

# Rome was built in 1776: A Case Study on Factual Correctness in Knowledge-Grounded Response Generation

Sashank Santhanam<sup>1</sup> Behnam Hedayatnia<sup>2</sup> Spandana Gella<sup>2</sup>  
Aishwarya Padmakumar<sup>2</sup> Seokhwan Kim<sup>2</sup> Yang Liu<sup>2</sup> Dilek Hakkani-Tur<sup>2</sup>

<sup>1</sup>University of North Carolina, Charlotte, <sup>2</sup>Amazon Alexa AI

ssanthan1@uncc.edu

{behnam, sgella, padmakua, seokhwk, yangliud, hakkanit}@amazon.com

## Abstract

Recently neural response generation models have leveraged large pre-trained transformer models and knowledge snippets to generate relevant and informative responses. However, this does not guarantee that generated responses are factually correct. In this paper, we examine factual correctness in knowledge-grounded neural response generation models. We present a human annotation setup to identify three different response types: responses that are factually consistent with respect to the input knowledge, responses that contain hallucinated knowledge, and non-verifiable chit-chat style responses. We use this setup to annotate responses generated using different state-of-the-art models, knowledge snippets, and decoding strategies. In addition, to facilitate the development of a factual consistency detector, we automatically create a new corpus called Conv-FEVER that is adapted from the Wizard of Wikipedia dataset and includes factually consistent and inconsistent responses. We demonstrate the benefit of our Conv-FEVER dataset by showing that the models trained on this data perform reasonably well to detect factually inconsistent responses with respect to the provided knowledge through evaluation on our human annotated data. We will release the Conv-FEVER dataset and the human annotated responses.

## 1 Introduction

One of the biggest challenges in the field of conversational AI is to develop agents that can generate coherent and grammatical responses whilst also seamlessly blending in knowledge to produce informative responses (Roller et al., 2020a). With the development of large scale language models, conversational AI systems have made significant progress in being able to generate coherent and grammatical text. To produce informative responses, prior neural response generation work proposed grounding on knowledge sentences that are relevant to the

---

### Dialog History:

**Speaker 1:** What do you think about Murray?

**Speaker 2:** I think Murray is a great player he just needs to stay healthy in order to compete more. Who do you like best?

---

### Knowledge:

In tennis, the term Big Four refers to the quartet of men’s singles players comprising Roger Federer, Rafael Nadal, Novak Djokovic, and Andy Murray.

---

### Non-verifiable Response:

I like Federer and Nadal. I think Federer is the best.

**Verifiable** = 1; **F.C** = N/A; **H** = N/A

---

### Factually Consistent Response:

Rafael Nadal is my favorite of the Big Four.

**Verifiable** = 3; **F.C** = 1; **H** = No

---

### Hallucinated Response:

I like Djokovic. **He has played in the top ten singles players of the world.**

**Verifiable** = 4; **F.C** = 1; **H** = Yes

---

Figure 1: *Non-verifiable Response* does not include any information that needs to be verified and cannot be evaluated as consistent or not consistent. *Factually Consistent Response* is consistent with the provided knowledge. *Hallucinated Response* is not consistent with the knowledge but may still be correct. **F.C** = Factual Consistency. **H** = Hallucination. Description of the annotation for these labels are in Section 3.1.

dialog context (Ghazvininejad et al., 2018; Yavuz et al., 2019; Zhou et al., 2018; Roller et al., 2020b). However, it is not guaranteed that these models generate responses that leverage the provided knowledge and are factually correct. Being able to generate factually correct responses is important because providing incorrect information can reduce response quality and can even be critical especially for certain domains such as politics, medicine, or finance.

Previous studies have evaluated a neural model’s factual correctness in several fields. For example, for large pre-trained models, it is shown that these generative models are able to memorize knowledge, but still generate factually incorrect responses (Petroni et al., 2019; Logan et al., 2019; Roberts et al., 2020).

Factual correctness has been studied in multiple generation tasks such as summarization (Maynez et al., 2020; Lux et al., 2020; Kryscinski et al., 2020; Nan et al., 2021; Zhou et al., 2020), document-level generation (Massarelli et al., 2020) and natural language generation (NLG) (Dušek and Kasner, 2020; Thomson and Reiter, 2020). While factual correctness has been studied in knowledge-grounded response generation (Li et al., 2019; Shuster et al., 2021; Honovich et al., 2021; Rashkin et al., 2021), there has been very limited work in creating a human annotated dataset and detecting factual consistency for open-domain dialog systems.

In this work, we conduct a thorough study to evaluate how factually consistent neural response generation models are with respect to the provided knowledge sentence or text retrieved from a knowledge base. We include a variety of GPT2-based models with different sizes, decoding strategies and quality of knowledge sentences. For our annotation of factual correctness, we categorize responses into three types as shown in Figure 1: non-verifiable responses that don’t contain information that need to be verified, factually consistent responses with respect to the input knowledge, and hallucinated responses that contain information not found in the input knowledge. Such annotations are different from previous work in that we are dealing with open-domain dialogs, where there are responses that don’t need to be evaluated for factual correctness. We form two datasets from these annotations: a large annotated dataset from public crowdworkers that labeled generated responses across different knowledge retrieval, decoding and model setups, and a smaller annotated dataset from internal crowdworkers to improve annotation quality.

After showing that state-of-the-art neural response generation models do produce factually inconsistent responses, we propose to build a factual consistency detection model. To train such a model, we created a corpus, Conv-FEVER, that includes the human written responses in the Wizard of Wikipedia dataset (Dinan et al., 2019) and

automatically generated factually inconsistent responses. We demonstrate competitive performance of our detector on human annotated data. Note that we make a distinction between factually **consistent** and factually **correct** responses. The former accurately portrays the input knowledge (assuming the provided knowledge input is correct), and the latter is accurate with respect to the “world knowledge”. Our detection model focuses specifically on factual consistency. Checking if a response is correct against “world knowledge” is an important problem that we leave for future exploration.

In summary, this is the first study with a thorough analysis of factual correctness for knowledge-grounded neural response generation models. In addition to the Conv-FEVER corpus, we are also releasing our annotated corpus of factual correctness on responses from multiple neural response generation models. We expect this study will increase awareness of the factual correctness problem in knowledge-grounded dialog systems and these datasets can help benchmark progress in this area.

## 2 Related Work

Previous work in detecting factual consistency have leveraged natural language inference datasets and models. Maynez et al. (2020) used a model trained on MNLI to detect consistency between abstractive summaries and their source documents. Dušek and Kasner (2020) also applied a model trained on MNLI to detect hallucinations and omissions for data-to-text generation. However, the MNLI dataset consists of sentences from a document while our work focuses on conversations.

Different decoding mechanisms have also been explored to improve factual correctness in generative models. Massarelli et al. (2020) found that delayed beam-search helps increase factual correctness for document level generation. Li et al. (2019) proposed a two-pass approach to generate factually consistent responses for knowledge-grounded response generation. Dziri et al. (2021) queried a knowledge graph using an entity retrieval model to improve factual correctness in response generation. Shuster et al. (2021) studied how different knowledge retrieval setups affect factual correctness in generated responses. However, their work focused on comparing factual correctness in non-knowledge grounded versus knowledge-grounded response generation models. Our work focuses purely on knowledge-grounded response genera-

tion. Additionally, we take into account the quality of the provided knowledge and annotate responses using different knowledge retrieval setups.

There has also been previous work on generating synthetic data to train a factual consistency detector. Kryscinski et al. (2020) generated factually incorrect examples by using techniques such as entity swapping and negation. This data was used to train a factual detector on generated summaries. Schuster et al. (2021a) generated negative examples for sentence level verification. We look at generating factually incorrect examples as well; however, our work focuses on factual consistency for the conversation domain.

Previous work also includes studies on human evaluation setups for factual correctness. Thomson and Reiter (2020) defined a methodology to annotate factual correctness for the data-to-text task. Mielke et al. (2020) proposed an annotation schema for the correctness and confidence of a model’s response, and calibrated the response such that its correctness aligns with its confidence. However, their work focused on a Q/A dataset and on knowledge that a model has learned during training.

Most recently, Honovich et al. (2021) proposed to use question generation/answering to evaluate factual consistency for knowledge-grounded dialogs. They annotated two model outputs for factual correctness. Within their annotation schema, they asked evaluators to ignore responses that are uninformative or chit-chat. Our annotation schema is similar in that we are also annotating if responses need to be verified before annotating for factual consistency. However, our dataset includes annotations for a diverse set of responses, generated across different knowledge retrieval, decoding and model setups to demonstrate the performance of different systems. We will be releasing our human annotations that consists of both if a sentence is verifiable and if it is factually consistent or hallucinates knowledge. We also propose to use different factual consistency detection approaches.

### 3 Are Neural Responses Consistent With Grounded Knowledge?

Knowledge-grounded dialog systems leverage knowledge snippets to have engaging conversations with users. In this study, we first evaluate how factually consistent the responses are in such systems. We identify three types of responses within open-domain dialogs, as shown in Figure 1.

**Non-verifiable responses** do not leverage the knowledge in a way that it needs to be verified. In Figure 1, the *Non-verifiable Response* extracts the entities *Federer* and *Nadal* from the *Knowledge*, but expresses the information about the entities as an opinion. Table 1 shows more examples of different types of non-verifiable responses.

**Factually consistent responses** contain detailed information and do not contradict the *Knowledge*. In Figure 1, the response implies that *Rafael Nadal* is a part of the Big Four. This requires verification to see if this information is valid based on the given knowledge.

**Hallucinated responses** are not consistent with the information provided in the knowledge. In Figure 1, the *Hallucinated Response* example leverages part of the *Knowledge* by including the entity *Djokovic*, but also hallucinates information (shown in the highlighted text). This piece of information is not present in the knowledge, and is not clear if it is correct. Some of the hallucination may be correct with respect to the “world knowledge”. This is similar to the hallucination problem in summarization and NLG tasks (Thomson and Reiter, 2020).

Responses can also contain a mixture of verifiable and non-verifiable spans. For the *Hallucinated Response* in Figure 1, the first sentence is an opinion on the entity while the second half is a factual statement on the same entity (not based on the knowledge sentence).

#### 3.1 Annotation Scheme

We propose a two-stage human annotation setup for factual correctness in the context of open-domain agents. The annotators are given a dialog context along with a knowledge sentence, and evaluate the system generated response. Stage 1 involves evaluating **Appropriateness** and **Verifiability** on a Likert scale of 1-5. For this stage we only show the dialog context and response. We define appropriateness as relevance to the dialog context and verifiability as to what degree the response needs to be verified. This stage is used to filter out incoherent and non-verifiable responses. If a response scores low on verifiability, it is categorized as a non-verifiable response.

Stage 2 of our setup involves evaluating **Factual Consistency** and **Hallucination**. We pose the Factual Consistency question: *Is the response generated factually accurate with regards to the input*

Type	Context	Response
Opinion	<b>B:</b> Ketchup and potato chips go well together too. <b>A:</b> I like the crinkle cut fries too.	I like poutine and chili cheese fries too.
Suggestion	<b>A:</b> If everybody liked the same things the world would be such a dull place. <b>B:</b> I'm glad our differences can make the world interesting.	I agree, and I'm sure you'd enjoy a good cup of coffee!
Topic Switch	<b>A:</b> i grew up in n.y. so i'm familiar with all that	I would love to visit the museums there.
Yes-No Question	<b>B:</b> How do you usually make pork steak, do you broil or pan fry them? <b>A:</b> I always fan fry them.	Do you like to flip them at all?
Personal information	<b>A:</b> Yes, in my area winter has a lot of snow and freezing temperatures. Do you have a cold winter where you are?	Yes, I do. I live in the Arctic region.

Table 1: Examples of different non-verifiable responses.

*knowledge?* with a three-point scale: factually incorrect (0), partially correct (0.5), and completely correct (1). For Hallucination, we pose the question: *Is the response generated making up more information than what is provided in the conversational context and input knowledge?*, with a binary label (Yes/No) for whether each response contains any hallucinated information. We present the annotation interface used for this collection in the Appendix. The example in Figure 1 shows the annotation for verifiability, factual consistency, and hallucination, for each type of responses.

### 3.2 Factual Correctness Study on Neural Response Generators

To quantitatively present the issues of factual correctness in open-domain dialog systems, we leverage our human annotation setup to label a set of outputs from different neural response generators. We use public crowd-workers in order to annotate a large amount of data under different configurations, namely model size, knowledge retrieval, and decoding setups. We vary knowledge retrieval setups as factual correctness in a response can vary depending on how accurate a knowledge retrieval model is (Schuster et al., 2021b). We denote this dataset as *Factual dataset-crowd*. We leverage Wizard of Wikipedia (WoW) (Dinan et al., 2019), a knowledge grounded dialog dataset generated through MTurkers who played the role of wizard and apprentice. The wizard has access to Wikipedia pas-

sages and the apprentice is given the role of learning more about a topic by engaging in a dialog with the wizard. At every turn the wizard selects a knowledge sentence from the Wikipedia passages to generate a knowledge-grounded turn.

For the response generation models, we leverage GPT2 (Radford et al., 2019) and follow the end-to-end-approach of (Dinan et al., 2019; Gopalakrishnan et al., 2019) and take in as input both the dialog context and a knowledge sentence and minimize cross-entropy loss on the ground-truth response. We fine-tune our models in a TransferTransfo fashion (Wolf et al., 2019) on the WoW dataset. The dialog context and the knowledge sentence are represented with the pretrained embeddings from the GPT2 model whose vocabulary is BPE tokenized. Hyperparameters are provided in the Appendix.

To evaluate factual consistency of different response generators, we vary the generator from three aspects: GPT2 model size, provided knowledge, and decoding strategy.

**Model size:** We use the four variants of GPT2 (Radford et al., 2019) based language models: small, medium, large, XL.

**Decoding strategies:** For decoding we use Nucleus sampling (NS) (Holtzman et al., 2020) with  $p=0.9$ ; Beam-Search (BS) with a beam size of 5, and Delayed Beam-Search (DBS) (Massarelli et al., 2020), which uses top-k sampling for  $n$  delay steps followed by beam search. We use  $k=10$ , 5 delay

Retriever	Decoding	GPT2-Small			GPT2-Medium			GPT2-Large			GPT2-XL		
		F ↑	H ↓	A ↑	F ↑	H ↓	A ↑	F ↑	H ↓	A ↑	F ↑	H ↓	A ↑
GT	BS	0.78	17.2%	3.9	0.80	18.4%	3.7	0.82	19.1%	<b>4.1</b>	0.84	8.0%	3.8
	NS	0.80	19.0%	3.8	0.84	11.6%	4.0	<b>0.88</b>	8.5%	4.0	<b>0.88</b>	<b>3.2%</b>	3.9
	DBS	0.81	22.2%	3.8	0.82	13.8%	4.0	0.81	11.7%	<b>4.1</b>	0.83	8.0%	3.8
KNN	BS	0.72	25.4%	3.6	0.74	21.7%	3.5	0.72	28.0%	3.6	0.77	22.7%	3.7
	NS	0.70	32.5%	3.4	<b>0.83</b>	<b>16.0%</b>	3.6	0.74	17.5%	<b>3.8</b>	0.75	28.0%	3.7
	DBS	0.71	21.4%	3.5	0.69	20.0%	3.7	0.70	28.3%	3.6	0.70	22.0%	<b>3.8</b>
DPR	BS	0.65	30.5%	3.1	0.64	26.8%	<b>3.5</b>	0.69	30.7%	3.4	0.62	35.1%	3.4
	NS	0.61	43.2%	3.1	0.65	30.2%	3.4	0.7	24.4%	<b>3.5</b>	0.67	27.0%	3.4
	DBS	0.59	39.0%	3.2	0.62	35.6%	<b>3.5</b>	0.7	24.4%	3.4	<b>0.71</b>	<b>18.2%</b>	3.3

Table 2: Human annotation results on data *Factual Dataset-crowd*. F = Factual Consistency score (0-1), H = Hallucination score (% of responses labeled “Yes”). A = Appropriateness (1-5)

steps and a beam size of 5.

**Knowledge:** we use three types of configurations for the knowledge sentence:

**(a) Ground Truth knowledge (GT):** This is the knowledge used by the “Wizards” to generate their response. It is explicitly available in the WoW test set and it is also referred to as the checked sentence in the WoW data.

**(b) K-Nearest Neighbours (K-NN)** : In this configuration, we use a  $K$  nearest neighbor search across all the 5 million Wikipedia articles used in the WoW dataset. In order to construct the knowledge base (**E**), we encode each individual sentence of the Wikipedia articles by leveraging the sentence transformers library (Reimers and Gurevych, 2019). Given an input dialog, we compute the representation for each turn using an encoder model. The model used to encode the knowledge base & input dialog is a DistilBert (Sanh et al., 2019) model finetuned on the NLI and STSb tasks. We take the average of the encoded embeddings and perform kNN search using the FAISS library (Johnson et al., 2017) for efficiency. In our experiments, we retrieve the top most relevant sentence from the knowledge base.

**(c) Dense Passage Retriever (DPR)** : For the last configuration, we use an open domain DPR retriever model from Karpukhin et al. (2020). In the DPR model, the knowledge base consists of 21 million Wikipedia passages. For our experiments, we use the pretrained encoder model made available from Karpukhin et al. (2020) to encode the last turn of the dialog. Using the encoded turn, we retrieve

top most relevant passage from the knowledge base by performing a similarity search using FAISS.

To create our annotation dataset we first sample 100 dialogs of length 5 turns from the WoW test sets and obtain responses using each of the configurations described above, resulting in 3600 responses in total (100 dialogs \* 4 model sizes \* 3 knowledge setups \* 3 decoding mechanisms). To compute inter-annotator agreement we collect three annotations for each system generated response based on the setup described in Section 3.1<sup>1</sup>. After the completion of stage 1, we filter out responses whose appropriateness and verifiable scores are below 3 to ensure we have responses that contain some form of knowledge and are relevant to the dialog context. After filtering out inappropriate responses there were 2781 responses remaining. After filtering out non-verifiable responses from this set there were 1869 remaining responses to be annotated in stage 2, that is, 33% of appropriate responses were non-verifiable. This highlights the significant amount of non-verifiable responses within open-domain dialog. On average 52 responses were annotated during stage 2 for each configuration. We show the exact number responses annotated for each configuration in the Appendix.

For each response we compute the factual consistency score by taking the average across the three annotation scores. For hallucination we do majority voting of the three annotation scores for each response. We use Krippendorff’s alpha for inter-annotator agreement (IAA), and found IAA for Factual Consistency is 0.15 and 0.21 for Hallucina-

<sup>1</sup>We leverage Mechanical Turk for our annotation https://www.mturk.com/

tion indicating slight to fair agreement (Landis and Koch, 1977). In Section 5 we will create a smaller dataset with higher IAA.

Table 2 shows the annotation results for factual consistency and hallucination under different configurations. Our human annotation results show two insights: we see that none of these state-of-the-art neural generation models generate factually consistent responses all the time and second this rate changes depending on the chosen configuration. As models get larger, there is an increase in factual consistency and a decrease in hallucinated responses. In a more realistic setting where DPR is being used to retrieve knowledge from a larger knowledge base, the highest average factual consistency score is 0.71, and at the lowest 18.2% of responses have hallucinated information, indicating room for improvement in factual consistency. Additionally, we see that in the DPR setting, DBS performs well for larger models. Massarelli et al. (2020) showed DBS performed the best for factual verification across all model sizes; however, that was for document-level generation.

Table 2 shows the average appropriateness score from the Stage 1 annotation. We can see that using DPR results in lower appropriateness scores than using the ground truth knowledge as it can return irrelevant knowledge sentences. Using NN has higher appropriateness scores than DPR since the size of knowledge base in NN is smaller (it is from the same set of Wikipedia articles used in WoW), and thus it is more likely to return a relevant knowledge sentence than DPR. These results show that the input knowledge quality can affect both the response quality and the factual consistency in model responses.

## 4 Factual Consistency Detector

After demonstrating the factual correctness issues in knowledge-grounded neural response generators, we develop a model to automatically detect the factual consistency of a generated response with respect to the knowledge fed into a response generation model. Since there are no existing labeled dialog datasets to train such detection models, we first develop a new dataset, Conv-FEVER, described below, and finetune our model using this dataset. We leverage the HuggingFace repo (Wolf et al., 2020) for all of our experiments.

### 4.1 Conv-FEVER Corpus

We create the Conv-FEVER dataset from the Wizard of Wikipedia corpus (Dinan et al., 2019) consisting of dialogs between paired crowd sourced workers that played the role of wizard and apprentice in a controlled environment. The conversations between the wizard and apprentice are around particular topics, and the wizard’s responses are based on the knowledge sentence chosen from Wikipedia documents about the topic. Therefore we hypothesize that the *responses generated by the wizard are consistent with respect to the provided knowledge*. To generate inconsistent responses, we leverage a few data augmentation strategies introduced in (Kryscinski et al., 2020), including random pairing, negation, entity swapping. Every data point in the Conv-FEVER dataset contains a Conversational Context ( $C_n$ ), Knowledge ( $K_n$ ), Wizard’s Response or Target Response ( $R_n$ ), where  $n$  represents the size of the dataset. Table 3 shows the statistics of the data set. Figure 2 shows examples of these data augmentation techniques.

Dataset	Num.	
	Consistent	Inconsistent
WoW	68957	-
Random Pairing	-	137914
Negation	-	107845
Entity	-	73178

Table 3: Conv-FEVER dataset statistics

**Random Pairing** : we perform two types of random pairing when given a conversational context ( $C_i$ ), knowledge ( $K_i$ ), wizard’s response ( $R_i$ ): (1) we replace the response ( $R_i$ ) associated with conversation ( $C_i$ ) with a response ( $R_j$ ) from a random conversation ( $C_j$ ) where  $j \neq i$ ; (2) We replace the knowledge ( $K_i$ ) associated with conversation ( $C_i$ ) with a Knowledge ( $K_j$ ) from a random conversation ( $C_j$ ) where  $j \neq i$ .

**Negation** : we perform two types of negations: (1) negation on the response whilst keeping the context and knowledge untouched; (2) negation applied to the knowledge whilst keeping the context and response untouched. In both the cases, we performed negation only if certain verbs<sup>2</sup> were present

<sup>2</sup>We apply negation for these tokens: are, is, was, were, have, has, had, do, does, did, can, ca, could, may, might, must, shall, should, will, would.

---

*Dialog History:*

**Speaker 1:** I couldn't imagine living when internet access was rare

**Speaker 2:** Oh me either! It seems like such a long time ago. I wonder when Internet was first created?

---

*Knowledge:*

Use by a wider audience only came in 1995 when restrictions on the use of the Internet to carry commercial traffic were lifted.

---

*Response:*

It used to be restricted but around 1995, the restrictions were lifted and commercial use of it began

---

*Negated Response:*

It used to be restricted but around 1995, the restrictions **were**n't lifted and commercial use of it began

---

*Entity Swap Knowledge:*

Use by a wider audience only came in **1971** when restrictions on the use of the Internet to carry commercial traffic were lifted.

---

Figure 2: Data augmentation examples in ConvFEVER corpus. Red text highlight shows the changes made to the original sentence.

in the response or knowledge. Further as a part of negation, we added or replaced the words with “not or n't” as part of negating the sentence.

**Entity Swapping** : We performed this data augmentation to ensure that the dataset contains representation of the most common mistakes made by current state of the art generation models. To perform entity swapping, we used SpaCY NER tagger (Honnibal et al., 2020) over the conversational context ( $C_i$ ) and knowledge ( $K_i$ ) to create two indexes of all entities mentioned in context ( $E_c$ ) and knowledge ( $E_k$ ) respectively. We then performed two types of entity swapping: (1) we performed entity swapping only on the context if there is a common entity ( $Y$ ) mentioned in the context and the response. In this case, we randomly sampled an entity from the context index ( $E_c$ ) and replaced the common entity ( $Y$ ); (2) we performed entity swapping only on the knowledge if there is a common entity ( $Y$ ) mentioned in the knowledge and the response. In this case, we randomly sampled an entity from the knowledge index ( $E_k$ ) and replaced the common entity ( $Y$ ).

---

*Dialog History:*

**Speaker 1:** Well I mainly make pecan because it is so delicious and usually includes vanilla, and salt to balance the taste

**Speaker 2:** i didnt know vanilla was added into it

---

*Knowledge:* Pecan pie Pecan pie is often served with whipped cream, vanilla ice cream, or hard sauce.

---

*Response:* Oh thats interesting! I love Pecan pie, it is often served with whipped cream, vanilla ice cream, or hard sauce.

---

*Ground-truth:* Consistent

*FactCC model:* Consistent

---

*ALBERT-base-VitaminC model:* Not Enough Info

---

*ALBERT-base-VitaminC-FEVER model:* Not Enough Info

---

*ALBERT-base-MNLI model:* Not Enough Info

---

*ALBERT-base-VitaminC-MNLI model:* Not Enough Info

---

*Bert-base-Conv-FEVER:* Consistent

---

Figure 3: Responses can contain spans of verifiable (things pecan pie is served with) and non verifiable (liking pecan pie) information. In this example our detector predicts correctly while the NLI models choose not enough info due to the extra information in the response.

## 4.2 Factual Consistency Detector Models

### 4.2.1 Baseline: ALBERT-base

For our NLI baseline models we use out-of-the box ALBERT-base (Lan et al., 2019) models<sup>3</sup> that were finetuned on three separate datasets: FEVER (Thorne et al., 2018), MNLI dataset (Williams et al., 2018) and VitaminC (Schuster et al., 2021a). The MNLI task is an entailment task with three labels: entailment, neutral, and contradiction. The FEVER and VitaminC tasks are used to verify whether claims are supported by a reference piece of knowledge and have three labels: supports, not enough info, and refutes. There is a clear mapping between the labels in such two tasks, i.e., supports corresponds to entailment, not enough info to neutral, and refutes to contradiction. Therefore we use these three labels: supports, not enough info, and refutes, from these finetuned baseline models. The models take in the knowledge

<sup>3</sup><https://huggingface.co/tals>

Retriever	Decoding	GPT2-Small		GPT2-Medium		GPT2-Large		GPT2-XL	
		F1(C)	F1(IC)	F1(C)	F1(IC)	F1(C)	F1(IC)	F1(C)	F1(IC)
GT	BS	86.7	0.0	87.0	0.0	86.6	10.5	91.2	0.0
	NS	87.1	13.3	90.3	14.3	93.1	0.0	<b>97.5</b>	0.0
	DBS	87.9	31.6	89.9	0.0	93.3	50.0	93.8	<b>36.4</b>
KNN	BS	80.5	26.1	82.7	23.5	80.0	20.0	78.9	11.8
	NS	81.9	40.0	87.1	26.7	<b>87.2</b>	40.0	78.2	29.6
	DBS	82.4	11.8	86.8	35.3	76.5	24.0	83.1	<b>43.5</b>
DPR	BS	72.0	36.4	71.2	26.1	81.4	42.1	78.4	52.2
	NS	66.7	26.1	78.8	30.0	80.6	22.2	<b>84.6</b>	<b>53.9</b>
	DBS	69.4	34.8	67.9	32.0	82.9	40.0	76.0	25.0

Table 4: Results on *Factual dataset-crowd* for BERT-base Conv-FEVER. F1 (C): F1 (Consistent) , F1 (IC): F1 (Inconsistent), BS: Beam-Search, NS: Nucleus Sampling, DBS: Delayed Beam-Search

Model	Consistent			Inconsistent		
	P	R	F1	P	R	F1
FactCC	81.3	73.0	77.0	30.1	41.1	34.7
ALBERT + VitC	93.5	50.7	65.8	33.5	87.6	48.5
+ FEVER	92.7	53.4	67.8	34.1	85.2	<b>48.8</b>
ALBERT + MNLI	<b>95.6</b>	43.1	59.4	31.6	<b>93.0</b>	47.2
+ VitC	93.2	48.5	63.8	32.5	87.6	47.5
BERT-base CONV-FEVER	78.7	<b>99.8</b>	<b>88.0</b>	<b>85.7</b>	4.60	8.80

Table 5: Results on *Factual dataset-expert*. P=Precision, R=Recall. VitC = VitaminC dataset.

along with the response, and predict one of the three output labels. To match our two classes for factual consistency, we consider the *supports* label as factually consistent, and bucket *not enough info* and *refutes* together as factually inconsistent. We choose to combine *not enough info* examples under factually inconsistent because these are typically cases where the response contains hallucinated information unrelated to the knowledge that does not directly contradict or refute the knowledge.

#### 4.2.2 Baseline: FactCC

FactCC (Kryscinski et al., 2020), is a BERT-based model used to predict factual consistency for neural summarization models. The model was trained on synthesized negative examples using the same data augmentation methods as described in Section 4.1. For our experiments we use the open-sourced version of the model.<sup>4</sup>

#### 4.2.3 BERT-base-Conv-FEVER

We propose a factual consistency detector using our automatically collected data Conv-FEVER. We use the BERT-base model (Devlin et al., 2019) as

<sup>4</sup><https://github.com/salesforce/factCC>

the backbone, and initialize it by first training on the FEVER dataset (Thorne et al., 2018) taken from the set of tasks presented in KILT (Petroni et al., 2020). The FEVER task is aimed at determining if a claim can be supported or refuted given a Wikipedia document. Claims that are labeled as supports can be thought of to be consistent and refutes can be thought of to be inconsistent. To create the initial training corpus, we extracted all data points in the FEVER corpus that contained a pointer to the ground truth Wikipedia documents as knowledge. In total we trained on 48,451 supports and 18,625 refutes<sup>5</sup>. Because FEVER is not a conversational dataset, we further finetune the model on Conv-FEVER. The set of hyperparameters used are shown in the Appendix.

## 5 Factual Consistency Detector Evaluation

Different knowledge retrieval models can greatly affect response quality as shown in Table 2. Therefore we create a *Factual dataset-expert* that only consists of the ground truth knowledge, allowing us to assume the selected knowledge is correct. Additionally due to the lower inter-annotator agreement on *Factual dataset-crowd*, we annotate *Factual dataset-expert* with three internal annotators.

For this dataset we generate 1000 outputs from the fine-tuned GPT2-small model, using the ground-truth knowledge and nucleus sampling with  $p=0.9$  and follow the same annotation setup described in Section 3.1. After filtering responses based on the stage 1 annotations, 584 responses were kept, meaning that 41.6% were deemed as

<sup>5</sup>This number differs from the original since we dropped data points without pointer to Wikipedia documents.

non-verifiable responses, following a similar pattern as *Factual dataset-crowd*. After stage 2, the inter-annotator agreement scores using Krippendorff’s alpha are 0.80 and 0.86 on factual consistency and hallucination respectively indicating substantial to almost perfect agreement (Landis and Koch, 1977). We compute factual consistency for each response by averaging the three scores from the internal annotators. We bucket the generated responses such that a factual consistency score  $\geq 0.5$  is labeled as consistent and  $< 0.5$  is labeled as inconsistent. This results in 502 responses labeled as consistent and 82 responses labeled as inconsistent.

Table 5 shows the factual detector’s performance on *Factual dataset-expert*. We see our model trained on Conv-FEVER outperforms just training on a document-level dataset on F1 for the consistent class, while NLI models outperform our model on the inconsistent class. This is because the NLI models predict the *not enough info* class 58% of the time, which we bucket into the inconsistent label.

To view our detector’s performance under different knowledge retrieval configurations we run our BERT-base Conv-FEVER model on *Factual dataset-crowd*. We bucket our responses in the same way as *Factual dataset-expert*, resulting in 1618 consistent responses and 251 inconsistent responses across all configurations.

Table 4 shows that we achieve similar F1-scores for responses generated by GPT2-small using nucleus sampling with ground truth knowledge, which is the same configuration as *Factual dataset-expert*. We see the general trend that the overall scores improve for responses generated by larger models but decrease when responses are not generated using the ground truth knowledge. The detector’s performance is in line with the appropriateness scores shown in Table 2 where the detector performed better when the model generated more relevant responses. This further motivates the need to first filter out non-relevant responses before annotating for factual correctness as it is much harder to determine if a response is factually consistent if it is not first appropriate for that dialog context. Figure 3 shows a qualitative example from our models. We present additional examples in the Appendix.

## 6 Conclusions and Future Work

In this work, we present a case study for factual correctness for knowledge-grounded response generation. We propose a human annotation setup to

identify non-verifiable, factually consistent, and hallucinated responses and use this setup to annotate responses from multiple neural generative models. Furthermore, we propose a detector to identify if a response is factually consistent with its respective input knowledge. We create a new dataset called Conv-FEVER to train a factual consistency detector and through evaluation show the benefit of this dataset. Both Conv-FEVER and our human annotated data will be released. Since our factual detector model relies heavily on pretrained models to learn representations for both knowledge and responses, we expect it can generalize well to other knowledge-grounded dialog datasets.

The factual consistency detector can be used to automatically evaluate system responses for factual correctness. Therefore our future work involves using factual consistency information as feedback to improve response generation models or utilizing the detector to re-rank responses from a response generator. Additionally we will look at evaluating factual correctness against the “world knowledge”.

## References

- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. [BERT: Pre-training of deep bidirectional transformers for language understanding](#). In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 4171–4186, Minneapolis, Minnesota. Association for Computational Linguistics.
- Emily Dinan, Stephen Roller, Kurt Shuster, Angela Fan, Michael Auli, and Jason Weston. 2019. [Wizard of wikipedia: Knowledge-powered conversational agents](#). In *7th International Conference on Learning Representations, ICLR 2019, New Orleans, LA, USA, May 6-9, 2019*. OpenReview.net.
- Ondřej Dušek and Zdeněk Kasner. 2020. [Evaluating semantic accuracy of data-to-text generation with natural language inference](#). In *Proceedings of the 13th International Conference on Natural Language Generation*, pages 131–137, Dublin, Ireland. Association for Computational Linguistics.
- Nouha Dziri, Andrea Madotto, Osmar Zaiane, and Avishek Joey Bose. 2021. Neural path hunter: Reducing hallucination in dialogue systems via path grounding. *arXiv preprint arXiv:2104.08455*.
- Marjan Ghazvininejad, Chris Brockett, Ming-Wei Chang, Bill Dolan, Jianfeng Gao, Wen-tau Yih, and Michel Galley. 2018. A knowledge-grounded neural conversation model. In *Thirty-Second AAAI Conference on Artificial Intelligence*.

- Karthik Gopalakrishnan, Behnam Hedayatnia, Qiang Chen, Anna Gottardi, Sanjeev Kwatra, Anu Venkatesh, Raefer Gabriel, and Dilek Hakkani-Tür. 2019. Topical-chat: Towards knowledge-grounded open-domain conversations. *Proc. Interspeech 2019*, pages 1891–1895.
- Ari Holtzman, Jan Buys, Li Du, Maxwell Forbes, and Yejin Choi. 2020. [The curious case of neural text degeneration](#). In *8th International Conference on Learning Representations, ICLR 2020, Addis Ababa, Ethiopia, April 26-30, 2020*. OpenReview.net.
- Matthew Honnibal, Ines Montani, Sofie Van Landeghem, and Adriane Boyd. 2020. [spaCy: Industrial-strength Natural Language Processing in Python](#).
- Or Honovich, Leshem Choshen, Roei Aharoni, Ella Neeman, Idan Szpektor, and Omri Abend. 2021. [Q2: Evaluating factual consistency in knowledge-grounded dialogues via question generation and question answering](#). *arXiv preprint arXiv:2104.08202*.
- Jeff Johnson, Matthijs Douze, and Hervé Jégou. 2017. Billion-scale similarity search with gpus. *arXiv preprint arXiv:1702.08734*.
- Vladimir Karpukhin, Barlas Oguz, Sewon Min, Patrick Lewis, Ledell Wu, Sergey Edunov, Danqi Chen, and Wen-tau Yih. 2020. [Dense passage retrieval for open-domain question answering](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 6769–6781, Online. Association for Computational Linguistics.
- Wojciech Kryscinski, Bryan McCann, Caiming Xiong, and Richard Socher. 2020. [Evaluating the factual consistency of abstractive text summarization](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 9332–9346, Online. Association for Computational Linguistics.
- Zhenzhong Lan, Mingda Chen, Sebastian Goodman, Kevin Gimpel, Piyush Sharma, and Radu Soricut. 2019. Albert: A lite bert for self-supervised learning of language representations. *arXiv preprint arXiv:1909.11942*.
- J Richard Landis and Gary G Koch. 1977. The measurement of observer agreement for categorical data. *biometrics*, pages 159–174.
- Zekang Li, Cheng Niu, Fandong Meng, Yang Feng, Qian Li, and Jie Zhou. 2019. [Incremental transformer with deliberation decoder for document grounded conversations](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 12–21, Florence, Italy. Association for Computational Linguistics.
- Robert Logan, Nelson F. Liu, Matthew E. Peters, Matt Gardner, and Sameer Singh. 2019. [Barack’s wife hillary: Using knowledge graphs for fact-aware language modeling](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 5962–5971, Florence, Italy. Association for Computational Linguistics.
- Klaus-Michael Lux, Maya Sappelli, and Martha Larson. 2020. [Truth or error? towards systematic analysis of factual errors in abstractive summaries](#). In *Proceedings of the First Workshop on Evaluation and Comparison of NLP Systems*, pages 1–10, Online. Association for Computational Linguistics.
- Luca Massarelli, Fabio Petroni, Aleksandra Piktus, Myle Ott, Tim Rocktäschel, Vassilis Plachouras, Fabrizio Silvestri, and Sebastian Riedel. 2020. [How decoding strategies affect the verifiability of generated text](#). In *Findings of the Association for Computational Linguistics: EMNLP 2020*, pages 223–235, Online. Association for Computational Linguistics.
- Joshua Maynez, Shashi Narayan, Bernd Bohnet, and Ryan McDonald. 2020. [On faithfulness and factuality in abstractive summarization](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 1906–1919, Online. Association for Computational Linguistics.
- Sabrina J Mielke, Arthur Szlam, Y-Lan Boureau, and Emily Dinan. 2020. Linguistic calibration through metacognition: aligning dialogue agent responses with expected correctness. *arXiv preprint arXiv:2012.14983*.
- Feng Nan, Cicero Nogueira dos Santos, Henghui Zhu, Patrick Ng, Kathleen McKeown, Ramesh Nallapati, Dejiao Zhang, Zhiguo Wang, Andrew O Arnold, and Bing Xiang. 2021. Improving factual consistency of abstractive summarization via question answering. *arXiv preprint arXiv:2105.04623*.
- Fabio Petroni, Aleksandra Piktus, Angela Fan, Patrick Lewis, Majid Yazdani, Nicola De Cao, James Thorne, Yacine Jernite, Vassilis Plachouras, Tim Rocktäschel, et al. 2020. Kilt: a benchmark for knowledge intensive language tasks. *arXiv preprint arXiv:2009.02252*.
- Fabio Petroni, Tim Rocktäschel, Sebastian Riedel, Patrick Lewis, Anton Bakhtin, Yuxiang Wu, and Alexander Miller. 2019. [Language models as knowledge bases?](#) In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 2463–2473, Hong Kong, China. Association for Computational Linguistics.
- Alec Radford, Jeffrey Wu, Rewon Child, David Luan, Dario Amodei, and Ilya Sutskever. 2019. Language models are unsupervised multitask learners. *OpenAI blog*, 1(8):9.

- Hannah Rashkin, David Reitter, Gaurav Singh Tomar, and Dipanjan Das. 2021. Increasing faithfulness in knowledge-grounded dialogue with controllable features. *arXiv preprint arXiv:2107.06963*.
- Nils Reimers and Iryna Gurevych. 2019. [Sentence-BERT: Sentence embeddings using Siamese BERT-networks](#). In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 3982–3992, Hong Kong, China. Association for Computational Linguistics.
- Adam Roberts, Colin Raffel, and Noam Shazeer. 2020. [How much knowledge can you pack into the parameters of a language model?](#) In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 5418–5426, Online. Association for Computational Linguistics.
- Stephen Roller, Y-Lan Boureau, Jason Weston, Antoine Bordes, Emily Dinan, Angela Fan, David Gunning, Da Ju, Margaret Li, Spencer Poff, et al. 2020a. Open-domain conversational agents: Current progress, open problems, and future directions. *arXiv preprint arXiv:2006.12442*.
- Stephen Roller, Emily Dinan, Naman Goyal, Da Ju, Mary Williamson, Yinhan Liu, Jing Xu, Myle Ott, Kurt Shuster, Eric M Smith, et al. 2020b. Recipes for building an open-domain chatbot. *arXiv preprint arXiv:2004.13637*.
- Victor Sanh, Lysandre Debut, Julien Chaumond, and Thomas Wolf. 2019. Distilbert, a distilled version of bert: smaller, faster, cheaper and lighter. *arXiv preprint arXiv:1910.01108*.
- Tal Schuster, Adam Fisch, and Regina Barzilay. 2021a. [Get your vitamin c! robust fact verification with contrastive evidence](#). In *NAACL 2021*.
- Tal Schuster, Adam Fisch, and Regina Barzilay. 2021b. [Get your vitamin c! robust fact verification with contrastive evidence](#). *arXiv preprint arXiv:2103.08541*.
- Kurt Shuster, Spencer Poff, Moya Chen, Douwe Kiela, and Jason Weston. 2021. Retrieval augmentation reduces hallucination in conversation. *arXiv preprint arXiv:2104.07567*.
- Craig Thomson and Ehud Reiter. 2020. [A gold standard methodology for evaluating accuracy in data-to-text systems](#). In *Proceedings of the 13th International Conference on Natural Language Generation*, pages 158–168, Dublin, Ireland. Association for Computational Linguistics.
- James Thorne, Andreas Vlachos, Christos Christodoulopoulos, and Arpit Mittal. 2018. [FEVER: a large-scale dataset for fact extraction and VERification](#). In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long Papers)*, pages 809–819, New Orleans, Louisiana. Association for Computational Linguistics.
- Adina Williams, Nikita Nangia, and Samuel Bowman. 2018. [A broad-coverage challenge corpus for sentence understanding through inference](#). In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long Papers)*, pages 1112–1122. Association for Computational Linguistics.
- Thomas Wolf, Lysandre Debut, Victor Sanh, Julien Chaumond, Clement Delangue, Anthony Moi, Pierric Cistac, Tim Rault, Rémi Louf, Morgan Funtowicz, Joe Davison, Sam Shleifer, Patrick von Platen, Clara Ma, Yacine Jernite, Julien Plu, Canwen Xu, Teven Le Scao, Sylvain Gugger, Mariama Drame, Quentin Lhoest, and Alexander M. Rush. 2020. [Transformers: State-of-the-art natural language processing](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing: System Demonstrations*, pages 38–45, Online. Association for Computational Linguistics.
- Thomas Wolf, Victor Sanh, Julien Chaumond, and Clement Delangue. 2019. Transfertransfo: A transfer learning approach for neural network based conversational agents. *arXiv preprint arXiv:1901.08149*.
- Semih Yavuz, Abhinav Rastogi, Guan-Lin Chao, and Dilek Hakkani-Tur. 2019. [DeepCopy: Grounded response generation with hierarchical pointer networks](#). In *Proceedings of the 20th Annual SIGdial Meeting on Discourse and Dialogue*, pages 122–132, Stockholm, Sweden. Association for Computational Linguistics.
- Chunting Zhou, Graham Neubig, Jiatao Gu, Mona Diab, Paco Guzman, Luke Zettlemoyer, and Marjan Ghazvininejad. 2020. Detecting hallucinated content in conditional neural sequence generation. *arXiv preprint arXiv:2011.02593*.
- Hao Zhou, Tom Young, Minlie Huang, Haizhou Zhao, Jingfang Xu, and Xiaoyan Zhu. 2018. Commonsense knowledge aware conversation generation with graph attention. In *IJCAI*, pages 4623–4629.

## A Appendices

### A.1 Qualitative Examples and Analysis

Qualitatively, we find that all models is able to correctly identify cases where the response states the fact clearly that is present in the knowledge.

**Context :**

**Apprentice:** *I love going for hikes in nature for exercise.*

**Wizard:** *It is one of my favorite hobbies as well, Hiking is the preferred term, in Canada and the United States*

**Apprentice:** *What do they call it elsewhere?*

**Wizard:** *in the United Kingdom, and the Republic of Ireland, the word "walking" is acceptable to describe "all forms" of walking*

**Apprentice:** *Well that's interesting, but seems like it could get confusing! I bet humans have been hiking ever since we could walk though!*

**Knowledge:** *Hiking The term bushwalking is endemic to Australia, having been adopted by the Sydney Bush Walkers club in 1927.*

**Response:** *Hiking The term bushwalking is endemic to Australia, having been adopted by the Sydney Bush Walkers club in 1927.*

**Human Annotation:** *Consistent*

**FactCC model prediction:** *Consistent*

**ALBERT-base-VitaminC model prediction:** *Consistent*

**ALBERT-base-VitaminC-FEVER model prediction:** *Consistent*

**ALBERT-base-MNLI model prediction:** *Consistent*

**ALBERT-base-VitaminC-MNLI model prediction:** *Consistent*

**Bert-base-Conv-FEVER prediction:** *Consistent*

Responses can contains spans of verifiable (things pecan pie is served with) and non verifiable (liking pecan pie) information. In the example below the our detector predicts correctly while the NLI models choose not enough info due to the extra information in the response.

**Context :**

**Wizard:** *I like to bake pecan pie it is east to make with pecans, eggs, butter, and sugar*

**Conversation Context:**  
\${context}

**Response:**  
\${responses}

On a scale of 1-5, how appropriate is the response with respect to the conversation context?  
 1  2  3  4  5

On a scale of 1-5, how informative is the response?  
 1  2  3  4  5

Figure 4: Interface provided for Stage 1

**Conversation Context:**  
\${context}

**External Knowledge:**  
\${knowledge}

**Response:**  
\${response}

Is the response Factually correct?  
 0  0.5  1

Did the response make up facts?  
 Yes  No

Figure 5: Interface provided for Stage 2

**Apprentice:** *Im allergic to pecans, unfortunately i cant eat those. i love apple pies though! those are delicious!*

**Wizard:** *that is terrible apple pie is so tart*

**Apprentice:** *What other flavors of pie are there?*

**Wizard:** *Well I mainly make pecan because it is so delicious and usually includes vanilla, and salt to balance the taste*

**Apprentice:** *i didnt know vanilla was added into it*

**Knowledge:** *Pecan pie Pecan pie is often served with whipped cream, vanilla ice cream, or hard sauce.*

**Response:** *Oh thats interesting! I love Pecan pie, it is often served with whipped cream, vanilla ice cream, or hard sauce.*

**Human Annotation:** *Consistent*

**FactCC model prediction:** *Consistent*

**ALBERT-base-VitaminC model prediction:** *Not Enough Info*

**ALBERT-base-VitaminC-FEVER model prediction:** *Not Enough Info*

**ALBERT-base-MNLI model prediction:** *Not Enough Info*

**ALBERT-base-VitaminC-MNLI model prediction:** *Not Enough Info*

**Bert-base-Conv-FEVER prediction:** *Consistent*

Example response showcasing hallucination.

NLI models predict Not Enough Info which can be thought of as a neutral class. Our detector predicts inconsistent which would be deemed correct as there are no statements in the response that contain any of the knowledge regardless of if it's correct or not.

**Context :**

*Apprentice:* I'm curious to know about Dylan's Candy Bar

*Wizard:* Dylan's candy bar is a chain of boutique candy shop

*Apprentice:* Who founded it?

*Wizard:* It is owned Dylan Lauren

*Apprentice:* What year was it founded?

*Knowledge:* Dylan's Candy Bar It is owned by Dylan Lauren, daughter of fashion designer Ralph Lauren.

*Response:* It was founded in the late 1800s

**Human Annotation:** Inconsistent

**FactCC model prediction:** Inconsistent

**ALBERT-base-VitaminC model prediction:** Not Enough Info

**ALBERT-base-VitaminC-FEVER model prediction:** Not Enough Info

**ALBERT-base-VitaminC-MNLI model prediction:** Not Enough Info

**ALBERT-base-MNLI model prediction:** Not Enough Info

**Bert-base-Conv-FEVER prediction:** Inconsistent

We find that our detector does not always capture incorrect dates in responses. The date 2005 does exist in the knowledge but for a different movie.

**Context :**

*Apprentice:* I'm curious to know about Dylan's Candy Bar

*Wizard:* Dylan's candy bar is a chain of boutique candy shop

*Apprentice:* Who founded it?

*Wizard:* It is owned Dylan Lauren

*Apprentice:* What year was it founded?

*Knowledge:* Jimmy Fallon He left the program for the film industry, starring in films such as "Taxi" (2004) and "Fever Pitch" (2005).

*Response:* I have read the book, I love "Taxi",

which was released in 2005.

**Human Annotation:** Inconsistent

**FactCC model prediction:** Inconsistent

**ALBERT-base-VitaminC model prediction:** Not Enough Info

**ALBERT-base-VitaminC-FEVER model prediction:** Not Enough Info

**ALBERT-base-VitaminC-MNLI model prediction:** Not Enough Info

**ALBERT-base-MNLI model prediction:** Inconsistent

**Bert-base-Conv-FEVER prediction:** Consistent

## A.2 Human Evaluation Interface

We include images of the interface provided to crowdworkers for our human annotation in figures 4 and 5.

## A.3 Response Generation Hyperparameters

We finetune the GPT2 model on the Wizard of Wikipedia dataset using dialog context and knowledge as input. We train the both the language modeling and multiple choice head of GPT2 in a TransferTransfo fashion where the loss from each head is weighted equally. We truncate the dialog context to use the past 128 tokens and we truncate the knowledge to use the first 64 tokens. We use ADAM as an optimizer with a learning rate of 6.25e-05 and a piecewise linear scheduler. We use a batch size of 2 and train for 10 epochs using perplexity for early stopping with a patience of 1.

## A.4 BERT-base Conv-FEVER model Hyperparameters

We finetune the BERT-base-uncased model on the Conv-FEVER dataset using knowledge and response as input. We use ADAM as an optimizer with a learning rate of 2e-05. We use a batch size of 16 examples and train for 4 epochs.

## A.5 Albert-base model Hyperparameters

All models were used out-of-the box and taken from the HuggingFace repo<sup>6</sup>. These models were trained for 50000 max steps using a batch size of 32 examples and a learning rate of 2e-5.

<sup>6</sup><https://huggingface.co/tals>

Retriever	Decoding	GPT2-Small		GPT2-Medium		GPT2-Large		GPT2-XL	
		F1(C)	F1(IC)	F1(C)	F1(IC)	F1(C)	F1(IC)	F1(C)	F1(IC)
Ground Truth	Beam Search	87.0	0.0	87.0	0.0	87.0	11.0	91.0	0.0
	Nucleus Sampling	87.0	13.0	90.0	14.0	93.0	0.0	97.0	0.0
	DBS	88.0	32.0	90.0	0.0	93.0	50.0	94.0	36.0
KNN	Beam Search	80.0	26.0	83.0	24.0	80.0	20.0	79.0	12.0
	Nucleus Sampling	82.0	40.0	87.0	27.0	87.0	40.0	83.0	43.0
	DBS	82.0	12.0	87.0	35.0	77.0	24.0	78.0	30.0
DPR	Beam Search	72.0	36.0	71.0	26.0	81.0	42.0	78.0	52.0
	Nucleus Sampling	67.0	26.0	79.0	30.0	81.0	22.0	85.0	54.0
	DBS	69.0	35.0	68.0	32.0	83.0	40.0	76.0	52.0

Table 6: Results on *Factual Dataset-crowd* for Bert-base-Conv-FEVER. F1 (C): F1 Consistent / F1(IC): F1 Inconsistent

Retriever	Decoding Strategy	GPT2-small	GPT2-medium	GPT2-large	GPT2-XL
Ground-Truth	Beam Search	64	65	73	62
	Nucleus Sampling	58	69	70	61
	DBS	63	65	68	62
KNN	Beam Search	55	46	50	44
	Nucleus Sampling	44	50	57	57
	DBS	51	50	53	50
DPR	Beam Search	36	41	39	37
	Nucleus Sampling	37	43	45	52
	DBS	36	39	45	33

Table 7: Number of responses annotated for Stage 2 under each knowledge retrieval, model, decoding condition

Retriever	Decoding	GPT2-Small		GPT2-Medium		GPT2-Large		GPT2-XL	
		F1(C)	F1(IC)	F1(C)	F1(IC)	F1(C)	F1(IC)	F1(C)	F1(IC)
Ground Truth	Beam Search	61.0	37.0	64.0	41.0	71.0	44.0	55.0	18.0
	Nucleus Sampling	68.0	47.0	65.0	30.0	58.0	27.0	72.0	13.0
	DBS	63.0	38.0	64.0	35.0	67.0	35.0	59.0	19.0
KNN	Beam Search	61.0	55.0	31.0	34.0	55.0	53.0	42.0	30.0
	Nucleus Sampling	42.0	54.0	61.0	44.0	63.0	47.0	60.0	60.0
	DBS	50.0	39.0	42.0	34.0	57.0	57.0	47.0	45.0
DPR	Beam Search	34.0	56.0	46.0	60.0	44.0	61.0	21.0	52.0
	Nucleus Sampling	38.0	67.0	49.0	53.0	44.0	44.0	39.0	52.0
	DBS	30.0	58.0	45.0	64.0	36.0	39.0	39.0	46.0

Table 8: Results on *Factual Dataset-crowd* for ALBERT-base MNLI. F1 (C): F1 Consistent / F1(IC): F1 Inconsistent

## A.6 FactCC model Hyperparameters

This model was taken out-of-the box <sup>7</sup>. The model was trained for 10 epochs using batch size of 12 examples and learning rate of 2e-5.

<sup>7</sup><https://github.com/salesforce/factCC>

## A.7 Factual Consistency Detector Evaluation

In table 6, 8, 9, 10, 11 and 12 we present results on *Factual Dataset-crowd* trained on various Bert or Albert based models trained on FEVER, ConvFever, MNLI, VitC or a combination of them and FactCC.

Retriever	Decoding	GPT2-Small		GPT2-Medium		GPT2-Large		GPT2-XL	
		F1(C)	F1 (IC)	F1(C)	F1 (IC)	F1(C)	F1 (IC)	F1(C)	F1 (IC)
Ground Truth	Beam Search	71.0	42.0	64.0	28.0	75.0	41.0	72.0	29.0
	Nucleus Sampling	74.0	47.0	77.0	41.0	71.0	29.0	76.0	15.0
	DBS	69.0	44.0	67.0	29.0	75.0	40.0	70.0	28.0
KNN	Beam Search	59.0	54.0	43.0	29.0	53.0	55.0	48.0	43.0
	Nucleus Sampling	46.0	55.0	62.0	33.0	51.0	55.0	67.0	59.0
	DBS	61.0	47.0	59.0	36.0	57.0	47.0	54.0	50.0
DPR	Beam Search	34.0	56.0	51.0	56.0	50.0	57.0	28.0	53.0
	Nucleus Sampling	50.0	43.0	41.0	62.0	50.0	43.0	52.0	56.0
	DBS	25.0	62.0	35.0	50.0	39.0	36.0	42.0	42.0

Table 9: Results on *Factual Dataset-crowd* for ALBERT-base-VitaminC. F1 (C): F1 Consistent / F1(IC): F1 Inconsistent

Retriever	Decoding	GPT2-Small		GPT2-Medium		GPT2-Large		GPT2-XL	
		F1(C)	F1 (IC)	F1(C)	F1 (IC)	F1(C)	F1 (IC)	F1(C)	F1 (IC)
Ground Truth	Beam Search	72.0	43.0	58.0	32.0	65.0	35.0	65.0	21.0
	Nucleus Sampling	82.0	55.0	70.0	36.0	71.0	21.0	74.0	7.0
	DBS	70.0	48.0	56.0	25.0	68.0	36.0	68.0	22.0
KNN	Beam Search	66.0	52.0	44.0	33.0	54.0	50.0	57.0	46.0
	Nucleus Sampling	75.0	68.0	73.0	47.0	63.0	47.0	71.0	61.0
	DBS	63.0	48.0	62.0	41.0	67.0	65.0	57.0	51.0
DPR	Beam Search	32.0	49.0	62.0	65.0	54.0	59.0	39.0	56.0
	Nucleus Sampling	29.0	57.0	57.0	43.0	46.0	38.0	56.0	52.0
	DBS	37.0	62.0	54.0	59.0	43.0	37.0	47.0	44.0

Table 10: Results on *Factual Dataset-crowd* for ALBERT-base-VitaminC-FEVER. F1 (C): F1 Consistent / F1(IC): F1 Inconsistent

Retriever	Decoding	GPT2-Small		GPT2-Medium		GPT2-Large		GPT2-XL	
		F1(C)	F1 (IC)	F1(C)	F1 (IC)	F1(C)	F1 (IC)	F1(C)	F1 (IC)
Ground Truth	Beam Search	60.0	33.0	62.0	37.0	69.0	43.0	54.0	22.0
	Nucleus Sampling	75.0	51.0	69.0	36.0	56.0	24.0	70.0	13.0
	DBS	67.0	46.0	58.0	26.0	67.0	38.0	55.0	18.0
KNN	Beam Search	67.0	55.0	43.0	29.0	48.0	48.0	47.0	39.0
	Nucleus Sampling	57.0	59.0	63.0	45.0	67.0	49.0	68.0	62.0
	DBS	51.0	43.0	54.0	34.0	58.0	59.0	51.0	49.0
DPR	Beam Search	39.0	54.0	54.0	62.0	38.0	57.0	44.0	57.0
	Nucleus Sampling	37.0	64.0	49.0	53.0	40.0	40.0	43.0	53.0
	DBS	30.0	58.0	55.0	67.0	43.0	41.0	42.0	42.0

Table 11: Results on *Factual Dataset-crowd* for ALBERT-base-VitaminC-MNLI. F1 (C): F1 Consistent / F1(IC): F1 Inconsistent

Retriever	Decoding	GPT2-Small		GPT2-Medium		GPT2-Large		GPT2-XL	
		F1(C)	F1(IC)	F1(C)	F1(IC)	F1(C)	F1(IC)	F1(C)	F1(IC)
Ground Truth	Beam Search	67.0	26.0	73.0	19.0	85.0	48.0	70.0	18.0
	Nucleus Sampling	75.0	21.0	85.0	26.0	78.0	25.0	83.0	0.0
	DBS	69.0	31.0	75.0	14.0	83.0	40.0	66.0	17.0
KNN	Beam Search	68.0	48.0	62.0	14.0	67.0	26.0	73.0	18.0
	Nucleus Sampling	62.0	21.0	79.0	33.0	74.0	27.0	72.0	30.0
	DBS	67.0	30.0	69.0	21.0	71.0	44.0	60.0	18.0
DPR	Beam Search	70.0	55.0	71.0	52.0	56.0	46.0	45.0	35.0
	Nucleus Sampling	71.0	55.0	72.0	32.0	69.0	34.0	76.0	48.0
	DBS	50.0	50.0	64.0	45.0	48.0	22.0	60.0	26.0

Table 12: Results on *Factual Dataset-crowd* for FactCC. F1 (C): F1 Consistent / F1(IC): F1 Inconsistent