

Dialog Inpainting: Turning Documents into Dialogs

Zhuyun Dai^{*1} Arun Tejasvi Chaganty^{*1} Vincent Zhao^{*1} Aida Amini¹ Qazi Mamunur Rashid¹
Mike Green¹ Kelvin Guu^{*1}

Abstract

Many important questions (e.g. “How to eat healthier?”) require conversation to establish context and explore in depth. However, conversational question answering (ConvQA) systems have long been stymied by scarce training data that is expensive to collect. To address this problem, we propose a new technique for synthetically generating diverse and high-quality dialog data: *dialog inpainting*. Our approach takes the text of any document and transforms it into a two-person dialog between the writer and an imagined reader: we treat sentences from the article as utterances spoken by the writer, and then use a dialog inpainter to predict what the imagined reader asked or said in between each of the writer’s utterances. By applying this approach to passages from Wikipedia and the web, we produce WikiDialog and WebDialog, two datasets totalling 19 million diverse information-seeking dialogs—1,000x larger than the largest existing ConvQA dataset. Furthermore, human raters judge the *answer adequacy* and *conversationality* of WikiDialog to be as good or better than existing manually-collected datasets. Remarkably, our approach shows strong zero-shot capability, generating high quality synthetic data without using any in-domain ConvQA data. Using our inpainted data to pre-train ConvQA retrieval systems, we significantly advance state-of-the-art across three benchmarks (QReCC, OR-QuAC, TREC CAsT) yielding up to 40% relative gains on standard evaluation metrics.

1. Introduction

Modern information-seeking tools such as web search and question answering (Karpukhin et al., 2020; Zhu et al.,

^{*}Equal contribution ¹Google Inc., Mountain View, USA. Correspondence to: Dialog Inpainting Core <dialog-inpainting-core@google.com>.

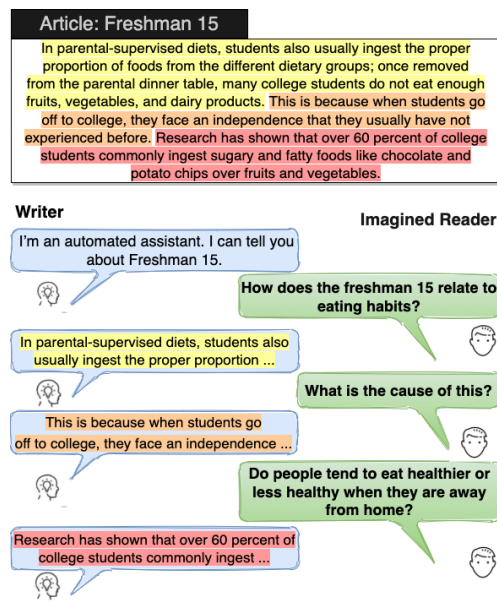


Figure 1. A real example of a dialog inferred from a Wikipedia passage using dialog inpainting. Highlighted utterances are original sentences from the article. All other utterances are generated by the dialog inpainter.

2021) excel at questions that have well-defined answers (e.g., “Where was Barack Obama born?”). But many important questions are more open-ended—e.g., “How to eat healthier?”—and require conversation to elicit context and explore in depth: “How do I eat more protein?”, “What about vegetarians?”. *Conversational* question answering systems (ConvQA) (Stede & Schlangen, 2004; Radlinski & Craswell, 2017; Culpepper et al., 2018), would empower users to answer these questions as if they could discuss with an expert at any time.

Despite this promising vision, progress has been stymied by scarce training data. While conversational data is abundant in online forums, much of it focuses on personal anecdotes and subjective opinions, and is thus unsuitable for an information-seeking system: we desire responses that minimize personal biases and cite reliable sources. Directly crowdsourcing dialogs is also hard: crowdworkers are rarely experts in the domain of interest and tend to overlook impor-

tant questions or provide shallow answers (Li et al., 2021). It is also expensive: the largest extant datasets contain only about 10,000 conversations each (Choi et al., 2018; Reddy et al., 2019; Dinan et al., 2018; Saeidi et al., 2018; Campos et al., 2020; Feng et al., 2020; Anantha et al., 2021).

On the other hand, high-quality documents, such as those in Wikipedia or PubMed, are abundant. These documents are often edited or written by experts who have invested significant time streamlining their discourse and anticipating a reader’s questions. What if we could rewrite these documents into dialogs between the writer and their imagined reader? This would yield an enormous corpus of information-seeking dialogs with attributable, expert answers that could then be used to train a ConvQA system. We aim to achieve this with dialog inpainting.

To transform any document into a dialog, our first observation is that we already know what the writer wants to discuss—that is reflected in the original text of the document. So, we pretend that the original document is the transcript of what the writer said to the reader in an imagined dialog. But we are still missing what the reader asked. This is like overhearing someone else’s phone call: you hear one side, but not the other. Oftentimes, one can still guess what the other side was saying — we call this prediction task *dialog inpainting*, because we are “inpainting” the missing parts of the dialog that we did not hear (inspired by the term’s usage in computer vision (Iizuka et al., 2017; Liu et al., 2018; Yu et al., 2018)). Drawing on this intuition, we train an *inpainter* model to predict missing utterances in a dialog, and use it to predict the unobserved questions in a document. By interleaving the generated questions and sentences from the document, we form a dialog (Figure 1).

We apply our inpainter to passages from Wikipedia and the web, yielding WikiDialog and WebDialog,¹ two datasets totalling 19M+ dialogs — 1,000x larger than the largest existing ConvQA dataset. When evaluated for *conversationality* and *answer adequacy*, we surprisingly find that our synthetically generated data is as good or better than previous crowd-sourced datasets (Section 3). Our generated dialogs inherit the good qualities of the professionally written documents we inpaint (topical diversity, coherent discourse, evidence-backed claims, etc.) without needing to train on dialog data of the same quality.

Importantly, we find that our inpainted datasets are powerful sources of training data for ConvQA systems (Section 4). When used to pre-train standard retriever and reranker architectures, they advance state-of-the-art across three different ConvQA retrieval benchmarks (QRECC, OR-QUAC, TREC-CAST), delivering up to 40% relative gains on stan-

dard evaluation metrics (Section 5). Remarkably, we find that just pre-training on WikiDialog enables strong zero-shot retrieval performance—up to 95% of a finetuned retriever’s performance—without using any in-domain ConvQA data.

2. Dialog Inpainting

The goal of dialog inpainting is to take a *partial dialog* (one where some of the speaker turns are unobserved), and generate a *complete dialog* (one where all unobserved turns have been filled in with the model’s predictions).

Formally, a complete dialog d is a sequence of speaker utterances, $d = (u_1, u_2, \dots, u_t, \dots, u_T)$. We use the same notation for partial dialogs, denoting unobserved utterances with the \diamond symbol. For example, $(u_1, u_2, \diamond, u_4, \diamond)$ is a partial dialog where utterances u_3 and u_5 are unobserved. We refer to these as “masked” utterances. We also use the shorthand $d_{m(3,5)}$ to denote a dialog d with utterances 3 and 5 masked.

To complete the partial dialog $d_{m(3,5)}$, we generate predictions for utterances 3 and 5, denoted \hat{u}_3 and \hat{u}_5 . The inpainted dialog is then:

$$\text{Inpaint}(d_{m(3,5)}) = (u_1, u_2, \hat{u}_3, u_4, \hat{u}_5).$$

2.1. Training: Dialog reconstruction

We train a *dialog inpainter* using the following dialog reconstruction task: Given a complete dialog, $d = (u_1, u_2, \dots, u_T)$, we randomly mask one utterance, u_t , yielding a partial dialog:

$$d_{m(t)} = (u_1, \dots, u_{t-1}, \diamond, u_{t+1}, \dots, u_T).$$

Given this partial dialog, we train our model to predict u_t , the original value of the masked utterance. This is similar to the masked language modeling task used by BERT (Devlin et al., 2019), but instead of reconstructing missing tokens in a passage, we reconstruct a missing utterance in a dialog.

More precisely, let us assume that the dialog inpainter is a generative model with parameters θ specifying a probability distribution $p_\theta(u_t | d_{m(t)})$. Then, our training objective is to minimize the following loss:

$$\mathcal{L}(\theta) = - \sum_{d \in \mathcal{D}} \mathbb{E}_{u_t \sim d} [\log p_\theta(u_t | d_{m(t)})], \quad (1)$$

where \mathcal{D} is a corpus of complete dialogs and u_t is a randomly sampled utterance from the dialog d .

We implement our inpainter using the T5 (Raffel et al., 2020) text-to-text encoder-decoder Transformer (Vaswani et al., 2017), where the input x and output y are represented as text strings. We convert each dialog d into an (x, y) pair as follows: First, we randomly sample a turn t to mask, yielding partial dialog $d_{m(t)}$ and target output u_t . To form

¹We released WikiDialog at <https://github.com/google-research/dialog-inpainting>

input string x , we prepend each utterance in $d_{m(t)}$ with its corresponding speaker id (e.g., 0 or 1 in a two-speaker dialog), and simply concatenate them together, for example, “0: u_1 1: u_2 0: \diamond 1: u_4 ”. The output string y is simply the masked utterance u_t . We optimize this model using a standard cross-entropy loss which is equivalent to Eq. 1.

2.2. Inference: Transforming documents into dialogs

We now show how to use a trained inpainter to transform a document into a dialog. Suppose we have a document or passage p consisting of sentences (s_1, s_2, \dots, s_m) . Imagine that each of these sentences was an utterance spoken by the writer in an imagined dialog with the reader: we want to know what the imagined reader said between each of the writer’s utterances. We can pose this question to our inpainter by asking it to complete the following partial dialog: $(\diamond, s_1, \diamond, s_2, \diamond, \dots, \diamond, s_m)$ —each utterance from the imagined reader starts masked and is responded to by the writer with a sentence from the passage. While the model knows the masked utterances come from another speaker, it still lacks any hint that the speaker should ask questions. To provide this hint, we prepend an utterance from the writer s_{prompt} that says: “Hello, I am an automated assistant and can answer questions about (document title)”. Hence, the overall partial dialog has the form:

$$\text{PartialDialog}(p) = (s_{\text{prompt}}, \diamond, s_1, \diamond, \dots, \diamond, s_m).$$

Note, however, that this partial dialog contains multiple masked utterances, while our inpainter is only trained to inpaint a single utterance at a time. To handle this, we use the model autoregressively: we begin by providing $(s_{\text{prompt}}, \diamond, s_1)$ as input to the inpainter and generate \hat{u}_1 via greedy decoding. We then replace the first mask with \hat{u}_1 and use a new input $(s_{\text{prompt}}, \hat{u}_1, s_1, \diamond, s_2)$ to generate \hat{u}_2 and so on until all masks are filled and the dialog is complete.

2.3. Case study: Applying dialog inpainting to generate an information seeking dialog dataset

We now have all the ingredients necessary to generate synthetic information-seeking dialogs at scale. In this section, we will outline how we trained inpainters and applied them on two document corpora to respectively generate two such datasets: WikiDialog and WebDialog.

Inpainter model training. We used PublicDialog, TaskMaster, OR-QuAC, and QReCC (Table 2) to train three different inpainters: Inpaint_{PT}, Inpaint_{OQ} and Inpaint_{PTOQ}, where the subscripts indicate which of the above datasets were used. We chose these subsets to study how the type and quantity of training data affect inpainter quality: Inpaint_{PT} only uses open-domain conversational datasets that do not contain any explicit question answering, while Inpaint_{OQ} only uses relatively small conversational question answer-

ing datasets, and Inpaint_{PTOQ} uses all of the above. Each inpainter model was initialized from a pre-trained T5-XXL (11B parameters) checkpoint and finetuned on its corresponding training data. See Appendix A.2 for details on the hyperparameters used.

Document processing. We apply dialog inpainting to two document corpora: WIKI, a collection of 11.4M passages from 5.9M English Wikipedia articles in the OR-QuAC retrieval corpus (Qu et al., 2020), and WEB, a collection of 8.4M English web passages from the MS Marco retrieval corpus (Nguyen et al., 2016). Both corpora were used as is without any further filtering. The passages were split into sentences using the Google Cloud Natural Language API. To limit computation, we only use the first 6 sentences of each passage.² The passages were then converted to partial dialogs and inpainted using the method described in Section 2.2. The final results are the datasets WikiDialog_{PT}, WikiDialog_{OQ} and WikiDialog_{PTOQ} and WebDialog_{PT}, where the subscripts indicate which inpainter model was used to generate the data (Table 2).

3. Evaluating WikiDialog as a Dataset

We now turn to a qualitative evaluation of the dialogs generated in Section 2. We will show that they largely contain *information-seeking* dialogs with well-matched questions and answers. This makes the data suitable for ConvQA systems, an application explored later in the paper. As WikiDialog is generated using *passages* from OR-QuAC, the corresponding OR-QuAC *dialogs* form a natural point of reference: both datasets rely on the same content, but are respectively automatically and manually generated. For this reason, our analysis focuses on the WikiDialog variants.

Our analysis combines automatic and human evaluation. The human evaluation was conducted by asking human raters the subjective questions listed in Table 1 for each turn of a dialog.³ We found substantial inter-annotator agreement on all four questions, with a Krippendorff’s α of at least 0.89, and report results aggregated over dialogs corresponding to a common set of 200 randomly chosen passages using the mean score of three raters.

How information seeking are the generated utterances?

Raters judge the generated utterances to almost always be information seeking and topically relevant. We note that, prompt aside, our method is not tailored to explicitly generate information seeking utterances in any way. It is remarkable then that raters found 94.5% of utterances in WikiDialog_{PT}, which was generated without using any

²We truncated about 64% of passages containing 7+ sentences.

³Additional details, including screenshots of the annotation interface, instructions and examples, are provided in Appendix B.

Question (α)	WikiDialog			OR-QuAC
	PT	OQ	PTOQ	
Is the question information seeking? (0.98)				
Yes	<u>94.5%</u>	100%	99.3%	100%
How relevant is question to the conversation? (0.94)				
Not at all	0.3%	0%	0%	0%
Topic only	<u>45.8%</u>	49.5%	42.1%	52.6%
Follows up	<u>53.9%</u>	50.5%	<u>57.9%</u>	47.4%
How specific is the question? (0.91)				
Not at all	6	<u>12%</u>	5.8%	5.4%
Somewhat	15%	<u>28.7%</u>	22.4%	12%
Very	79%	<u>59.3%</u>	<u>71.7%</u>	82.6%
How well answered is the question? (0.89)				
Not at all	0.1%	0.1%	0.0%	0.8%
Incompletely	<u>15.6%</u>	19.7%	25.4%	22.8%
Sufficiently	<u>52.4%</u>	46.8%	<u>50.2%</u>	36.6%
Perfectly	<u>31.9%</u>	33.4%	<u>24.4%</u>	39.8%

Table 1. Results from a human evaluation of the generated utterances in three WikiDialog variants vs. a dataset manually collected on the same passages, OR-QuAC. The table also presents inter-annotator agreement using Krippendorff’s α . Underlined numbers statistically differ from corresponding OR-QuAC ones at a $p < 0.05$ level based on a paired randomization test.

ConvQA data, to be information seeking; by training the inpainter using ConvQA data, this number goes to 99–100%. Almost all of these are single sentences phrased as questions, though WikiDialog_{PT} includes some exceptions, e.g., “Exciting! I wonder if they got a #1 debut with that song.”.

What types of questions are generated? Next, we study the distribution of questions generated by clustering utterances on their first two words. Figure 2 shows how this distribution changes over turns: dialogs start with more definitional questions (e.g., what is, who is, where is, etc.) but then diversify into a range of follow-up questions (what happened, did, is, how, why, etc.). A good conversation should have a mix of both specific and broad questions. Here, raters find that the WikiDialog variants primarily differ from OR-QuAC in having more “somewhat specific” questions (e.g., “Is there anything else interesting about his guitar?”) and fewer “very specific” questions (e.g., “Did people enjoy their music?”).

How well answered are the generated questions? So far, we have shown that one speaker in WikiDialog consistently asks questions, playing the role of the imagined reader well. We now turn to the other speaker, the writer: how often are their responses—sentences from a Wikipedia passage—adequate answers to the questions? Raters find that questions in all three variants of WikiDialog are suffi-

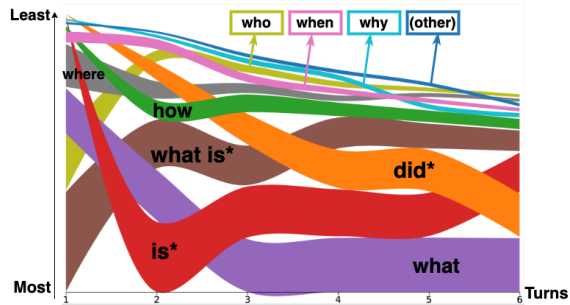


Figure 2. Bump chart showing how questions are distributed across turns in WikiDialog_{PTOQ}: the width and ordering of each line is proportional to the fraction of questions of that type in each turn. Dialogs tend to start with definitional questions (what is, who, where, etc.) and diversify into a wide range of follow questions (what happened, did, is, etc).

ciently answered about as often as in OR-QuAC,⁴ though questions in OR-QuAC are often answered better.

How conversational are the data? A key distinguishing feature of dialog data is its dependence on prior dialog context. Following prior work (Choi et al., 2018), we use the degree of context dependence as a measure of conversation. Raters judge that questions in the WikiDialog variants follow-up on dialog context significantly more often than those in OR-QuAC. We hypothesize that the tendency to follow-up likely arises from the structure of the underlying passages from which WikiDialog is derived: when a paragraph in Wikipedia explores a single topic in depth, so too does the dialog generated from it in WikiDialog.

Do the data contain potentially sensitive language that may perpetuate unfair bias? Evaluating a dataset of this size for potential unfair bias is a daunting task that requires diligent and thorough investigation. We approached the problem by curating a set of 700 terms from the literature (Bolukbasi et al., 2016; Garg et al., 2018; May et al., 2019; Nadeem et al., 2020; Abid et al., 2021) related to sensitive characteristics—such as race, ethnicity, gender, and sexual orientation. Many instances of these terms are well-motivated: for example, a dialog from a passage about transgender rights in Canada includes the question “What does anti-discrimination act mean in relation to transgender people?”. We further refined the approach to instead look at co-occurrences between these terms and adjectives that may have negative connotations, focusing on instances where the terms were not explicitly mentioned in the passage. We find that 0.2–0.5% of dialogs in the dataset contain such potentially sensitive interactions, but it is difficult to establish if

⁴When considering the net percentage of questions sufficiently or perfectly answered, the differences between the WikiDialog variants and OR-QuAC are not statistically significant at the $p < 0.05$ level.

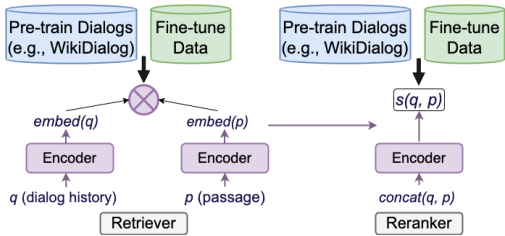


Figure 3. We use a two-stage ConvQA retrieval system. We first retrieve top-K passages from the corpus using a dual-encoder model and then rerank them using a cross-attention model.

they perpetuate unfair bias without expert manual review. Therefore, we advise users to note these observations and exercise care while using the dataset.⁵

4. Application: Open-domain Conversational Retrieval

In this section, we show how our inpainted datasets can serve as a powerful source of training data for ConvQA systems.

A ConvQA system engages with a user through multi-turn dialog, where typically the user poses questions and the system answers (there can be exceptions, e.g., the system asks a clarifying question). During a dialog, whenever it is the system’s turn to speak (at some time t), it looks at all previous dialog turns $d_{1:t} = (u_1, u_2, \dots, u_t)$ which we call the *dialog history*, and outputs a new utterance, u_{t+1} .

Because ConvQA dialogs are knowledge-intensive, many systems decompose the task into a two-part retrieve-then-generate process (Qu et al., 2020; Anantha et al., 2021). First, they employ a *conversational retriever* to retrieve passages that are relevant to the conversation based on the dialog history $d_{1:t}$. Second, they employ a *generator* which uses both the dialog history ($d_{1:t}$) and the retrieved passages to generate a response, u_{t+1} . While both steps are important, the *conversational retriever* is key to helping the model access the right knowledge and also for showing people evidence for an answer.

This work focuses on the conversational retriever, showing how to improve it by pre-training on our inpainted data, leaving improvements to the generator for future work.

Models. The input to a conversational retriever is the dialog history ($d_{1:t}$) and a passage (p). The output is a score, $s(d_{1:t}, p)$, indicating the passage’s relevance. Retrieval is performed by selecting the passages with the highest scores.

We also refer to the dialog history as the “query” and denote it as q . In some benchmarks that we study, the “dialog his-

⁵Further details of our approach and some examples are provided in Appendix H.

tory” is defined to be all previous utterances, while in others the history is defined to only include the user’s questions but not the system’s responses.

We employ two standard models for retrieval: first, we use a *dual encoder* (Reimers & Gurevych, 2019; Karpukhin et al., 2020; Ni et al., 2021) to select an initial set of candidates. We then rescore those candidates using a cross-attention *reranker* (Nogueira & Cho, 2019; Nogueira et al., 2020). Model architectures and training objectives are detailed in Appendix C. To train these models, we need a corpus of (q, p^*) pairs, where q is a dialog history and p^* is a relevant passage. The following sections describe such data. Figure 3 illustrates our system.

Pre-training and Fine-tuning. As outlined in Section 2.2, each dialog generated by our inpainter tends to consist of alternating question and answer utterances: $d = (s_{\text{prompt}}, \hat{u}_1, s_1, \dots, \hat{u}_m, s_m)$, where inpainted utterances \hat{u}_i are questions, and their subsequent answers s_i are sentences from the original passage p . Intuitively, for each question in the dialog, p is a highly relevant passage that should be retrieved. Based on this observation, we generate examples as follows: first, we randomly select a dialog prefix that ends in a question to be the dialog history: $q_i = (\hat{u}_1, s_1, \dots, \hat{u}_i)$.⁶ We then wish to mark the original passage p as a positive passage to retrieve. However, directly using p as a positive example will not yield good results: the dialog history (q_i) includes exact sentences from p , which would cause our retriever to simply learn to string-match, rather than to generalize. To eliminate this problem, we form a new passage that consists only of the remaining sentences in p that haven’t appeared in q_i yet: $p_i^* \stackrel{\text{def}}{=} \text{Concat}(s_j \text{ where } j > i)$.

After pre-training (q_i, p_i^*) pairs from the inpainted data, our retriever is fine-tuned on a downstream ConvQA dataset.

5. Evaluation

We report quantitative evaluation of dialog inpainting by measuring the impact of WikiDialog and WebDialog on ConvQA retrieval systems.

5.1. Experimental setup

The following presents a summary of our experimental setup; see Appendix D for full details.

Datasets. We use three open-domain conversational QA retrieval benchmarks: OR-QuAC (Qu et al., 2020), QReCC (Anantha et al., 2021), and TREC CAsT19 and CAsT20 (Dalton et al., 2019; 2020). Table 2 summarizes basic statistics of the each dataset; further details are pro-

⁶We omit the leading prompt utterance (s_{prompt}) from the history. Also, we optionally omit answer turns, depending on the task setup.

Dataset	Dialogs	Turns	Ques.
<i>Non-conversational question answering datasets</i>			
MS Marco			0.5M
PAQ			65M
<i>Non-question answering conversational datasets</i>			
PublicDialog	2.0M	1:2:4	
TaskMaster	19.8K	9:10:14	
<i>Conversational question answering datasets</i>			
OR-QuAC	5.6K	4:8:12	40.5K
QReCC	13.6K	1:6:12	80.0K
CAsT-19	80	7:10:12	748
CAsT-20	25	6:8:13	216
WikiDialog	11.4M	1:6:6	56.1M
WebDialog	8.4M	1:3:6	26.9M

Table 2. The number of dialogs, turns and questions (Ques.) of datasets considered in this paper; turns and question lengths are reported using (1%:50%:99%) percentiles.

vided in Appendix D.3.

Retrieval corpora. We use standard retrieval corpora and relevance labels from each dataset for evaluation evaluation: OR-QuAC uses 11.1M passages from English Wikipedia. TREC CAsT-19 and CAsT-20 use a shared corpus of 8M web passages from MS Marco (Nguyen et al., 2016) and 30M Wikipedia passages from Dusart et al. (2019); QReCC uses 54M passages from the Internet Archive Wayback Machine and the English pages in CommonCrawl. Following prior work (Qu et al., 2020; Yu et al., 2021; Wu et al., 2021), we use only the current and previous questions as inputs for OR-QuAC, TREC CAsT-19 and CAsT-20, and use previous answers in addition to these as input for QReCC. We report mean reciprocal rank at rank 5 (MRR@5) for OR-QuAC and mean reciprocal rank at full rank (MRR) for QReCC and TREC CAsT. Additional metrics are reported in Appendix F.

Model implementation. We initialize our dual-encoder retrievers and rerankers from pre-trained T5 checkpoints following prior work (Ni et al., 2021; Nogueira et al., 2020). Both retrievers and rerankers are pre-trained on our inpainted datasets. We use the notation T5-Large DE \triangleright WikiD to denote a dual-encoder that was initialized from a T5-Large checkpoint and pre-trained on WikiDialog (we use WikiDialog_{PTOQ} unless otherwise specified), and T5-Large DE \triangleright WikiD+WebD to denote a model pre-trained on the union of WikiDialog and WebDialog. Similar notation is used for T5-Base initialized models and rerankers.

During fine-tuning, we separately train retrievers and rerankers on OR-QuAC and QReCC, using their validation sets to select checkpoints. Because CAsT19 and CAsT20 are extremely small datasets and do not include a training split, we do not fine-tune *dual-encoder retrievers* on these

datasets, instead using a retriever finetuned on QReCC data⁷. We follow Yu et al. (2021) and use 5-fold cross-validation to finetune *rerankers* on CAsT19 and CAsT20: for each fold, we split the data into 5 splits based on dialogs, train a reranker on 3 splits of the data, select a checkpoint on one split and test on the remaining split. Similar to Yu et al. (2021) and Lin et al. (2021), we use hard negatives mined from a base retriever during fine-tuning (See Appendix C.3 for details). Finally, during inference, we use MIPS (Ram & Gray, 2012) to index the passage embeddings, and use exhaustive search to retrieve and then rerank the top 100 passages for all datasets.

Baselines. We compare our pre-trained dual-encoder retrievers against five published *retrieval-only* baselines: BM25-Query Rewriter (Yu et al., 2020), BM25-T5QR (Wu et al., 2021), ANCE-Query Rewriter (Yu et al., 2021), CONQRR (Wu et al., 2021) and ConvDR (Yu et al., 2021).⁸ We also evaluate our base model, a T5-Large DE model that is only finetuned on each dataset: this is a strong baseline that outperforms previous state-of-the-art on QReCC. We also compare our pre-trained *retrieval+reranking* systems against the best performing systems in TREC CAsT⁹: CFDA_CLIP_RUN7 (Yang et al., 2019) and h2oloo_RUN4 (Dalton et al., 2020), and a state-of-the-art retriever+reranker method ConvDR \rightarrow BERT (Yu et al., 2021).

5.2. Main Results

Table 3 compares our models with baselines on four ConvQA retrieval benchmarks. We first note that our relatively simple base retriever model, T5-Large DE, is a strong baseline and significantly outperforms the existing state-of-the-art results on QReCC. By simply pre-training on WikiDialog, we observe a 9–30% gain, and outperform all baselines on QReCC and OR-QuAC with large margins. Including WebDialog increases this gain by a further 3–15% MRR, with the most significant gains accruing to CAsT-19 and CAsT-20. With this model, we observe a 43% relative MRR gain over the current state-of-the-art on QReCC, a 12% relative MRR@5 gain on OR-QuAC, and comparable or better performance on CAsT-19 and CAsT-20.

The gains achieved in the retrieval stage remain when the results are reranked: we outperformed the current state-of-the-art on OR-QuAC and TREC CAsT-19, and achieved comparable performance on CAsT-20. We could not find existing published reranker results for QReCC, but observe

⁷The QReCC training data does not include TREC CAsT.

⁸Appendix D.3 includes brief summaries of these methods.

⁹TREC CAsT has several settings based on if a model uses additional inputs such as manual query rewrites. We follow the *automatic* setting that does not use additional inputs, and compare against top *automatic* runs reported in Dalton et al. (2019; 2020).

System	QReCC		TREC CAsT	
	MRR	MRR@5	19	20
<i>Retrieval</i>				
BM25-QR	–	20.2	58.1	25.0
ANCE-QR	–	45.7	66.5	37.5
ConvDR	–	61.6	74.0	50.1
BM25-T5QR	32.8	–	–	–
CONQRR	41.8	–	–	–
T5-Large DE	55.7	56.9	61.0	34.3
▷ WikiD	60.4	66.5	68.1	43.7
▷ WikiD+WebD	60.7	68.7	74.1	51.3
<i>Retrieval + Reranking</i>				
CFDA_CLIP_RUN7	–	–	71.4	–
h2olo0_RUN4	–	–	–	59.3
ConvDR→BERT	–	77.3	79.9	54.5
T5-Large DE	–	–	–	–
→ reranker	68.9	72.6	75.3	55.1
▷ WikiD	70.7	79.7	79.3	60.3
▷ WikiD+WebD	71.8	81.2	82.0	59.7

Table 3. We evaluate models trained using our inpainted datasets against baselines on three different ConvQA retrieval tasks: QReCC, OR-QuAC, and TREC CAsT. WikiD, WikiD+WebD: We use WikiDialog_{PTOQ} for all tasks except when WikiD+WebD is used in TREC CAsT, where WikiDialog_{PT} gives higher performance. Additional metrics are reported in Appendix F.

a 5% relative MRR gain over reranking the non pre-trained T5-Large DE model, which we found to be a strong baseline.

5.3. Analysis

Next, we dive into our results in greater detail and study factors contributing to performance:

How much does in-domain training matter for inpainter quality? Table 4 compares the performance of retrievers pre-trained on three variants of WikiDialog introduced in Section 3: PT, OQ, and PTOQ. Remarkably, using WikiDialog_{PT} is sufficient to significantly outperform current state-of-the-art methods, despite being generated using only open-domain conversational data. Next, we observe that using an inpainter trained on OR-QuAC and QReCC data (WikiDialog_{OQ}) results in slightly better performance; it is notable that an inpainter trained using only 20K dialogs is able to generate such high-quality data, and can improve performance over a retriever just fine-tuned on the same data by up to 20%. Finally, we observed that an inpainter trained on both types of data (WikiDialog_{PTOQ}) is able to generalize better to CAsT-19, an out-of-domain task we evaluate on using a QReCC retriever.

How does WikiDialog compare to other retriever pre-training datasets? Table 4 also reports results for re-

System	QReCC	OR-QuAC	CAsT-19
	MRR	MRR@5	MRR
T5-Base DE	53.4	53.6	55.2
▷ MS Marco	55.0	57.5	57.5
▷ PAQ	57.0	62.7	48.4
▷ PT	56.8	59.5	60.1
▷ WikiD _{PT}	58.5	64.0	61.3
▷ WikiD _{OQ}	58.9	65.3	65.5
▷ WikiD _{PTOQ}	59.0	64.1	66.5

Table 4. Retriever performance when T5-Base DE is pre-trained (▷) on various datasets, then fine-tuned on QReCC and OR-QuAC; results on CAsT-19 use a retriever finetuned on QReCC.

trievers pre-trained on several alternatives to WikiDialog: MS Marco (Nguyen et al., 2016), a search dataset that is widely used for retriever pre-training; PAQ (Lewis et al., 2021), the largest existing (non-conversational) question-answering dataset, also automatically generated from Wikipedia; and PublicDialog + TaskMaster (PT), the open-domain dialog data introduced in Section 2 as a training dataset for the inpainter—we use it here to directly pre-train a conversational retriever using a next utterance retrieval task. Table 2 lists their characteristics.

We observe that pre-training on any of the WikiDialog variants significantly outperforms the two classic non-conversational QA datasets, MS Marco and PAQ. WikiDialog also outperforms open-domain dialog data PublicDialog+TaskMaster (PT). Despite being generated using PT, pre-training on WikiDialog_{PT} far outperforms it, showing that the proposed dialog inpainting recipe provides additional signal not present in PT.

In addition, we found that pre-training a retriever on PT does surprisingly well, significantly outperforming MS Marco on all datasets. Prior ConvQA retrieval systems have focused on traditional QA data for pre-training, but we find that open-domain conversational data can be just as useful!

Does pre-training on WikiDialog enable zero/few-shot learning? We now explore how much fine-tuning data is needed after pre-training on WikiDialog. Figure 4 (a) plots the retrieval results of a T5-Base retriever pre-trained on WikiDialog_{PT}—which does not use any in-domain data—and finetuned on varying percentages of the QReCC training dataset¹⁰.

Surprisingly, we observe that the zero-shot performance of the pre-trained retriever is already quite good: achieving nearly 95% the MRR of a model that uses the full dataset. Furthermore, pre-trained retriever needs fewer fine-tuning

¹⁰Unlike the other experiments presented in this paper that uses mined hard negatives to fine-tune retrievers, these results do not use hard negatives because our hard negative mining model was fine-tuned on the whole QReCC dataset (Appendix C).

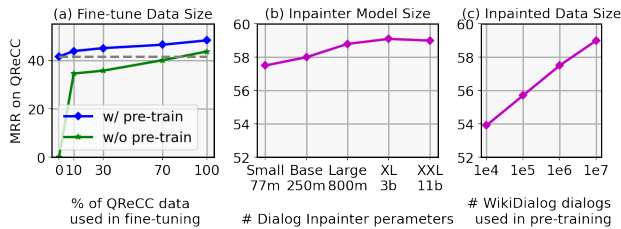


Figure 4. Retriever performance on QReCC when T5-Base DE \triangleright WikiDialog_{PT} is trained with (a) varying fine-tuning data sizes, (b) different sizes inpainter models, and (c) varying pre-training data sizes. Results in (a) do not include mined hard-negatives.

data – by fine-tuning on just 10% of the data, the pre-trained retriever exceeds the performance of the full-data baseline.

Does our method scale with inpainting model size and data size? We now explore if our dialog inpainting method can benefit from scaling up along two dimensions: the inpainter model size, and the inpainted WikiDialog data size. Results are shown in Figure 4 (b) and (c).

From Figure 4 (b), we observe that retriever performance increases with inpainter model size with one exception: the T5-XL model slightly outperforms T5-XXL; we hypothesize this is due to insufficient hyperparameter search for T5-XXL. Surprisingly, the quality of data generated by T5-Small is already sufficient to significantly outperform current state-of-the-art methods.

In Figure 4 (c), we evaluate how retrievers pre-trained with 10K–11M dialogs sampled from WikiDialog perform on QReCC. We observe a roughly log-linear relationship between performance and pre-training data size that has not yet plateaued: simply inpainting more passages may further increase retrieval performance.

6. Related Work

Conversational question answering retrieval. Several manually collected conversational question answering datasets have been proposed to address the scarcity of high-quality training data (Choi et al., 2018; Reddy et al., 2019; Dinan et al., 2018; Saeidi et al., 2018; Dalton et al., 2019; Campos et al., 2020; Dalton et al., 2020; Qu et al., 2020; Feng et al., 2020; Anantha et al., 2021). However, because they are relatively small, existing retrieval systems all depend on rewriting queries to use with a *non-conversational* retrieval system (Yang et al., 2019; Dalton et al., 2020; Yu et al., 2020; 2021; Wu et al., 2021). Query rewriting is a hard problem in itself (Vakulenko et al., 2020)—prior systems rely on proprietary search logs (Yu et al., 2020), reinforcement learning (Wu et al., 2021) or distillation recipes (Yu et al., 2021)—and some queries cannot be rewritten at all, e.g., “What *else* were they famous for?”. Here, we use

WikiDialog, a 1000x larger ConvQA dataset, to train a standard retriever without requiring query rewrites.

Conversational language models. Large conversational language models such as DialogGPT (Zhang et al., 2020) and Meena (Adiwardana et al., 2020) have shown impressive open-ended conversational capabilities, and even the ability to directly answer many questions. However, they often hallucinate answers and amplify unfair biases present in their training data. Subsequent work address this problem by instead posing queries to a *non-conversational* retriever and using its answers (Roller et al., 2021; Komeili et al., 2021; Nakano et al., 2021; Thoppilan et al., 2022; Elgohary et al., 2019). In this paper, we instead use a *masked* conversational language model to generate ConvQA data: the data can be readily audited for unfair biases and to train a *conversational* retrieval system.

Data augmentation and synthetic data generation. Data augmentation has been widely used to improve the performance of document retrieval and related tasks (Lee et al., 2019; Chang et al., 2020; Gao et al., 2021) by using training data more efficiently. However, these methods typically do not result in high-quality data that can be used outside of a training recipe. More recently, Ma et al. (2021) and Lewis et al. (2021) train special-purpose question generation models to synthetically generate millions of high-quality question answer pairs. We extend this line of work to generate millions *information seeking dialogs*, and show that even a general purpose inpainter model can generate high-quality data.

7. Discussion

In this paper, we have presented *dialog inpainting*, a novel approach to generating synthetic *conversational* data. We showed that it is possible to generate compelling information-seeking dialogs using only general-purpose data, suggesting applications to other conversational tasks. While synthetic data cannot entirely replace real data, it can help bootstrap interactive conversation systems and create a virtuous cycle wherein users find it natural to engage with and improve the system. We are particularly optimistic about applying the dialog inpainting data to (1) distillation, where the inpainted datasets serve as large-scale distillation sets, (2) end-to-end conversational question answering, and (3) *zero-shot* conversational QA, which is motivated by the zero-shot retrieval capabilities shown in this work.

It is important to be aware of the biases that generating data can introduce or amplify. We want to encourage good inductive biases that make conversations conversational—e.g., use of anaphora or elision of context—and to introduce further control over the dialogs generated—e.g., persona or dialog acts. At the same time, we must interrogate the

generated data and work towards minimizing instances of potentially sensitive language that may perpetuate unfair biases.

Acknowledgements

We thank Jianmo Ni and T5X Retrieval project¹¹ for helping us set up the dual-encoder training framework. We thank the LaMDA team (Thoppilan et al., 2022) for providing help with inpainter training. Finally, we thank Alex Salcianu, Victoria Fossum, Raphael Hoffmann, Kenton Lee, and Maithra Raghu for helpful discussions and feedback.

References

- Abid, A., Farooqi, M., and Zou, J. Large language models associate muslims with violence. *Nature Machine Intelligence*, 3(6):461–463, 2021.
- Adiwardana, D., Luong, M.-T., So, D. R., Hall, J., Fiedel, N., Thoppilan, R., Yang, Z., Kulshreshtha, A., Nemade, G., Lu, Y., et al. Towards a human-like open-domain chatbot. *arXiv preprint arXiv:2001.09977*, 2020.
- Anantha, R., Vakulenko, S., Tu, Z., Longpre, S., Pulman, S., and Chappidi, S. Open-domain question answering goes conversational via question rewriting. In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, NAACL-HLT*. Association for Computational Linguistics, 2021.
- Bolukbasi, T., Chang, K.-W., Zou, J. Y., Saligrama, V., and Kalai, A. T. Man is to computer programmer as woman is to homemaker? debiasing word embeddings. *Advances in neural information processing systems*, 29, 2016.
- Byrne, B., Krishnamoorthi, K., Sankar, C., Neelakantan, A., Goodrich, B., Duckworth, D., Yavuz, S., Dubey, A., Kim, K., and Cedilnik, A. Taskmaster-1: Toward a realistic and diverse dialog dataset. In Inui, K., Jiang, J., Ng, V., and Wan, X. (eds.), *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing, EMNLP-IJCNLP*. Association for Computational Linguistics, 2019.
- Campos, J. A., Otegi, A., Soroa, A., Deriu, J. M., Cieliebak, M., and Agirre, E. Doqa-accessing domain-specific faqs via conversational qa. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, 2020.
- Chang, W., Yu, F. X., Chang, Y., Yang, Y., and Kumar, S. Pre-training tasks for embedding-based large-scale retrieval. In *8th International Conference on Learning Representations, ICLR*, 2020.
- Choi, E., He, H., Iyyer, M., Yatskar, M., Yih, W., Choi, Y., Liang, P., and Zettlemoyer, L. Quac: Question answering in context. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing, EMNLP-IJCNLP*. Association for Computational Linguistics, 2018.
- Culpepper, J. S., Diaz, F., and Smucker, M. D. Research frontiers in information retrieval: Report from the third strategic workshop on information retrieval in lorne (SWIRL). *SIGIR Forum*, 2018.
- Dalton, J., Xiong, C., and Callan, J. TREC cast 2019: The conversational assistance track overview. In *Proceedings of the Twenty-Ninth Text REtrieval Conference, TREC*, NIST Special Publication. National Institute of Standards and Technology (NIST), 2019.
- Dalton, J., Xiong, C., and Callan, J. Cast 2020: The conversational assistance track overview. In *Proceedings of the Twenty-Ninth Text REtrieval Conference, TREC*, NIST Special Publication. National Institute of Standards and Technology (NIST), 2020.
- Devlin, J., Chang, M., Lee, K., and Toutanova, K. BERT: pre-training of deep bidirectional transformers for language understanding. In Burstein, J., Doran, C., and Solorio, T. (eds.), *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, NAACL-HLT*. Association for Computational Linguistics, 2019.
- Dinan, E., Roller, S., Shuster, K., Fan, A., Auli, M., and Weston, J. Wizard of wikipedia: Knowledge-powered conversational agents. *arXiv preprint arXiv:1811.01241*, 2018.
- Dusart, A., Hubert, G., and Pinel-Sauvagnat, K. IRIT at TREC 2019: Incident streams and complex answer retrieval tracks. In *Proceedings of the Twenty-Eighth Text REtrieval Conference, TREC*, NIST Special Publication. National Institute of Standards and Technology (NIST), 2019.
- Elgohary, A., Peskov, D., and Boyd-Graber, J. L. Can you unpack that? learning to rewrite questions-in-context. In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing, EMNLP-IJCNLP*. Association for Computational Linguistics, 2019.

¹¹https://github.com/google-research/t5x_retrieval

- Feng, S., Wan, H., Gunasekara, C., Patel, S., Joshi, S., and Lastras, L. Doc2dial: A goal-oriented document-grounded dialogue dataset. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, 2020.
- Gao, T., Yao, X., and Chen, D. Simcse: Simple contrastive learning of sentence embeddings. In Moens, M., Huang, X., Specia, L., and Yih, S. W. (eds.), *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing, EMNLP*. Association for Computational Linguistics, 2021.
- Garg, N., Schiebinger, L., Jurafsky, D., and Zou, J. Word embeddings quantify 100 years of gender and ethnic stereotypes. *Proceedings of the National Academy of Sciences*, 115(16):E3635–E3644, 2018.
- Iizuka, S., Simo-Serra, E., and Ishikawa, H. Globally and locally consistent image completion. *ACM Trans. Graph.*, 36(4):107:1–107:14, 2017.
- Karpukhin, V., Oguz, B., Min, S., Lewis, P. S. H., Wu, L., Edunov, S., Chen, D., and Yih, W. Dense passage retrieval for open-domain question answering. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing, EMNLP*. Association for Computational Linguistics, 2020.
- Komeili, M., Shuster, K., and Weston, J. Internet-augmented dialogue generation. *arXiv preprint arXiv:2107.07566*, 2021.
- Kwiatkowski, T., Palomaki, J., Redfield, O., Collins, M., Parikh, A. P., Alberti, C., Epstein, D., Polosukhin, I., Devlin, J., Lee, K., Toutanova, K., Jones, L., Kelcey, M., Chang, M., Dai, A. M., Uszkoreit, J., Le, Q., and Petrov, S. Natural questions: a benchmark for question answering research. *Trans. Assoc. Comput. Linguistics*, 2019.
- Lee, K., Chang, M., and Toutanova, K. Latent retrieval for weakly supervised open domain question answering. In Korhonen, A., Traum, D. R., and Màrquez, L. (eds.), *Proceedings of the 57th Conference of the Association for Computational Linguistics, ACL*. Association for Computational Linguistics, 2019.
- Lewis, P., Wu, Y., Liu, L., Minervini, P., Küttler, H., Piktus, A., Stenetorp, P., and Riedel, S. PAQ: 65 million probably-asked questions and what you can do with them. *arXiv preprint arXiv:2102.07033*, 2021.
- Li, H., Gao, T., Goenka, M., and Chen, D. Ditch the gold standard: Re-evaluating conversational question answering. *arXiv preprint arXiv:2112.08812*, 2021.
- Lin, S.-C., Yang, J.-H., and Lin, J. In-batch negatives for knowledge distillation with tightly-coupled teachers for dense retrieval. In *Proceedings of the 6th Workshop on Representation Learning for NLP (RepL4NLP)*, Online, August 2021. Association for Computational Linguistics.
- Liu, G., Reda, F. A., Shih, K. J., Wang, T., Tao, A., and Catanzaro, B. Image inpainting for irregular holes using partial convolutions. In Ferrari, V., Hebert, M., Sminchisescu, C., and Weiss, Y. (eds.), *Computer Vision - ECCV 2018 - 15th European Conference, Munich, Germany, September 8-14, 2018, Proceedings, Part XI*, Lecture Notes in Computer Science. Springer, 2018.
- Ma, J., Korotkov, I., Yang, Y., Hall, K. B., and McDonald, R. T. Zero-shot neural passage retrieval via domain-targeted synthetic question generation. In Merlo, P., Tiedemann, J., and Tsarfaty, R. (eds.), *Proceedings of the 16th Conference of the European Chapter of the Association for Computational Linguistics: Main Volume, EACL*. Association for Computational Linguistics, 2021.
- May, C., Wang, A., Bordia, S., Bowman, S., and Rudinger, R. On measuring social biases in sentence encoders. In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, NAACL*, 2019.
- Nadeem, M., Bethke, A., and Reddy, S. Stereoset: Measuring stereotypical bias in pretrained language models. *arXiv preprint arXiv:2004.09456*, 2020.
- Nakano, R., Hilton, J., Balaji, S., Wu, J., Ouyang, L., Kim, C., Hesse, C., Jain, S., Kosaraju, V., Saunders, W., et al. WebGPT: Browser-assisted question-answering with human feedback. *arXiv preprint arXiv:2112.09332*, 2021.
- Nguyen, T., Rosenberg, M., Song, X., Gao, J., Tiwary, S., Majumder, R., and Deng, L. MS MARCO: A human generated machine reading comprehension dataset. In *Proceedings of the Workshop on Cognitive Computation: Integrating neural and symbolic approaches 2016 co-located with the 30th Annual Conference on Neural Information Processing Systems, NIPS*, 2016.
- Ni, J., Ábrego, G. H., Constant, N., Ma, J., Hall, K. B., Cer, D., and Yang, Y. Sentence-t5: Scalable sentence encoders from pre-trained text-to-text models. *arXiv preprint arXiv:2108.08877*, 2021.
- Nogueira, R. and Cho, K. Passage re-ranking with bert. *arXiv preprint arXiv:1901.04085*, 2019.
- Nogueira, R., Jiang, Z., Pradeep, R., and Lin, J. Document ranking with a pretrained sequence-to-sequence model. In *Findings of the Association for Computational Linguistics: EMNLP*. Association for Computational Linguistics, 2020.

- Qu, C., Yang, L., Chen, C., Qiu, M., Croft, W. B., and Iyyer, M. Open-retrieval conversational question answering. In *Proceedings of the 43rd International ACM SIGIR conference on research and development in Information Retrieval, SIGIR*. ACM, 2020.
- Qu, Y., Ding, Y., Liu, J., Liu, K., Ren, R., Zhao, W. X., Dong, D., Wu, H., and Wang, H. RocketQA: An optimized training approach to dense passage retrieval for open-domain question answering. In Toutanova, K., Rumshisky, A., Zettlemoyer, L., Hakkani-Tür, D., Beltagy, I., Bethard, S., Cotterell, R., Chakraborty, T., and Zhou, Y. (eds.), *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, NAACL-HLT*. Association for Computational Linguistics, 2021.
- Radlinski, F. and Craswell, N. A theoretical framework for conversational search. In *Proceedings of the 2017 conference on conference human information interaction and retrieval*, 2017.
- Raffel, C., Shazeer, N., Roberts, A., Lee, K., Narang, S., Matena, M., Zhou, Y., Li, W., and Liu, P. J. Exploring the limits of transfer learning with a unified text-to-text transformer. *J. Mach. Learn. Res.*, 2020.
- Ram, P. and Gray, A. G. Maximum inner-product search using cone trees. In Yang, Q., Agarwal, D., and Pei, J. (eds.), *The 18th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, KDD '12, Beijing, China, August 12-16, 2012*, pp. 931–939. ACM, 2012.
- Reddy, S., Chen, D., and Manning, C. D. Coqa: A conversational question answering challenge. *Transactions of the Association for Computational Linguistics*, 7:249–266, 2019.
- Reimers, N. and Gurevych, I. Sentence-bert: Sentence embeddings using siamese bert-networks. In Inui, K., Jiang, J., Ng, V., and Wan, X. (eds.), *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing, EMNLP-IJCNLP*. Association for Computational Linguistics, 2019.
- Roller, S., Dinan, E., Goyal, N., Ju, D., Williamson, M., Liu, Y., Xu, J., Ott, M., Smith, E. M., Boureau, Y.-L., et al. Recipes for building an open-domain chatbot. In *Proceedings of the 16th Conference of the European Chapter of the Association for Computational Linguistics: Main Volume*, 2021.
- Saeidi, M., Bartolo, M., Lewis, P., Singh, S., Rocktäschel, T., Sheldon, M., Bouchard, G., and Riedel, S. Interpretation of natural language rules in conversational machine reading. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pp. 2087–2097, 2018.
- Santhanam, K., Khattab, O., Saad-Falcon, J., Potts, C., and Zaharia, M. Colbertv2: Effective and efficient retrieval via lightweight late interaction. *arXiv preprint arXiv:2112.01488*, 2021.
- Stede, M. and Schlangen, D. Information-seeking chat: Dialogues driven by topic-structure. In *Proceedings of Catalog (the 8th workshop on the semantics and pragmatics of dialogue; SemDial04)*. Citeseer, 2004.
- Thoppilan, R., Freitas, D. D., Hall, J., Shazeer, N., Kulshreshtha, A., Cheng, H.-T., Jin, A., Bos, T., Baker, L., Du, Y., Li, Y., Lee, H., Zheng, H. S., Ghafouri, A., Mene-gali, M., Huang, Y., Krikun, M., Lepikhin, D., Qin, J., Chen, D., Xu, Y., Chen, Z., Roberts, A., Bosma, M., Zhou, Y., Chang, C.-C., Krivokon, I., Rusch, W., Pickett, M., Meier-Hellstern, K., Morris, M. R., Doshi, T., Santos, R. D., Duke, T., Soraker, J., Zevenbergen, B., Prabhakaran, V., Diaz, M., Hutchinson, B., Olson, K., Molina, A., Hoffman-John, E., Lee, J., Aroyo, L., Rajakumar, R., Butryna, A., Lamm, M., Kuzmina, V., Fenton, J., Cohen, A., Bernstein, R., Kurzweil, R., Aguera-Arcas, B., Cui, C., Croak, M., Chi, E., and Le, Q. Lamda: Language models for dialog applications, 2022.
- Vakulenko, S., Longpre, S., Tu, Z., and Anantha, R. A wrong answer or a wrong question? an intricate relationship between question reformulation and answer selection in conversational question answering. In *Proceedings of the 5th International Workshop on Search-Oriented Conversational AI (SCAI)*, 2020.
- Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A. N., Kaiser, Ł., and Polosukhin, I. Attention is all you need. In *Advances in neural information processing systems*, pp. 5998–6008, 2017.
- Wu, Z., Luan, Y., Rashkin, H., Reitter, D., and Tomar, G. S. CONQRR: Conversational query rewriting for retrieval with reinforcement learning. *arXiv preprint arXiv:2112.08558*, 2021.
- Xiong, L., Xiong, C., Li, Y., Tang, K., Liu, J., Bennett, P. N., Ahmed, J., and Overwijk, A. Approximate nearest neighbor negative contrastive learning for dense text retrieval. In *9th International Conference on Learning Representations, ICLR*, 2021.
- Yang, J., Lin, S., Wang, C., Lin, J., and Tsai, M. Query and answer expansion from conversation history. In Voorhees, E. M. and Ellis, A. (eds.), *Proceedings of the Twenty-Eighth Text REtrieval Conference, NIST Special Publication*. National Institute of Standards and Technology (NIST), 2019.

Yu, J., Lin, Z., Yang, J., Shen, X., Lu, X., and Huang, T. S. Generative image inpainting with contextual attention. In *2018 IEEE Conference on Computer Vision and Pattern Recognition, CVPR*. Computer Vision Foundation / IEEE Computer Society, 2018.

Yu, S., Liu, J., Yang, J., Xiong, C., Bennett, P. N., Gao, J., and Liu, Z. Few-shot generative conversational query rewriting. In Huang, J., Chang, Y., Cheng, X., Kamps, J., Murdock, V., Wen, J., and Liu, Y. (eds.), *Proceedings of the 43rd International ACM SIGIR conference on research and development in Information Retrieval, SIGIR*. ACM, 2020.

Yu, S., Liu, Z., Xiong, C., Feng, T., and Liu, Z. Few-shot conversational dense retrieval. In *SIGIR '21: The 44th International ACM SIGIR Conference on Research and Development in Information Retrieval*. ACM, 2021.

Zhang, Y., Sun, S., Galley, M., Chen, Y., Brockett, C., Gao, X., Gao, J., Liu, J., and Dolan, B. DIALOGPT: Large-scale generative pre-training for conversational response generation. In Celikyilmaz, A. and Wen, T. (eds.), *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics: System Demonstrations, ACL*. Association for Computational Linguistics, 2020.

Zhu, F., Lei, W., Wang, C., Zheng, J., Poria, S., and Chua, T.-S. Retrieving and reading: A comprehensive survey on open-domain question answering. *arXiv preprint arXiv:2101.00774*, 2021.

A. Inpainting Models

A.1. Datasets for dialog inpainter training

1. **PublicDialog** is an open-domain dialog dataset proposed by LaMDA (Thoppilan et al., 2022). Each dialog is mined from public forums and then scored and filtered using LaMDA’s SSI and safety model. While the original dataset contains multi-speaker dialogs, we only use a subset consisting of two-speaker dialogs. Note that most of the dialogs in PublicDialog are short and have only two turns.
2. **TaskMaster**¹² (Byrne et al., 2019) is a crowd-sourced dataset of task-oriented dialogs between two speakers. We only use the subset of dialogs related to movie ticket and restaurant reservations. While TaskMaster is significantly smaller than PublicDialog, its dialogs have many more turns, which helps mitigate turn bias.

¹²There are three datasets, Taskmaster-1, Taskmaster-2, and Taskmaster-3; we use Taskmaster-1. Also see <https://github.com/google-research-datasets/Taskmaster>.

3. **OR-QuAC** (Qu et al., 2020) is a version of the conversational question answering dataset QuAC (Choi et al., 2018) that consists of real human-human QA dialogs about a given Wikipedia article. OR-QuAC extends QuAC to the open-domain retrieval setting by making the first question context-independent and pairing each answer with a relevant passage from Wikipedia.
4. **QReCC** (Anantha et al., 2021) is another conversational question answering dataset constructed using question sequences from QuAC, TREC Conversational Assistant Track, and NaturalQuestions (Choi et al., 2018; Dalton et al., 2019; Kwiatkowski et al., 2019).¹³ For each question sequence, annotators search a document corpus to find relevant passages and answers. Though the questions in QReCC overlap with those in OR-QuAC, their answers can be completely different.

A.2. Training details

Unless otherwise specified, all our dialog inpainters are initialized from T5-XXL (11B parameters)¹⁴ and finetuned using 64 TPU v3 chips¹⁵ with constant learning rate 0.01, dropout rate 0.1 and batch size 128. We trained the Inpaint_{PT} and Inpaint_{PTOQ} using 100k steps, and Inpaint_{OQ} using 10k steps because its training set is significantly smaller.

A.3. Inference details

To generate the inpainted datasets, we used <https://beam.apache.org> to parallelize our computation. On average, it took from 20ms (for T5-Small) to - 141ms (for T5-XXL) to inpaint each utterance in a dialog, and required between 100 TPUv3-hours (for T5-Small) and 1900 TPUv3-hours (for T5-XXL) to inpaint each WikiDialog variant, and would cost between \$240 and \$4560 using preemptible TPUs on Google Cloud.

B. Human Evaluation Protocol

In this section we overview the human evaluation protocol used to evaluate WikiDialog variants and OR-QuAC in Section 3; the results of this evaluation were presented in Table 1.

Task design and iteration. Figure 5 provides a screenshot of the annotation interface and a description of its features. The task was established as follows:

¹³The NaturalQuestions dataset does not originally contain question sequences, so the authors asked human annotators to come up with follow-up questions first.

¹⁴We use t5.1.1 checkpoints from https://github.com/google-research/text-to-text-transfer-transformer/blob/main/released_checkpoints.md

¹⁵<https://cloud.google.com/tpu/>

Conversational Question Quality

[? Instructions](#)
[🗨️ Ask question](#)

Round 2 of 6

For each highlighted round of the conversation show to the right, fill in the questionnaire below. Numbers like 9 are hotkeys.

- ▶ **How is the query relevant to the conversation?**
 - Follows up on a previous query or response.** It is difficult to correctly understand the query without reading the conversation history.
 - Only related to the topic of the conversation.** The query is typically similar to previous queries or responses, but can be understood without reading them.
 - Not relevant.** The query doesn't appear to be relevant to the topic or a previous query or response.
- ▶ **How specific is the query?**
 - Very.** Only a specific answer would satisfy the user.
 - Somewhat.** A variety of answers of a specific kind would satisfy the user.
 - Not at all.** Many typically different answers would satisfy the user.
- ▼ **How well does the response answer the query?**
 - Perfectly.** The response completely satisfies the user's information need. 1
 - Sufficiently.** The response mostly answers the user's information need, though some additional information could be provided. 2
 - Incompletely.** The response provides some information relevant to the user, but doesn't adequately answer the query. 3
 - Not at all.** The response does not provide any relevant information for user's query or is not intelligible. 4

< Previous
> Next
✔ Submit
🗨️ Comment

Topic: Olivia Dunham

USER

How does Torv feel about Fauxlivia?

SYSTEM

Comparing Olivia to Fauxlivia, Torv felt that the two had many different characteristics.

USER

How do they compare personality wise?

SYSTEM

ANSWERS QUESTION?

Anna Torv described them as "remarkably alike."

Figure 5. A screenshot of the annotation interface. On the right, raters see the conversation history and the highlight turn they must rate in the form on the left. When all questions for a given turn are completed, they are allowed to move forward to the next turn (round in the figure). The task can be submitted when all turns are complete.

Response Options (1 / 4)

Is the query **information-seeking**?

- ✔
 - **Yes.** The user is looking to learn some *information* from the system.
 - Note: Information-seeking queries don't have to be phrased as questions.
- ✘
 - **No.** The query is unclear, difficult to understand or not seeking information.
 - Note: Not all questions are information seeking, e.g. questions directed at the system (“how are you”, “what do you think”) or ones that are nonsensical in the context (“Brian, how is Jill doing?”).

Who is Maya Jupiter? ✔

Tell me something about Maya Jupiter. ✔

Thanks. I'm looking for a place to visit in the Katha area. ✔

Hello, how are you? ✘

What do you think? ✘

Yes! Brian, how is Jill doing? ✘

Figure 6. Instructions and example provided for the question “Is the query information-seeking?”

Response Options (2 / 4)

How is the **query relevant** to the conversation?

- ✔
 - **Follows up on a previous query or response.** It is difficult to correctly understand the query without reading the conversation history.
- +
 - **Only related to the topic of the conversation.** The query is topically similar to previous queries or responses, but can be understood without reading them.
- ✘
 - **Not relevant.** The query doesn't appear to be relevant to the topic or a previous query or response.
 - Rule of thumb: if you are surprised by a query, it is probably not relevant.

Who is Maya Jupiter? ✔

... Maya Jupiter is a rapper... ✔

When did she start rapping? ✔

Tell me something interesting about her. +

Why did Dr. Dre win a Grammy? ✘

Figure 7. Instructions and example provided for the question “How is the query relevant to the conversation?”

Response Options (3 / 4)

How **specific** is the query?

- ✓ • **Very.** Only a specific answer would satisfy the user.
 - Example: "Why did she make the news in 1999?" likely requires a very specific answer.
- + • **Somewhat.** A variety of answers of a specific kind would satisfy the user.
 - Example: While there are many possible answers to "What else does she do?", they are all likely to be a job or activity.
- ✗ • **Not at all.** Many topically different answers would satisfy the user.
 - Example: "Tell me something interesting about her." can be answered in many different ways.

Who is Maya Jupiter? ✓

How was her debut album received? ✓

What else does she **do**? +

Thanks. I'm looking for a place to visit in the Katha area. +

Tell me something interesting about her. ✗

What else? ✗

That's interesting, tell me more. ✗

Figure 8. Instructions and example provided for the question "How specific is the query?"

Response Options (4 / 4)

How well does the response **answer** the query?

- ★ • **Perfectly.** The response completely satisfies the user's information need.
- ✓ • **Sufficiently.** The response mostly answers the user's information need, though some additional information could be provided.
- + • **Incompletely.** The response provides some information relevant to the user, but doesn't adequately answer the question.
- ✗ • **Not at all.** The response does not provide any relevant information for the user's query or is not intelligible.

Who else does she rap with? ★

Over her career, she has collaborated with MC Trey, DJ Nick Toth and Triple J. ✓

She was a member of hip-hop group, Foreign Heights, with MC Trey and DJ Nick Toth. +

She is best known for her collaborations with other rappers. +

From 2004 to 2008, Jupiter hosted the national radio station, Triple J's weekly "Hip-Hop Show". ✗

Figure 9. Instructions and example provided for the question "How well does the response answer the query?"

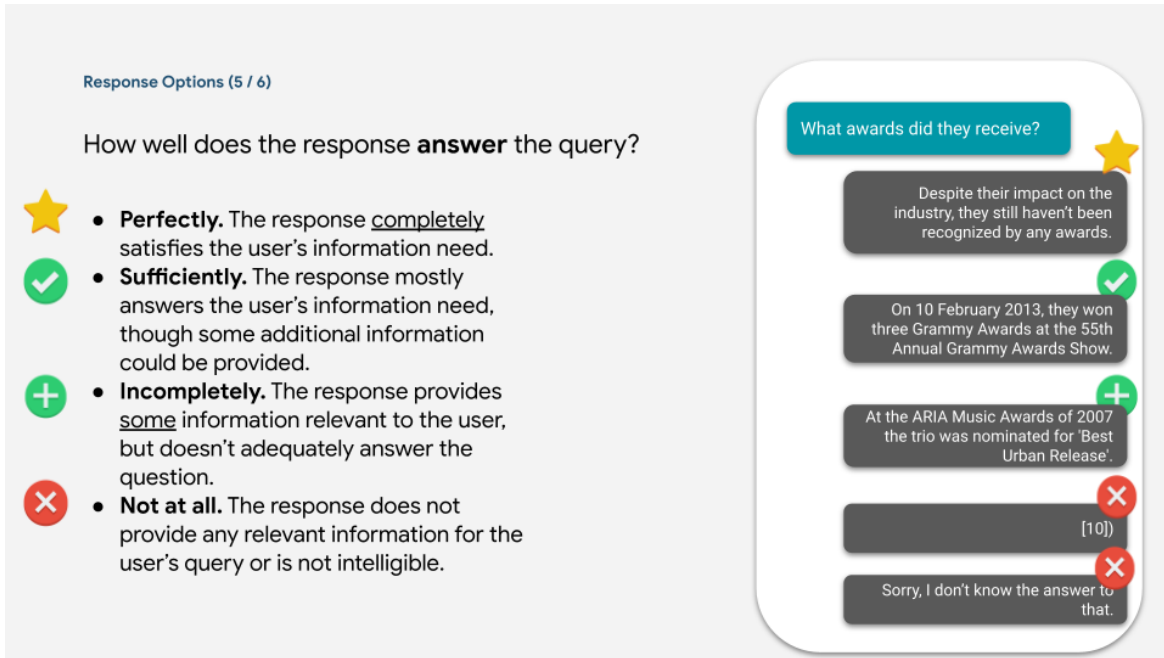


Figure 10. An additional example for the question “How well does the response answer the query?”

In this task, you will be spotting nonsensical or factually incorrect messages in artificially generated information-seeking conversations between a user and a system.

In each turn, raters were asked to answer the subjective questions in Table 1 and were provided the instructions and adjoining examples in Figures 6-10.

We iterated on our task design and instructions over two pilot annotations runs. Feedback from these pilots led to two main changes: (i) we extended the specificity and answer adequacy questions to be Likert scales rather than yes/no questions and (ii) we omitted a laborious question that asked raters to attest whether the passage supports the answer when interpreted in the context of the conversation history—we found that this perfectly correlated with whether or not the question was adequately answered or not.

Rater recruitment and training. We engaged with a vendor supplier of full-time crowd workers to recruit human annotators for our task. Raters were asked to review the above instructions and were provided direct feedback on their responses during the pilot annotation runs.

C. Retrieval models

C.1. Dual encoder retriever

The dual-encoder maps a query (q) and a passage (p) into dense embedding vectors $\text{embed}_\gamma(q)$ and $\text{embed}_\gamma(p)$,

where γ denotes model parameters. The relevance score between the two is their vector cosine similarity:

$$s_\gamma(q, p) = \frac{\text{embed}_\gamma(q)^\top \text{embed}_\gamma(p)}{\|\text{embed}_\gamma(q)\| \cdot \|\text{embed}_\gamma(p)\|}$$

This particular function enables one to retrieve the top-K highest-scoring passages for a given query using fast similarity search methods that run in sub-linear time (Ram & Gray, 2012).

For training, we minimize a standard contrastive loss with temperature τ :

$$\mathcal{L}(\gamma) = -\log \frac{\exp(s_\gamma(q, p^*)/\tau)}{\sum_{p \in p^* \cup \mathcal{N}(q)} \exp(s_\gamma(q, p)/\tau)}, \quad (2)$$

where p^* is a positive passage for q and $\mathcal{N}(q)$ denotes negative passages.

We implement the dual-encoder retriever following recent work (Ni et al., 2021): in particular, we use a shared Transformer encoder initialized from a T5 checkpoint, take the mean pooling of the top-most encoder layer, and project it to a fixed 768-dimensional embedding.

C.2. Reranker

The reranker model takes the same inputs as the dual encoder, but instead of encoding q and p into two separate vectors, it jointly encodes them into a single vector, $\text{embed}_\psi(q, p)$, where ψ denotes model parameters. It out-

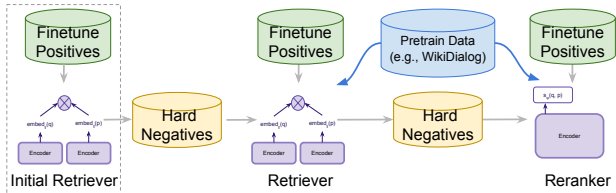


Figure 11. Our pipeline for training conversational retrieval system. It follows the standard multi-stage training scheme used in previous work (Lin et al., 2021). We first train an initial retriever with in-batch negatives. We then train a second retriever on hard negatives from the initial retriever. Finally, we train a reranker on hard negatives from Retriever. Performance is further improved by pre-training Retriever and Reranker on document derived-dialogs.

puts a relevance score:

$$s_\psi(q, p) = w^\top \text{embed}_\psi(q, p)$$

where w is also a model parameter. Unlike the dual encoder, this function does not support fast top-K retrieval in sub-linear time. Therefore, for computational tractability, we only use it to rerank a short-list of candidates retrieved by the dual encoder. However, the joint embedding of q and p permits a more expressive relevance function (e.g. cross-attention between q and p), so it can improve over the dual encoder’s relevance scores.

For training, we minimize a weighted binary classification loss:

$$\mathcal{L}(\psi) = -\log \sigma(s_\psi(q, p^*)) - \mathcal{N}(q)^{-1} \sum_{p \in \mathcal{N}(q)} \log[1 - \sigma(s_\psi(q, p))],$$

where σ denotes the sigmoid function.

Similar to our retriever, the reranker is also initialized from a T5 encoder. Our reranker implementation follows the implementation described in Nogueira et al. (2020).

C.3. Model Training

A training example from a typical retrieval dataset consists of a query paired with a positive passage, (q, p^*) . However, negative passages $\mathcal{N}(q)$ are usually not provided. Hence, we need to generate our own. We use two types of commonly used negatives: in-batch negatives, and “hard” negatives mined with a multi-stage training strategy.

In-batch negatives. When training with a batch, we treat the positive passage for example i as a negative for all other examples $\neq i$ in the same batch.

“Hard” negatives. Previous work has identified several weakness of in-batch negatives and identified the importance of selecting challenging negative examples for training retrievers (Qu et al., 2021; Xiong et al., 2021; Santhanam et al.,

	OR-QuAC	QReCC	CAsT-19	CAsT-20
# train/dev Dialogs	4k/0.5k	10k/0.5k	0	0
# train/dev Ques.	31k/3.4k	28k/1.5k	0	0
# test Dialogs	771	774	20	25
# test Ques.	5571	8209	173	208
# Doc.	11M	54M	38M	38M

Table 5. Statistics of ConvQA Retrieval datasets. CAsT-19 have 30 train and 50 test dialogs in total, but only 173 questions from 20 test dialogs have the relevance labels for evaluating retrieval model. Similarly, CAsT-20 only has 208 questions from 25 dialogs with relevance labels.

2021; Lin et al., 2021). Hence, following prior work (Lin et al., 2021), we adopt a multi-stage training strategy where new negatives are mined at each stage (Figure 11):

1. We train an initial retriever using only positives and in-batch negatives.
2. We run top- K retrieval using the initial retriever, and then randomly sample a subset of those to serve as negatives (when K is large, a high percentage of these are true negatives). The hard negatives are then combined with positives to train a second retriever.
3. We use the second retriever to again generate hard negatives. We train our reranker on these hard negatives combined with any positives retrieved by the second retriever.

When pre-training on WikiDialog, we use in-batch negatives. When fine-tuning for a downstream task, we use the multi-stage hard negative strategy.

Note that at inference time, we use the second retriever for top-K retrieval, and use the reranker to refine the top-K ranking (the initial retriever is not used at inference time).

D. Detailed Experimental Setup

D.1. Datasets

We use four open-domain conversational QA retrieval benchmarks: OR-QuAC (Qu et al., 2020), TREC CAsT-19 (Byrne et al., 2019), TREC CAsT-20 (Dalton et al., 2020), and QReCC (Anantha et al., 2021). Table 5 lists their statistics.

OR-QuAC and QReCC were introduced in Appendix A.1.

TREC CAsT-19 and CAsT-20 are two datasets from the The TREC Conversational Assistance Track (CAsT) shared task (Dalton et al., 2019; 2020) with small numbers of dialogues for evaluating information-seeking conversational search systems. Questions in a dialogue are constructed manually to mimic a “real” dialogue on a certain topic. The retrieval corpus includes web passages from MS

MARCO (Nguyen et al., 2016) and wikipedia passages from TREC Complex Answer Retrieval (CAR) (Dusart et al., 2019). CAsT-19 provides human relevance labels for 173 questions in 20 test dialogues. CAsT-20 provides human relevance labels for 208 questions in 25 test dialogues.

Note, QReCC reuses question sequences from QuAC and TREC CAsT-19 (Byrne et al., 2019). However, although the questions are the same as existing datasets, *the relevant passages and answers are different*, as it asked human raters to retrieve passages with a search engine and generate answers. In addition, TREC questions only appear in the *test* set.

In OR-QuAC, at least 99.5% of dialogs contain answers from the same answer passage, while in CAsT and QReCC, each question turn can be answered by a different passage.

We follow the official *automatic* setting (Byrne et al., 2019; Dalton et al., 2020) for both datasets, which only uses the questions as retrieval inputs. We follow the authors of QReCC (Anantha et al., 2021) and (Wu et al., 2021) and use the both questions and gold answers from conversation history as retrieval inputs.

D.2. Implementation

We implement dual-encoder retrievers and rerankers in JAX.

Retrievers. For pre-training on our inpainted datasets, we used a softmax temperature τ of 0.01, batch size 2048, and dropout rate 0.1. The models were trained with Adafactor optimizer with learning rate $1e^{-3}$ and 1k warm up steps. For checkpoint selection, we tested checkpoints at 50k and 100k steps and reported the better one based on each finetune datasets’ dev set performance. We fine-tuned the retrievers for 500 steps on OR-QuAC and QReCC. We did not fine-tune the retrievers on TREC CAsT-19 and CAsT-20 due to the small data size. We report TREC performance by retrieving with a QReCC retriever (trained without answers in the inputs).

Questions and passages are always lowercased. Maximum query length was set to 128 for all pretrain and fine-tune datasets except for QReCC, which uses 512 because it allows using previous answers in the queries. Maximum passage length was set to 256.

Rerankers. Rerankers were implemented using T5 encoder-decoder architecture. We pre-trained the rerankers on synthetic dialogs for 1.7M steps. For fine-tuning on OR-QuAC and QReCC, we used 10 hard negatives for each question sampled from the top 100 passages returned by the retriever. For TREC CAsT-19 and CAsT-20, we follow prior work (Yu et al., 2021) and used 5-fold cross-validation to fine-tune the reranker¹⁶. We used 20 hard negatives per question ran-

domly sampled from the top 100 retrieved results. We warm up the TREC rerankers on QReCC.

All reranker training used the Adafactor optimizer with constant learning rate $1e^{-3}$ and dropout rate 0.1. We used batch size 512 for pre-training, 128 for fine-tuning of OR-QuAC and QReCC, and 32 for fine-tuning of TREC CAsT-19 and CAsT-20.

D.3. Published baselines

We include five published *retrieval-only* baselines to compare with our DI retrievers.

BM25-Query Rewriter (Yu et al., 2021) and BM25-T5QR (Wu et al., 2021) are two query rewriting approaches that trains a model to rewrite the dialog history into a contextualized, keyword-like query. The former trains a GPT-2 query rewriter on ad hoc search sessions (Yu et al., 2020). The latter trains a T5 query rewriter on human-generated query rewrites from QReCC. Both systems issue the rewritten query to a classic BM25 lexical retrieval system.

ANCE-Query Rewriter (Yu et al., 2021) uses the GPT-2 query rewriter from (Yu et al., 2021), but it runs the rewritten query with ANCE (Xiong et al., 2021), is a strong non-conversational dense retriever trained on MS Marco.

CONQRR (Wu et al., 2021) is the previous state-of-the-art retrieval system on QReCC. CONQRR uses reinforcement learning to optimize the query rewriter for the retriever. Its retriever is a t5-base dual-encoder trained on MS Marco.

ConvDR (Yu et al., 2021) is the previous state-of-the-art conversational dense retrieval system on OR-QuAC and TREC CAsT tasks. Unlike the above approaches which all use query rewriting, ConvDR learns a conversational query encoder that directly maps the entire dialog history into a dense embedding. To address the bottleneck of limited conversational training data, ConvDR uses a teacher-student framework that trains the student conversational query encoder to “mimic” the representation of the oracle query rewrite from a non-conversational teacher.

In addition to the retrieval baselines, we include three published state-of-the-art conversational search systems that uses the *retrieval+reranker* pipeline, serving as baselines for our DI retriever + reranker system.

CFDA_CLIP_RUN7 (Yang et al., 2019) is the best performing system participated in TREC CAsT-19 automatic setting (Dalton et al., 2019). It uses heuristics to expand queries, BM25 with T5 doc2query for retrieval, and BERT to rerank.

retrievers on TREC CAsT. However, it is very expensive due to inference on the retrieval corpus, so we did not fine-tune retriever in this work.

¹⁶Theoretically, one could also use cross-validation to fine-tune

h2olo0_RUN4 is the best participating system in TREC CAsT-20 automatic setting (Dalton et al., 2020). It uses a T5 model for query rewriting, a dense-sparse hybrid retriever, and a T5 reranker trained on MS MARCO.

ConvDR→BERT (Yu et al., 2021) uses a BERT reranker to rerank ConvDR (Yu et al., 2021)'s retrieval results. The BERT reranker is pre-trained on MS MARCO and then fine-tuned for specific datasets on both the ranking loss and the query rewrite distillation loss.

Article: How Long Does Tylenol Take To Work

1: How long does it take for tylenol to work ?

0: If you have taken the medicine without having food (the stomach is empty), the duration it will make its mark is fifteen to twenty minutes.

1: What if you have taken it with food ?

0: If your stomach is filled (after a meal), and you take Tylenol, the time it will take to produce effect is half an hour to sixty minutes.

1: Does it work well ?

0: Similar to any medicine, it cannot do wonders.

1: Am happy i took it.

0: Kindly wait and the healing process will fall in place.

Article: How to get your iTunes music on your Windows phone

1: How do I get my music from iTunes on to my windows phone?,

0: 1 Locate the place on your computer where your music is stored. 2 Open Windows Explorer (on Windows 7) or File Explorer (Windows 8 or 10). 3 Your Windows phone will show up as a device in the left pane. 4 Click on your phone on the left, then on the hard drive icon. 5 Inside you will see a Music folder.ou will need to know where your music files are stored.

1: I don't know where they are stored.

0: Chances are, you will see them in your Music Library.

1: I don't see them there.

0: If your music was all purchased from iTunes, you may also find it in C:\\Users\\your name\\My Music\\iTunes\\iTunes Media.

Article: How Your Immune System Works

1: What is the most obvious part of the immune system?

0: The most obvious part of the immune system is what you can see.

1: Like what?

0: For example, skin is an important part of the immune system.

1: Why is it important?

0: It acts as a primary boundary between germs and your body.

1: How does it act as a boundary?

0:Part of your skin's job is to act as a barrier in much the same way we use plastic wrap to protect food.

Table 6. Examples from WebDialog. 1 is model inpainted utterances, and 0 is from the original passage.

E. Examples of Inpainting Non-Wikipedia Documents

WebDialog was generated from the MS Marco passage retrieval corpus (Nguyen et al., 2016). We used all 8.4M passages from the MS Marco dataset, without any filtering. Table 6 shows a few examples from WebDialog.

Model: How can you have so many pancakes?
 Document: I love pancakes like a fat kid loves cake.
 Model: I guess your body requires them as your emotional needs.
 Document: You want to know the sad part?
 Model: Go ahead.
 Document: I'm TERRIBLE at making pancakes.

Table 7. Example of inpainting a less-formal document.

In Table 7, we would also like to share an example where inpainting yields more natural dialog when applied to a less-formal document.

F. Additional Experimental Results

Table 8 reports additional retrieval metrics in addition to the MRR reported in Table 3. Specifically, we report recall and mean reciprocal rank at rank 5 (R@5 and MRR@5) following previous work Qu et al. (2020); Yu et al. (2021). On QReCC, we report recall at rank 10 (R@10) and mean reciprocal rank without rank cut off (MRR) following Anantha et al. (2021); Wu et al. (2021). On TREC CAsT-19 and CAsT-20, we use the official metrics MRR and NDCG@3 suggested by Dalton et al. (2019; 2020). Note that TREC CAsT-19 uses relevant grade ≥ 1 as positive for MRR but TREC CAsT-20 uses relevance grade ≥ 2 as positive for MRR (Dalton et al., 2020)

System	QReCC		OR-QuAC		TREC CAsT-19		TREC CAsT-20	
	MRR	R@10	MRR@5	R@5	MRR	NDCG@3	MRR	NDCG@3
	<i>Retrieval</i>							
BM25-QR	-	-	20.2	30.2	58.1	27.7	25.0	15.9
ANCE-QR	-	-	45.7	58.4	66.5	40.9	37.5	25.5
ConvDR	-	-	61.6	75.0	74.0	46.6	50.1	34.0
BM25-T5QR	32.8	52.5	-	-	-	-	-	-
CONQRR	41.8	65.1	-	-	-	-	-	-
T5-Large DE	55.7	78.7	56.9	69.7	61.0	32.5	34.3	19.9
▷ WIKI	60.4	83.3	66.5	77.9	68.1	40.3	43.7	27.9
▷ WIKI+WEB	60.7	83.1	68.7	80.3	74.1	47.0	51.3	33.2
	<i>Retrieval + Reranking</i>							
CFDA_CLIP_RUN7	-	-	-	-	71.4	43.6	-	-
h2olo0o_RUN4	-	-	-	-	-	-	59.3	45.8
ConvDR→BERT	-	-	77.3	85.0	79.9	54.1	54.5	39.2
T5-Large DE → reranker	68.9	86.0	72.6	78.6	75.3	48.3	55.1	36.7
▷ WIKI	70.7	88.8	79.7	86.8	79.3	50.7	60.3	43.0
▷ WIKI+WEB	71.8	88.9	81.2	87.9	82.0	53.6	59.7	42.5

Table 8. Full evaluation on QReCC, OR-QuAC and TREC CAsT. This table reports additional metrics for these datasets following prior work (Yu et al., 2021; Wu et al., 2021) in addition to MRR reported in Table 3.

G. Inpainted Dialog Examples

Tables 9-12 show additional examples of inpainted dialogs.

H. Studying Sensitive Language in the Dataset

Prior to evaluating the dataset for fairness implications, we found it necessary to define “fairness” in the context of the model’s generated queries. We began our inquiry by considering conversational fairness – that is, what constitutes insensitive, offensive, or prejudiced speech in a dialogue? We then qualitatively assessed a small sample of the dataset and described the following potential fairness related concerns that may arise in the dataset:

- Model generates a query that contains sensitive characteristics that were not explicitly mentioned in the passage.

Article: European School, Munich

Inpaint _{PT}	<p>1: What is the European School, Munich?</p> <p>0: The European School, Munich (ESM) is one of thirteen European Schools and one of three in Germany.</p> <p>1: What did the European School, Munich do?</p> <p>0: First established in 1977, it moved to its current location in Neuperlach, a district in the south-east of Munich, in the state of Bavaria in 1981.</p> <p>1: Are there any other interesting aspects about this article?</p> <p>0: The ESM was principally established to serve the schooling needs of children of the staff of the European Patent Office (EPO) – the executive body of the European Patent Organisation.</p> <p>1: Did it serve any other students?</p> <p>0: However, enrolment is open to other prospective students.</p> <p>1: What other services did it offer?</p> <p>0: The school offers the European Baccalaureate as its secondary leaving qualification.”</p>
Inpaint _{OQ}	<p>1: What is the European School in Munich?</p> <p>0: The European School, Munich (ESM) is one of thirteen European Schools and one of three in Germany.</p> <p>1: Where is it located?</p> <p>0: First established in 1977, it moved to its current location in Neuperlach, a district in the south-east of Munich, in the state of Bavaria in 1981.</p> <p>1: Who attended/teaches there?</p> <p>0: The ESM was principally established to serve the schooling needs of children of the staff of the European Patent Office (EPO) – the executive body of the European Patent Organisation.</p> <p>1: Was it later expanded to serve the needs of the general public?</p> <p>0: However, enrolment is open to other prospective students.</p> <p>1: what types of classes are taught there?</p> <p>0: The school offers the European Baccalaureate as its secondary leaving qualification.”</p>
Inpaint _{PTOQ}	<p>1: What is the purpose of the European School, Munich?</p> <p>0: The European School, Munich (ESM) is one of thirteen European Schools and one of three in Germany.</p> <p>1: Where is it located?</p> <p>0: First established in 1977, it moved to its current location in Neuperlach, a district in the south-east of Munich, in the state of Bavaria in 1981.</p> <p>1: Is it open to anyone?</p> <p>0: The ESM was principally established to serve the schooling needs of children of the staff of the European Patent Office (EPO) – the executive body of the European Patent Organisation.</p> <p>1: Are there any other interesting aspects about this article?</p> <p>0: However, enrolment is open to other prospective students.</p> <p>1: What are some of the courses taught?</p> <p>0: The school offers the European Baccalaureate as its secondary leaving qualification.”</p>

Table 9. Synthetic dialogs about *European School, Munich* generated by Inpaint_{PT}, Inpaint_{OQ}, and Inpaint_{PTOQ}.

Dialog Inpainting: Turning Documents into Dialogs

Article: Ageing, Disability and Home Care NSW

Inpaint _{PT}	<p>1: What is Ageing, Disability and Home Care NSW?</p> <p>0: Ageing, Disability and Home Care NSW (ADHC) is a division of the Department of Family and Community Services in the Government of New South Wales that is responsible for the provision of services to older people, people with a disability, and their families and carers in the state of New South Wales, Australia.</p> <p>1: What services does it provide?</p> <p>0: The ADHC division supports approximately 300,000 vulnerable people from its annual budget of 2.5 billion.</p> <p>1: Does it have a large staff?</p> <p>0: It employs more than 13,000 staff.</p> <p>1: What is it known for?</p> <p>0: ADHC delivers services and manages 900 non-government and local government service providers that are funded to deliver services such as early intervention, therapy, respite, supported accommodation, case management and post-school programs.</p> <p>1: Who runs it?</p> <p>0: The division is led by a Deputy Secretary, presently Jim Longley, who reports to the Minister for Ageing and Minister for Disability Services, presently the Hon. John Ajaka MLC.</p>
Inpaint _{OQ}	<p>1: What is Ageing, Disability and Home Care NSW?</p> <p>0: Ageing, Disability and Home Care NSW (ADHC) is a division of the Department of Family and Community Services in the Government of New South Wales that is responsible for the provision of services to older people, people with a disability, and their families and carers in the state of New South Wales, Australia.</p> <p>1: How many people do they serve?</p> <p>0: The ADHC division supports approximately 300,000 vulnerable people from its annual budget of 2.5 billion.</p> <p>1: How many employees do they have?</p> <p>0: It employs more than 13,000 staff.</p> <p>1: What kind of services do they provide?</p> <p>0: ADHC delivers services and manages 900 non-government and local government service providers that are funded to deliver services such as early intervention, therapy, respite, supported accommodation, case management and post-school programs.</p> <p>1: Who are the executives?</p> <p>0: The division is led by a Deputy Secretary, presently Jim Longley, who reports to the Minister for Ageing and Minister for Disability Services, presently the Hon. John Ajaka MLC.</p>
Inpaint _{PTOQ}	<p>1: Who does Ageing, Disability and Home Care NSW (ADHC) work for?</p> <p>0: Ageing, Disability and Home Care NSW (ADHC) is a division of the Department of Family and Community Services in the Government of New South Wales that is responsible for the provision of services to older people, people with a disability, and their families and carers in the state of New South Wales, Australia.</p> <p>1: How much does it cost?</p> <p>0: The ADHC division supports approximately 300,000 vulnerable people from its annual budget of 2.5 billion.</p> <p>1: How many people work for the division?</p> <p>0: It employs more than 13,000 staff.</p> <p>1: What services does the division provide?</p> <p>0: ADHC delivers services and manages 900 non-government and local government service providers that are funded to deliver services such as early intervention, therapy, respite, supported accommodation, case management and post-school programs.</p> <p>1: Who is in charge of the division?</p> <p>0: The division is led by a Deputy Secretary, presently Jim Longley, who reports to the Minister for Ageing and Minister for Disability Services, presently the Hon. John Ajaka MLC.</p>

Table 10. Synthetic dialogs about *Ageing, Disability and Home Care NSW* generated by Inpaint_{PT}, Inpaint_{OQ}, and Inpaint_{PTOQ}.

Inpaint _{PT}	<p>1: What is FAQ?</p> <p>0: An FAQ is a list of frequently asked questions (FAQs) and answers on a particular topic (also known as Questions and Answers [Q&A] or Frequently Answered Questions).</p> <p>1: Are there any other interesting aspects about this article?</p> <p>0: The format is often used in articles, websites, email lists, and online forums where common questions tend to recur, for example through posts or queries by new users related to common knowledge gaps.</p> <p>1: Is there anything else important?</p> <p>0: The purpose of an FAQ is generally to provide information on frequent questions or concerns; however, the format is a useful means of organizing information, and text consisting of questions and their answers may thus be called an FAQ regardless of whether the questions are actually "frequently" asked.</p> <p>1: Any other information?</p> <p>0: Since the acronym "FAQ" originated in textual media, its pronunciation varies.</p> <p>1: How does it vary?</p> <p>0: FAQ is most commonly pronounced as an initialism, "F-A-Q", but may also be pronounced as an acronym, "FAQ".</p>
Inpaint _{OQ}	<p>1: What is meant by a "Frequently asked question"?</p> <p>0: An FAQ is a list of frequently asked questions (FAQs) and answers on a particular topic (also known as Questions and Answers [Q&A] or Frequently Answered Questions).</p> <p>1: Who asks the questions?</p> <p>0: The format is often used in articles, websites, email lists, and online forums where common questions tend to recur, for example through posts or queries by new users related to common knowledge gaps.</p> <p>1: Is the information intended for a specific audience?</p> <p>0: The purpose of an FAQ is generally to provide information on frequent questions or concerns; however, the format is a useful means of organizing information, and text consisting of questions and their answers may thus be called an FAQ regardless of whether the questions are actually "frequently" asked.</p> <p>1: How is the term "FAQ" pronounced?</p> <p>0: Since the acronym "FAQ" originated in textual media, its pronunciation varies.</p> <p>1: What does it vary in pronunciation?</p> <p>0: FAQ is most commonly pronounced as an initialism, "F-A-Q", but may also be pronounced as an acronym, "FAQ".</p>
Inpaint _{PTOQ}	<p>1: what is the purpose of a faq</p> <p>0: An FAQ is a list of frequently asked questions (FAQs) and answers on a particular topic (also known as Questions and Answers [Q&A] or Frequently Answered Questions).</p> <p>1: Are these often used on a website?</p> <p>0: The format is often used in articles, websites, email lists, and online forums where common questions tend to recur, for example through posts or queries by new users related to common knowledge gaps.</p> <p>1: What questions are asked most frequently?</p> <p>0: The purpose of an FAQ is generally to provide information on frequent questions or concerns; however, the format is a useful means of organizing information, and text consisting of questions and their answers may thus be called an FAQ regardless of whether the questions are actually "frequently" asked.</p> <p>1: Is FAQ pronounced as an acronym?</p> <p>0: Since the acronym "FAQ" originated in textual media, its pronunciation varies.</p> <p>1: How is it pronounced?</p> <p>0: FAQ is most commonly pronounced as an initialism, "F-A-Q", but may also be pronounced as an acronym, "FAQ".</p>

Table 11. Synthetic dialogs about FAQ generated by Inpaint_{PT}, Inpaint_{OQ}, and Inpaint_{PTOQ}.

Article: George B. McClellan

Inpaint _{PT}	<p>1: Did George B. McClellan subordinates blame him for the failures on the campaign of 1862?</p> <p>0: Although McClellan’s subordinates can claim their share of responsibility for delays (such as Ambrose Burnside’s misadventures at Burnside Bridge) and blunders (Edwin V. Sumner’s attack without reconnaissance), these were localized problems from which the full army could have recovered.</p> <p>1: What other blunder happened during the campaign?</p> <p>0: Historians have faulted McClellan for accepting the cautious advice about saving his reserves, such as when a significant breakthrough in the center of the Confederate line could have been exploited, but Fitz John Porter is said to have told McClellan, ”Remember, General , I command the last reserve of the last Army of the Republic.” However, the veracity of this supposed statement is in doubt.</p> <p>1: Why is it in doubt?</p> <p>0: Porter in the post-war period never claimed to have made the statement and it also fails to note the several Union corps at that time defending Washington.</p> <p>1: Did the campaign end in a draw?</p> <p>0: Despite being a tactical draw, Antietam is considered a turning point of the war and a victory for the Union because it ended Lee’s strategic campaign (his first invasion of the North) and it allowed President Lincoln to issue the Emancipation Proclamation on September 22, taking effect on January 1, 1863.</p>
Inpaint _{OQ}	<p>1: What was George B. McClellan’s biggest failure?</p> <p>0: Although McClellan’s subordinates can claim their share of responsibility for delays (such as Ambrose Burnside’s misadventures at Burnside Bridge) and blunders (Edwin V. Sumner’s attack without reconnaissance), these were localized problems from which the full army could have recovered.</p> <p>1: What was his biggest success?</p> <p>0: Historians have faulted McClellan for accepting the cautious advice about saving his reserves, such as when a significant breakthrough in the center of the Confederate line could have been exploited, but Fitz John Porter is said to have told McClellan, ”Remember, General , I command the last reserve of the last Army of the Republic.” However, the veracity of this supposed statement is in doubt.</p> <p>1: Why is the veracity of this statement in doubt?</p> <p>0: Porter in the post-war period never claimed to have made the statement and it also fails to note the several Union corps at that time defending Washington.</p> <p>1: What was the biggest battle he was involved in?</p> <p>0: Despite being a tactical draw, Antietam is considered a turning point of the war and a victory for the Union because it ended Lee’s strategic campaign (his first invasion of the North) and it allowed President Lincoln to issue the Emancipation Proclamation on September 22, taking effect on January 1, 1863.</p>
Inpaint _{PTOQ}	<p>1: What did George B. McClellan do during the Seven Days’ Campaigns?</p> <p>0: Although McClellan’s subordinates can claim their share of responsibility for delays (such as Ambrose Burnside’s misadventures at Burnside Bridge) and blunders (Edwin V. Sumner’s attack without reconnaissance), these were localized problems from which the full army could have recovered.</p> <p>1: Are there any other interesting aspects about this article?</p> <p>0: Historians have faulted McClellan for accepting the cautious advice about saving his reserves, such as when a significant breakthrough in the center of the Confederate line could have been exploited, but Fitz John Porter is said to have told McClellan, ”Remember, General , I command the last reserve of the last Army of the Republic.” However, the veracity of this supposed statement is in doubt.</p> <p>1: Why is it in doubt?</p> <p>0: Porter in the post-war period never claimed to have made the statement and it also fails to note the several Union corps at that time defending Washington.</p> <p>1: Was it a successful campaign?</p> <p>0: Despite being a tactical draw, Antietam is considered a turning point of the war and a victory for the Union because it ended Lee’s strategic campaign (his first invasion of the North) and it allowed President Lincoln to issue the Emancipation Proclamation on September 22, taking effect on January 1, 1863.</p>

Table 12. Synthetic dialogs about *George B. McClellan* generated by Inpaint_{PT}, Inpaint_{OQ}, and Inpaint_{PTOQ}.

Dialog Inpainting: Turning Documents into Dialogs

- Model overemphasizes specific sensitive characteristics or identities in a negative or offensive way. For example: the model continues to generate questions related to a writer’s biracial identity when it is not the intended focus of the passage.

- Model ascribes negative values to sensitive characteristics without that being explicitly mentioned in the passage. For example, the model generates a question about “how barbaric” a society is for a passage about a non-Western society.

- Model produces “unhelpful” or repetitive questions when the source passage contains sensitive characteristics as compared to passages with non-sensitive characteristics.

The primary approach to evaluating fairness in the dataset involved scanning for sensitive terms and interactions between those terms. To this end, we curated a list of sensitive terms from the literature that includes nouns and adjectives that are considered harmful or insensitive when applied to particular groups (citation). When a combination of these terms appeared in a question in the dataset (Table 13), we considered it to be potentially sensitive. From this initial analysis, we observed that 0.2-0.5% of the dataset contains such potentially sensitive interactions of terms in generated queries. See Tables 14–18 for selected examples that matched this filter.

	Race/Ethnicity	Nationality	Gender	Region	Political	Religion
Adjective	✓	✓	✓			
Violence Related	✓	✓		✓		✓
Competence Adjective	✓	✓	✓			✓
Physical Appearance Adjectives	✓		✓			
Outsider Adjectives	✓	✓	✓			✓
Harmful Racial Stereotypes	✓		✓			
Race/Ethnicity Adjectives	✓	✓				
Gender Adjectives	✓		✓			
Political				✓		
Sexual Orientation	✓		✓			
Socio-economic/class	✓	✓			✓	
Occupation	✓	✓	✓			

Table 13. The categories of sensitive terms and their interactions.

Dialog Inpainting: Turning Documents into Dialogs

Article: Model minority

- 1: The term model minority is often used in reference to **African** Americans, **Hispanic** Americans, and Asian Americans. But what does the term actually mean?
- 0: The model minority label also includes South Asian communities, in particular, Indian Americans, because of their high socioeconomic success.
- 1: And what can Indian Americans do to help other Asian Americans?
- 0: But, Indian-Asians face a type of racism and discrimination, despite their Model Minority status, that other Asians do not face.
- 1: Can you tell me more about this racism?
- 0: One of the forms of this is discrimination based on religion; often times, Indians, specifically Sikhs, in America face racism that that is a result of mislabeling like being called Middle-Eastern or Muslim by the way they look, whether they actually identify with those groups, with Muslim Americans being markedly poorer than Hindu Americans.
- 1: And how does this impact other Asian Americans?
- 0: The mass grouping of all different Asian ethnicities into one “positive” stereotype delegitimizes the problems that certain Asian immigrants face, because the majority might not face those same negative stereotypes.
- 1: What can the Sikhs do to help other Asian immigrants?
- 0: Another result of the Model Minority status of Indian-Asians is their complacency, expressed more often than not.
- 1: How does this affect other Asian immigrants?
- 0: There is conflict within the Indian-Asian immigrant population in the US, as some exalt their “positive” minority status and wish to uphold it, resulting in complacency when attacked racially and even sometimes going as far as to support political candidates that are anti-immigration.
-

Table 14. This dialog was flagged as potentially sensitive because it contained sensitive terms “african” and “hispanic”.

Article: The Queers

- 1: What was the band The Queers’ response to the November 2015 Woolsey Fire disaster?
- 0: In an interview with Noisey, after their February 2016 show was cancelled due to a call for a petition to boycott the show from a local collective, Babely Shades, Joe Queer voiced his opinion when asked if he “supports racist, misogynistic or homophobic values?” he replied “Of course not.”
- 1: Did Joe Queer think that the Woolsey Fire was man-made?
- 0: “What normal minded person does?”
- 1: Did Joe Queer have any black friends?
- 0: “My closest pals here in Atlanta are Chris and Chester, two gay black guys.”
- 1: Was Joe Queer in school when the Woolsey Fire happened?
- 0: “I was actually working a part time job.”
- 1: Was Joe Queer working with any black people at the time?
- 0: “I was one of two white people on the whole job.”
- 1: Did Joe Queer think that **African Americans** were **lazy** and lacked work ethic?
- 0: “If I was so racist, I wouldn’t have worked there or be living in Atlanta!”
-

Table 15. This dialog was flagged as potentially sensitive because it contains sensitive terms “african” and “lazy”.

Dialog Inpainting: Turning Documents into Dialogs

Article: Female intrasexual competition

- 1: How does a woman's ovulation cycle affect the way she competes with other women?
0: This indirect form of competition appears exclusive toward females as findings also showed that women, irrespective of ovulation status (high or low), showed no difference in the rating of male faces.
1: What other behaviors or factors affect the way women compete with each other?
0: Supporting research has also found that younger women who are considered as having high fertility, gossip about other women more than older women, who are no longer at their most fertile stage.
1: Do women tend to be more **aggressive** toward each other than **men**?
0: Indeed, indirect aggression appears more prevalent amongst (or exclusive to) females than males who are said to engage in more direct forms of competition.
1: Are there any other interesting aspects about this article?
0: Research studying the relationship between indicators of attractiveness, such as physical attractiveness and indirect victimisation, showed that the likelihood of experiencing indirect victimization increased by 35% for females who perceived themselves as physically attractive.
1: Is this something that also happens with men?
0: In contrast, being a male who is physically attractive decreased the chances of experiencing such indirect victimization.
1: Why is this difference in behavior between men and women important?
0: This also highlights how the physical attractiveness a female is a trigger for indirect aggression and forms a core part of intersexual selection between the sexes.
-

Table 16. The dialog is potentially sensitive due to containing sensitive terms “aggressive” and “men”.

Article: Virginia M. Alexander

- 1: who was the first **african american** woman to become a **doctor** in the united states?
0: Virginia M. Alexander Virginia M. Alexander (February 4, 1899 – July 24, 1949) was an American physician, public health researcher, and the founder of the Aspiranto Health Home in Philadelphia, Pennsylvania.
1: what was her childhood like?
0: Virginia M. Alexander was born in Philadelphia, Pennsylvania on February 4, 1899 to Hilliard Alexander and Virginia Pace.
1: did she have siblings?
0: She had four siblings, including the prominent attorney Raymond Pace Alexander.
1: what were her parents like?
0: Alexander's mother died when she was 4 years old, and at age 13 her father's riding academy closed.
1: what did she do after her parents' academy closed?
0: Alexander withdrew from school to help relieve the resulting economic strain on her family, but her father insisted that she finish her education.
1: where did she go to school?
0: Alexander attended high school at the William Penn High School for Girls, where she graduated with honors before receiving a scholarship that allowed her to attend the University of Pennsylvania to complete her undergraduate education.
-

Table 17. This dialog is flagged as potentially sensitive because it contains the sensitive terms “African-American” and “doctor”. The first answer is not factually accurate—Virginia M. Alexander is not the first African-American physician.

Article: Symbolism of terrorism

1: “**Allah** is most gracious to the oppressed (Quran passage)” is mentioned in the analysis of a terrorist **attack**. Is this quote from the Quran a part of the analysis?

0: The analysis does not discuss the quote on the Quran.

1: What is the image in the analysis?

0: The central figure of the image is a yellow tiger, which is surrounded by a ring of bullets and crossed by two rifles.

1: What is the significance of the image?

0: The script both around the tiger, in Tamil, and below the tiger, in English, reads “Liberation Tigers of Tamil Eelam.” Eelam is Tamil for the Island of Sri Lanka, which the groups seeks to liberate from the control of the Sinhalese majority.

1: What is the significance of the tiger?

0: The tiger is a culturally important animal that symbolizes heroism, militancy, and patriotism.

1: What is the significance of the tiger in this image?

0: Its centrality in this image indicates the desire of the Liberation Tigers to be self-determined and independent.

1: What is the significance of the bullets and rifles?

0: The bullets and rifles indicate their employment of violence as the means to achieve these goals.

Table 18. This dialog is flagged as potentially sensitive because it contains the sensitive terms “allah” and “attack”.