSTM}_{\mathcal{H}, \text{conv}}([h_{\mathcal{H}_1}^0, \dots, h_{\mathcal{H}_{|\mathcal{H}|}}^0]). \quad (11)$$

Note that for modeling simplicity, we use the section title as the 0th “turn” for each conversation. We similarly obtain a summary representation of the conversation with MLP self-attention

$$h_{\mathcal{H}}^{\text{attn}} = \text{MLP}_{\text{SelfAttn}}(h_{\mathcal{H}}), \quad (12)$$

and concatenate it with $h_{\mathcal{T}}^{\text{attn}}$ to initialize the decoder.

³<https://spacy.io/>

To represent the input words in the decoder, we use the same embedding matrix as the encoder. We also employ weight tying between the input embeddings and the output weights for word prediction to reduce parameter budget (Press and Wolf, 2017). For each word in the decoder input, we concatenate its embedding with $h_{\mathcal{T}}^{\text{attn}}$ for topical context. We provide the decoder access through attention to all of the representations of encoded tokens, *i.e.*, $[h_{\mathcal{T}}^{\text{attn}}, h_{\mathcal{H}_0}^0, \dots, h_{\mathcal{H}_{|\mathcal{H}|}}^0]$. Finally, the weighted average of encoder representations is combined with the decoder LSTM’s representation of the decoded sequence to yield a probabilistic distribution over words in the vocabulary.

A.2 Informativeness/Specificity Model

For informativeness, we follow closely the open implementation of BiDAF++ for QuAC that is available in AllenNLP (Gardner et al., 2018). For each word, we concatenate its word representations with character representations derived from a convolutional neural network from its character spelling. We replace the ELMo embeddings with GloVe ones for computational efficiency, which results in a relatively small drop in QA performance compared to AllenNLP’s implementation (by about 2–3 F₁ on the official dev set). Note that following Choi et al. (2018), we use gated recurrent units (GRUs; Cho et al., 2014) in this part of the model.

For the specificity model, we first encode the topic and conversation history in a similar fashion as we did for the encoder in the question generator. Then, this representation is combined with the question representation from the BiGRU encoder in the QA model via a bidirectional attention (bi-attention) mechanism. The resulting representation is combined with the question representation from the bi-attention in the QA model, and max pooled over time, before an affine transform is applied to convert the representation into a score.

B Contexts for Case Study Examples

We include in Figures 5 and 6 the contexts that contain the answer for the examples we studied in Section 6, with gold answers in the case study highlighted in the paragraphs. Following Choi et al. (2018), we concatenate an artificial CANNOTANSWER token to the end of the paragraph for the question answering model to abstain from answering the question.

The band released their third album, *True*, in March 1983. Produced by Tony Swain and Steve Jolley, the album featured a slicker pop sound. It was at this point that Steve Norman began playing saxophone for the band. Preceded by the title track which reached number one in various countries, the album also reached number one in the UK. Their next single, “Gold”, reached number 2. [The follow-up album, *Parade*, was released in June 1984, and its singles were again big successes in the charts in Europe, Oceania and Canada.]_{Ans1} [The album’s opening song, “Only When You Leave”]_{Ans2}, [became the band’s last American hit.]_{Ans3} At the end of 1984, the band performed on the Band Aid charity single and in 1985 performed at Wembley Stadium as part of Live Aid. During this same year, Spandau Ballet achieved platinum status with the compilation *The Singles Collection*, which kept the focus on the band between studio albums and celebrated its five years of success. However, the album was released by Chrysalis Records without the band’s approval and the band instigated legal action against the label. In 1986, Spandau Ballet signed to CBS Records and released the album *Through the Barricades*, in which the band moved away from the pop and soul influences of *True* and *Parade* and more toward rock. Though the first single, “Fight for Ourselves” peaked at 15 in the UK, the title track and the album both reached the Top 10 in the UK and Europe. After a hiatus from recording, the band released their next album, *Heart Like a Sky*, in September 1989. The album and its singles were unsuccessful in the UK, and the album itself was not released in the United States. It did, however, do well in Italy (where its singles “Raw” and “Be Free with Your Love” reached the Top 10) and also in Belgium, Germany and the Netherlands. CANNOTANSWER

Figure 5: Private context that contains the answers to questions in our case study example in Figure 3.

C Specificity Analysis

We examine the outputs of our model to assess whether finetuning on the specificity reward results in more specific questions rather than generic and repetitive ones. To measure this, we compute the n -gram overlap between generated questions and all questions in the conversation history for all systems. The lower this repetition is, the more likely the system is bringing up new entities or topics in its questions, and thus more specific to the given conversation history. As can be seen in Figure 7, our system improves upon the baseline system by reducing this repetition noticeably in longer n -grams ($n \geq 3$). When n is very large ($n \geq 8$), our pragmatic system is less repetitive even compared to the human reference, which often contains long and repetitive questions like *Are there any other interesting aspects about this article?* as a generic inquiry for more information.

Following a period of personal and career evaluation, [Hoodoo]_{Ans2} was released in 1991. The album sold respectably in the UK, [and Moyet was nominated for a Grammy for the single “It Won’t Be Long”]._{Ans3} However, the release of *Hoodoo* marked the beginning of an eight-year fight for Moyet to secure complete control of her artistic direction. Like many similar artists (including Aimee Mann and the late Kirsty MacColl), Moyet was reluctant to record a radio-friendly “pop” album simply for the sake of creating chart hits. Moyet’s next album, *Essex* (1994), was also a source of controversy for her; in order for the album to be released, her label (now Sony) insisted that certain *Essex* tracks be re-recorded and re-produced, and that there be additional material remixed to create a more ‘commercial’ package. The video for the single “Whispering Your Name” again featured Dawn French. Following the release of *Essex*, Sony released a greatest hits compilation of Moyet’s work. Singles entered the UK charts at No. 1 and, following a UK tour, was re-issued as a double CD set which included “Live (No Overdubs)”, a bonus live CD. Upon re-issue, Singles charted again, this time in the Top 20. Due to prolonged litigation with Sony, Moyet did not record or release a new studio album for over eight years after the release of *Essex*. During this time, however, she recorded vocals for Tricky, Sylk-130, Ocean Colour Scene, The Lightning Seeds, and King Britt, and was featured on the British leg of the Lilith Fair tour. 2001 saw the release of *The Essential Alison Moyet* CD, and in 2002 *The Essential Alison Moyet* DVD. In 1995, she sang back-up vocals with Sinéad O’Connor for one of Dusty Springfield’s last television appearances, singing “Where Is a Woman to Go ?” on the music show *Later With Jools Holland*. [CANNOTANSWER]_{Ans1}

Figure 6: Private context that contains the answers to questions in our case study example in Figure 4.




Figure 7: Proportion of repeated n -grams in questions from the conversation history. As can be seen from the plot, our pragmatic system reduces the amount of n -grams repeated from previous questions especially for longer n -grams.

D Human Evaluation Details

In this section, we include further details about how the human evaluation is carried out to compare different systems.

We begin by randomly sampling 200 turns of questions from the 7354 question-answer pairs in the test set, and collect the questions from the human reference, the baseline system, and our fine-tuned system. Then, for each turn, we shuffle the order of the three candidate questions, and present them in a group to the annotators. The questions are accompanied with the entity and topic under discussion, as well as the conversation history from the QuAC dataset that led up to the turn under evaluation. We ask the annotators to provide a ranking amongst the questions from these unidentified systems, and allow ties when the annotators cannot observe a qualitative difference between two or more question candidates. An example annotation task for a turn of questions can be found in Figure 8. To encourage high inter-annotator agreement, we first conducted a trial annotation on 20 examples on the dev set, and composed annotation guidelines (see Figure 9) with some minimal examples to clarify edge cases.

Q/D183. Prince Philip , Duke of Edinburgh

Prince Philip of Greece and Denmark was born in Mon Repos on the Greek island of Corfu on 10 June 1921 , the only son and fifth and final child of Prince Andrew of Greece and Denmark and Princess Alice of Battenberg . Philip 's four elder sisters were Margarita , Theodora , Cecilie , and Sophie . He was baptised into the Greek Orthodox Church . His godparents were Queen Olga of Greece (his paternal grandmother) and the Mayor of Corfu .

Section discussed: Naval and wartime service

Q: when did philip join the military ?
A: After leaving Gordonstoun in 1939 , Philip joined the Royal Navy ,

Q: did he serve in any war ?
A: Yes , During the Second World War , he continued to serve in the British forces ,

Q: when did he retire from the forces ?
A: In January 1946 , Philip returned to the United Kingdom on the Whelp , and was posted as an instructor at HMS Royal Arthur , the Petty Officers ' School in Corsham , Wiltshire .

Q: did philip go to college ?
A: Yes , graduating the next year from the Royal Naval College , Dartmouth , as the best cadet in his course .

Q: were there other wars that he was involved in ?
A: After the invasion of Greece by Italy in October 1940 , he was transferred from the Indian Ocean to the battleship HMS Valiant in the Mediterranean Fleet .

Q: was he in the force with other royals ?
A: Yes , two of his brothers - in - law , Prince Christoph of Hesse and Berthold , Margrave of Baden , fought on the opposing German side .

Q: was philip ever promoted while in the force ?
A: Yes , Philip was appointed as a midshipman in January 1940 .

Q: are there any interesting aspects in the article ?
A: Yes , On 1 February 1941 , he was commissioned as a sub - lieutenant after a series of courses at Portsmouth ,

Q: what was the highest grade that he got to while in the navy ?
A: Promotion to lieutenant followed on 16 July 1942 .

Q: could you name a ship that he had been on when serving ?
A: Yes , stoking the boilers of the troop transport ship RMS Empress of Russia .

Question A: was he doing the job of stoking the boilers ?
Question B: what was his rank in the navy ?
Question C: what did he do after that ?

Q/D184. Rank the Overall Quality of each question

	1	2	3
Question A	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
Question B	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
Question C	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>

Q/D185. Rank the Informativeness of each question

	1	2	3
Question A	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
Question B	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
Question C	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>

Q/D186. Rank the Specificity of each question

	1	2	3
Question A	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
Question B	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
Question C	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>

Figure 8: The annotation interface for the human evaluation, which is built with Qualtrics.com. The annotators are given the title of the Wikipedia article under discussion, a short introductory paragraph, the title of the Wikipedia section discussed, the QuAC conversation history that leads up to the current turn, and the candidate questions to evaluate. They are asked to rank these candidate questions on overall quality, informativeness, and specificity, as we outline in the guidelines.

Evaluating Questions in an Information-Gathering Conversation

In this task, you will be asked to read a conversation between two agents on a given topic (an entity from Wikipedia, e.g., "Albert Einstein"), and evaluate a set of follow-up questions as candidates for the next utterance in the conversation. More specifically, the agents discuss about a given section in that Wikipedia article (e.g., "Early Life").

Only one of the two agents, the **teacher**, or answerer, has access to the text of the section, from which answers are provided. The **student's** (asker's) goal is to have a meaningful conversation and gather information from this unseen section of text through the conversation.

Setting

You will be provided the same information that is available to the **student**, i.e., the shared conversational topic (Wikipedia page title, a short introductory paragraph), the section title under discussion, as well as the entire history of conversation between the teacher and the student.

Task

Your task is to evaluate the quality of three candidate questions for each combination of topic under discussion, section title, and conversation history. You will be ranking these questions on three different evaluation metrics, where ties are allowed for any metric (and encouraged if there isn't a clear signal setting candidate questions apart). Specifically, you will be evaluating these questions on their

- **Overall Quality.** A good question should be fluent, specific, and moves the conversation forward. Does this question seem relevant to the conversation? Does it move the conversation forward by gathering more information? Is it grammatical and/or fluent?
 - If you had to choose one of these questions to ask as the student, in which order will you choose these questions (ties are allowed)?
- **Informativeness.** A good question in this setting should gather new information that hasn't already been revealed by the teacher. Does this question attempt to gather new information from the section under discussion?
 - Note that a question doesn't truly gather new information from the section if references in it are phrased too vaguely to be resolved to anything specific in the conversation history, or if it asks about something completely irrelevant to the (unseen) section under discussion.
 - Depending on the context, a seemingly repetitive question can actually gather more information (e.g., asking about other films an actor/actress has appeared in given the knowledge of some of his/her films). Use your best judgement in these cases.
- **Specificity.** A good question should also be tightly related to the topic under discussion, as well as what has just been discussed. Is this question specific for the current conversation, merely applicable to general discussions about this topic, applicable to discussions about virtually any topic, or worse, obviously irrelevant to the current discussion?
 - Note that pronoun use (e.g., "her", "it") shouldn't be discounted as less specific than mentioning the specific entities, as they are commonly used to refer to topics or entities under discussion.

Figure 9: Human evaluation guidelines to compare system-generated questions with the human reference.