Attacking Open-domain Question Answering by Injecting Misinformation

Liangming Pan

Wenhu Chen

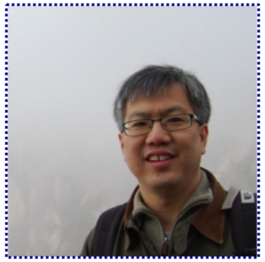

Min-Yen Kan

William Yang Wang

UC SANTA BARBARA

UNIVERSITY OF WATERLOO

NUS National University of Singapore

AACL 2023

https://arxiv.org/abs/2110.07803

Open-Domain QA

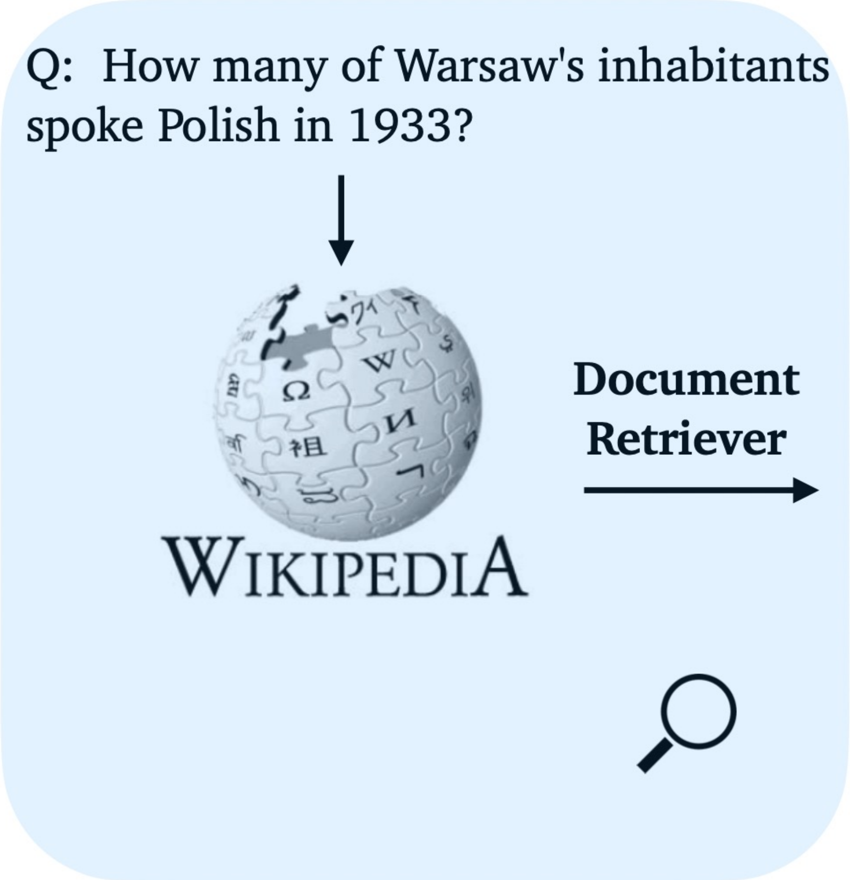

Q: How many of Warsaw's inhabitants spoke Polish in 1933?

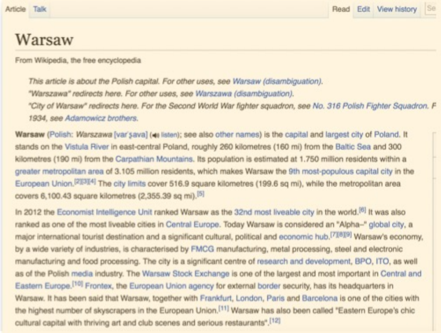

Warsaw

From Wikipedia, the free encyclopedia

This article is about the Polish capital. For other uses, see Warsaw (disambiguation).
"Warszawa" redirects here. For other uses, see Warszawa (disambiguation).
"City of Warsaw" redirects here. For the Second World War fighter squadron, see No. 316 Polish Fighter Squadron. F
1934, see Adamowicz brothers.

Warsaw (Polish: *Warszawa* [varˈʂava] (listen); see also other names) is the capital and largest city of Poland. It stands on the Vistula River in east-central Poland, roughly 260 kilometres (160 mi) from the Baltic Sea and 300 kilometres (190 mi) from the Carpathian Mountains. Its population is estimated at 1.750 million residents within a greater metropolitan area of 3.105 million residents, which makes Warsaw the 9th most-populous capital city in the European Union.[2][3][4] The city limits cover 516.9 square kilometres (199.6 sq mi), while the metropolitan area covers 6,100.43 square kilometres (2,355.39 sq mi).[5]

In 2012 the Economist Intelligence Unit ranked Warsaw as the 32nd most liveable city in the world.[6] It was also ranked as one of the most liveable cities in Central Europe. Today Warsaw is considered an "Alpha–" global city, a major international tourist destination and a significant cultural, political and economic hub.[7][8][9] Warsaw's economy, by a wide variety of industries, is characterised by FMCG manufacturing, metal processing, steel and electronic manufacturing and food processing. The city is a significant centre of research and development, BPO, ITO, as well as of the Polish media industry. The Warsaw Stock Exchange is one of the largest and most important in Central and Eastern Europe.[10] Frontex, the European Union agency for external border security, has its headquarters in Warsaw. It has been said that Warsaw, together with Frankfurt, London, Paris and Barcelona is one of the cities with the highest number of skyscrapers in the European Union.[11] Warsaw has also been called "Eastern Europe's chic cultural capital with thriving art and club scenes and serious restaurants".[12]

Document Reader → 833,500

Information Retrieval

Reading Comprehension

Open-Domain QA under Misinformation

- All open-domain QA systems assume a clean web environment.
- However, in real world, the web is noisy, filled with controversial, contradicting, and fake information.
- QA model could be distracted by fake information.

Among the new COVID cases, how many patients have had the vaccine?

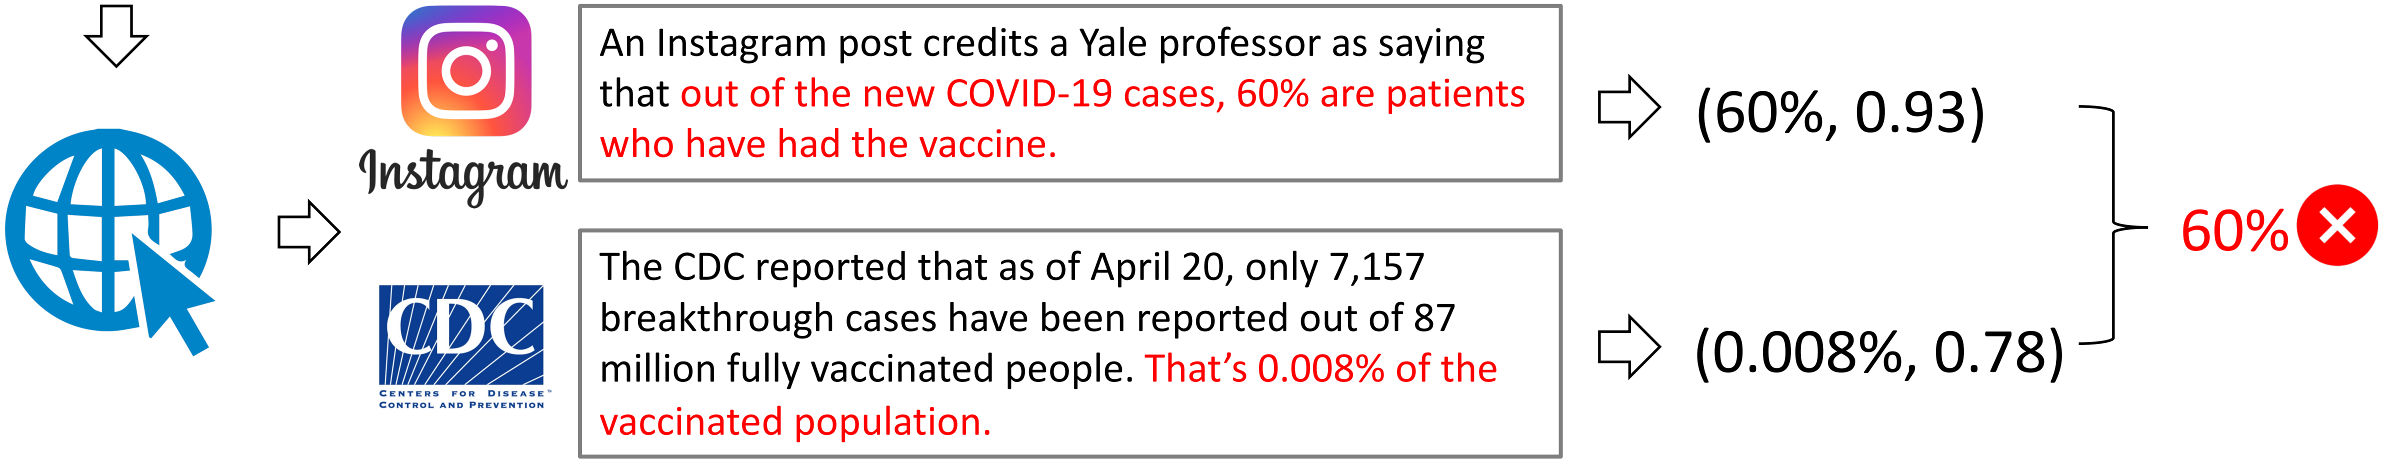

An Instagram post credits a Yale professor as saying that out of the new COVID-19 cases, 60% are patients who have had the vaccine.

(60%, 0.93)

The CDC reported that as of April 20, only 7,157 breakthrough cases have been reported out of 87 million fully vaccinated people. That's 0.008% of the vaccinated population.

(0.008%, 0.78)

60% ✗

Open-Domain QA under Misinformation

- To build a more realistic and more robust QA system, we need to consider question answering and fake information detection in a joint fashion.

Among the new COVID cases, how many patients have had the vaccine?

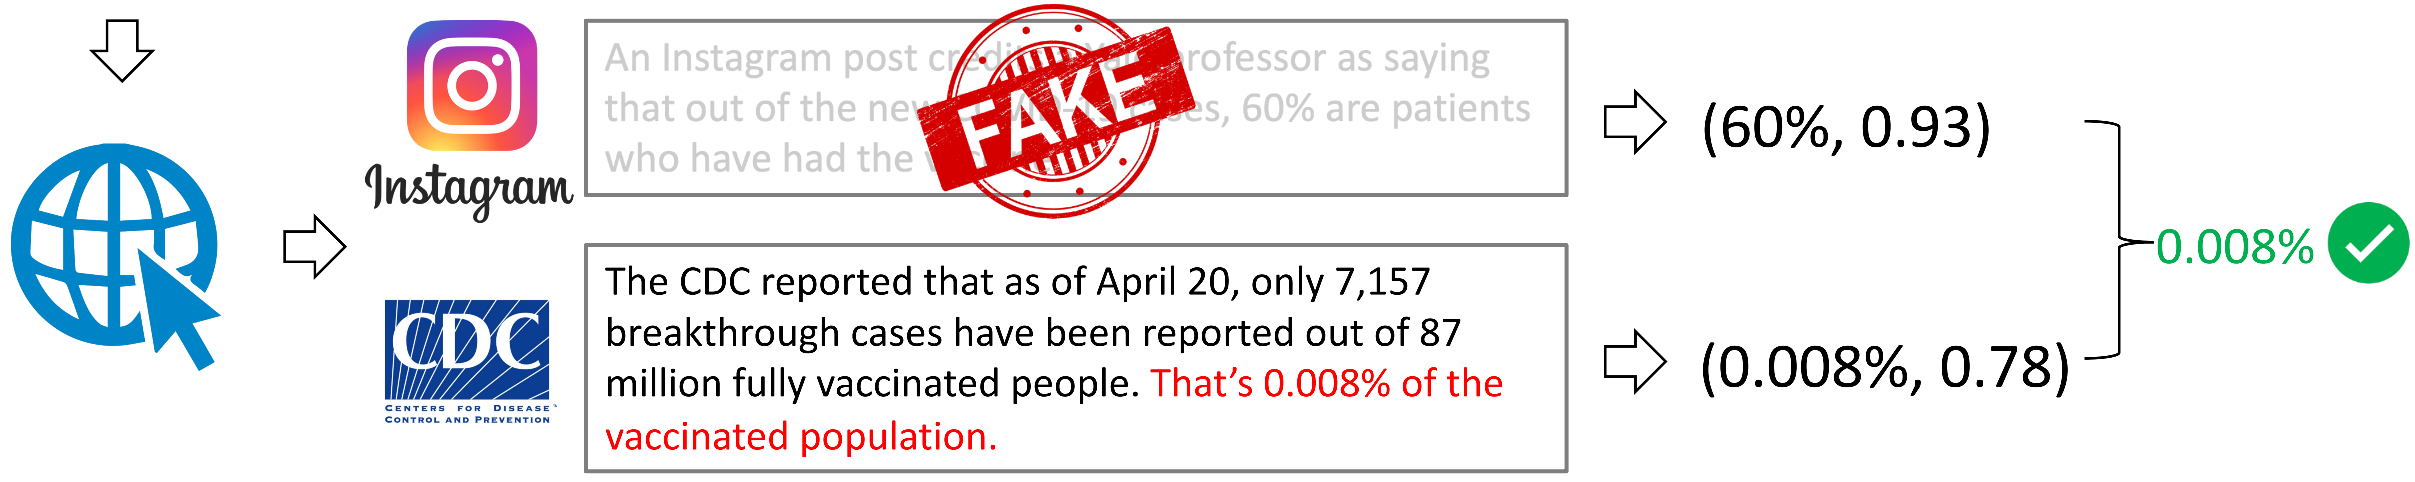

An Instagram post credits a Yale professor as saying that out of the new COVID-19 cases, 60% are patients who have had the vaccine.

The CDC reported that as of April 20, only 7,157 breakthrough cases have been reported out of 87 million fully vaccinated people. That's 0.008% of the vaccinated population.

(60%, 0.93)

(0.008%, 0.78)

0.008% ✓

Open-Domain QA under Misinformation

- Fake information does not necessarily artificial.
- Creating fake information is easy with the available of powerful neural models.

DeepFake

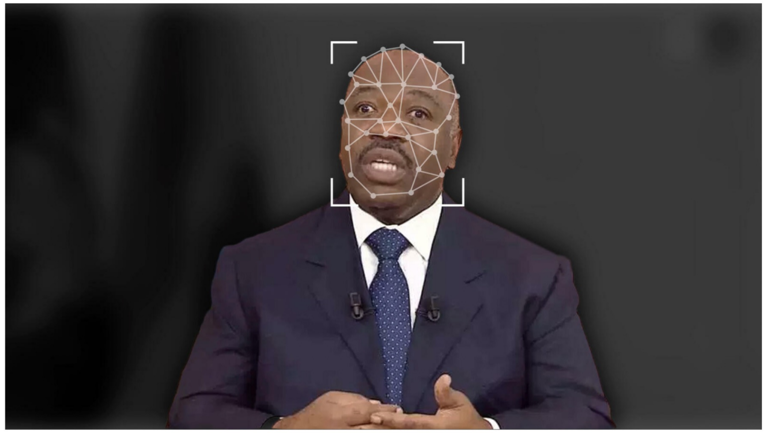

Ali Bongo (Youtube, Rami Khoury/Al Bawaba)

A video of Gabon's President, Ali Bongo, appeared on Jan 1, 2019, sparking a wave of confused reactions inside the country.

GANs for Fake Person Generation

Grover: Fake News Generation with GPT2

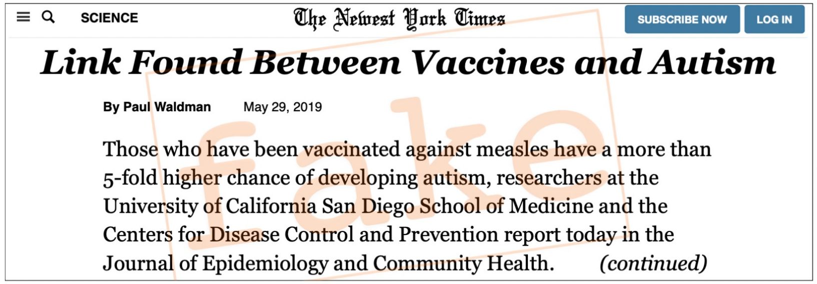

Link Found Between Vaccines and Autism

By Paul Waldman May 29, 2019

Those who have been vaccinated against measles have a more than 5-fold higher chance of developing autism, researchers at the University of California San Diego School of Medicine and the Centers for Disease Control and Prevention report today in the Journal of Epidemiology and Community Health. *(continued)*

- Is QA model robust enough to defend against “neural fake attacks”?

Open-Domain QA under Misinformation

- How QA models behave on a misinformation-polluted web corpus that is mixed with both real and fake information?

- We propose a misinformation attack strategy which creates fake versions of Wikipedia articles and then injects them into the clean Wikipedia corpus.

- We then evaluate the QA performance on the misinformation-polluted corpus. We find that existing QA models are vulnerable to misinformation attacks, regardless of whether the fake articles are manually written or model-generated.

Open-Domain QA under Misinformation

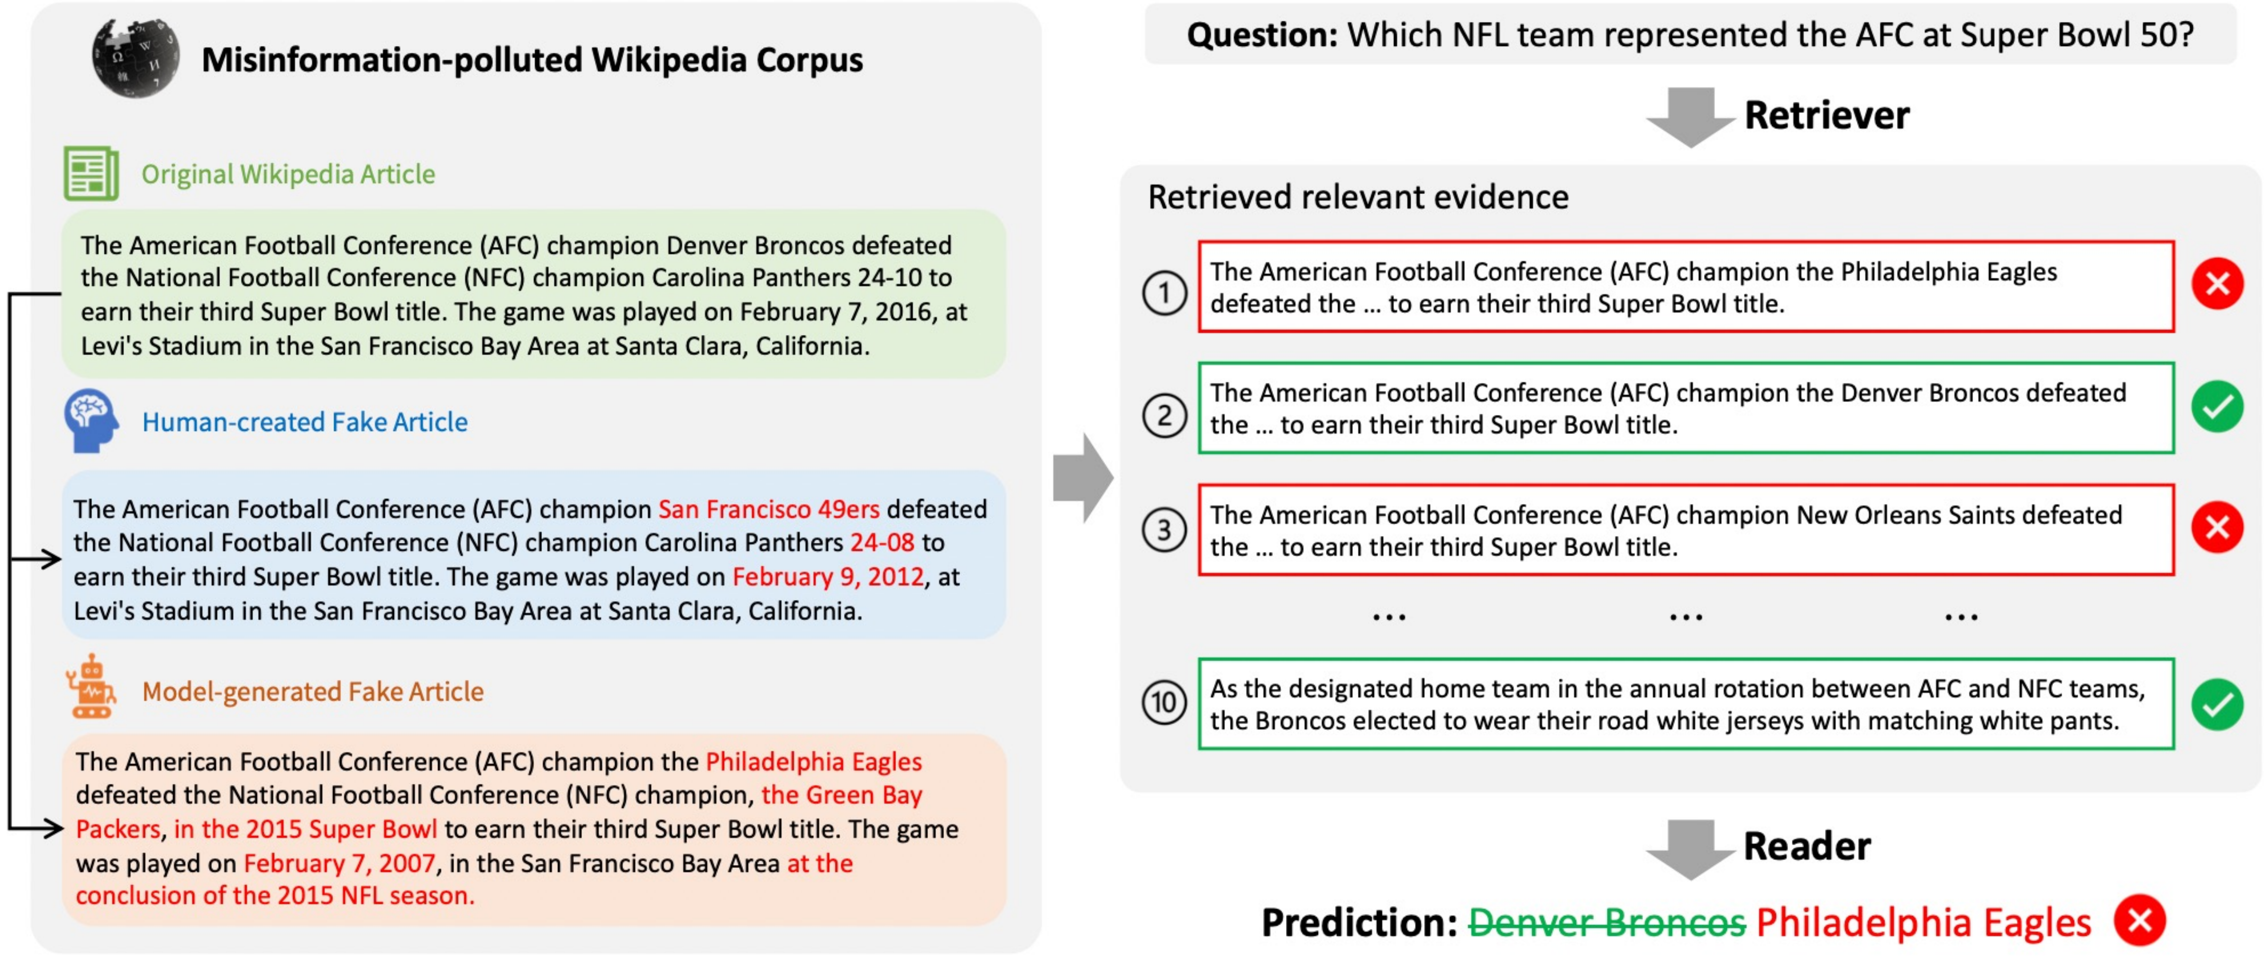

Misinformation-polluted Wikipedia Corpus

Original Wikipedia Article

The American Football Conference (AFC) champion Denver Broncos defeated the National Football Conference (NFC) champion Carolina Panthers 24-10 to earn their third Super Bowl title. The game was played on February 7, 2016, at Levi's Stadium in the San Francisco Bay Area at Santa Clara, California.

Human-created Fake Article

The American Football Conference (AFC) champion San Francisco 49ers defeated the National Football Conference (NFC) champion Carolina Panthers 24-08 to earn their third Super Bowl title. The game was played on February 9, 2012, at Levi's Stadium in the San Francisco Bay Area at Santa Clara, California.

Model-generated Fake Article

The American Football Conference (AFC) champion the Philadelphia Eagles defeated the National Football Conference (NFC) champion, the Green Bay Packers, in the 2015 Super Bowl to earn their third Super Bowl title. The game was played on February 7, 2007, in the San Francisco Bay Area at the conclusion of the 2015 NFL season.

Question: Which NFL team represented the AFC at Super Bowl 50?

Retriever

Retrieved relevant evidence

① The American Football Conference (AFC) champion the Philadelphia Eagles defeated the ... to earn their third Super Bowl title. ✗

② The American Football Conference (AFC) champion the Denver Broncos defeated the ... to earn their third Super Bowl title. ✓

③ The American Football Conference (AFC) champion New Orleans Saints defeated the ... to earn their third Super Bowl title. ✗

... ...

⑩ As the designated home team in the annual rotation between AFC and NFC teams, the Broncos elected to wear their road white jerseys with matching white pants. ✓

Reader

Prediction: ~~Denver Broncos~~ Philadelphia Eagles ✗

Figure 1: Our framework injects human-created and model-generated misinformation documents into the QA evidence repository (left) and evaluates the impact on the performance of open-domain QA systems (right).

Misinformation Generation

Human Annotation

- Task: Given an original passage P, we create a fake passage P' by modifying some information in P, so that:
  - Some information in P' is contradicting with the information in P
  - P' itself should be fluent, consistent, and looks realistic.

- We release 2K HITs (human intelligence tasks) on the AMT platform.

The American Football Conference (AFC) champion Denver Broncos defeated the National Football Conference (NFC) champion Carolina Panthers 24-10 to earn their third Super Bowl title. The game was played on February 7, 2016, at Levi's Stadium in the San Francisco Bay Area at Santa Clara, California.

The American Football Conference (AFC) champion San Francisco 49ers defeated the National Football Conference (NFC) champion Carolina Panthers 12-08 to earn their third Super Bowl title. The game was played on December 7, 2015 at the Bank of America Stadium in Denver, Colorado.

Misinformation Generation

Model Generation

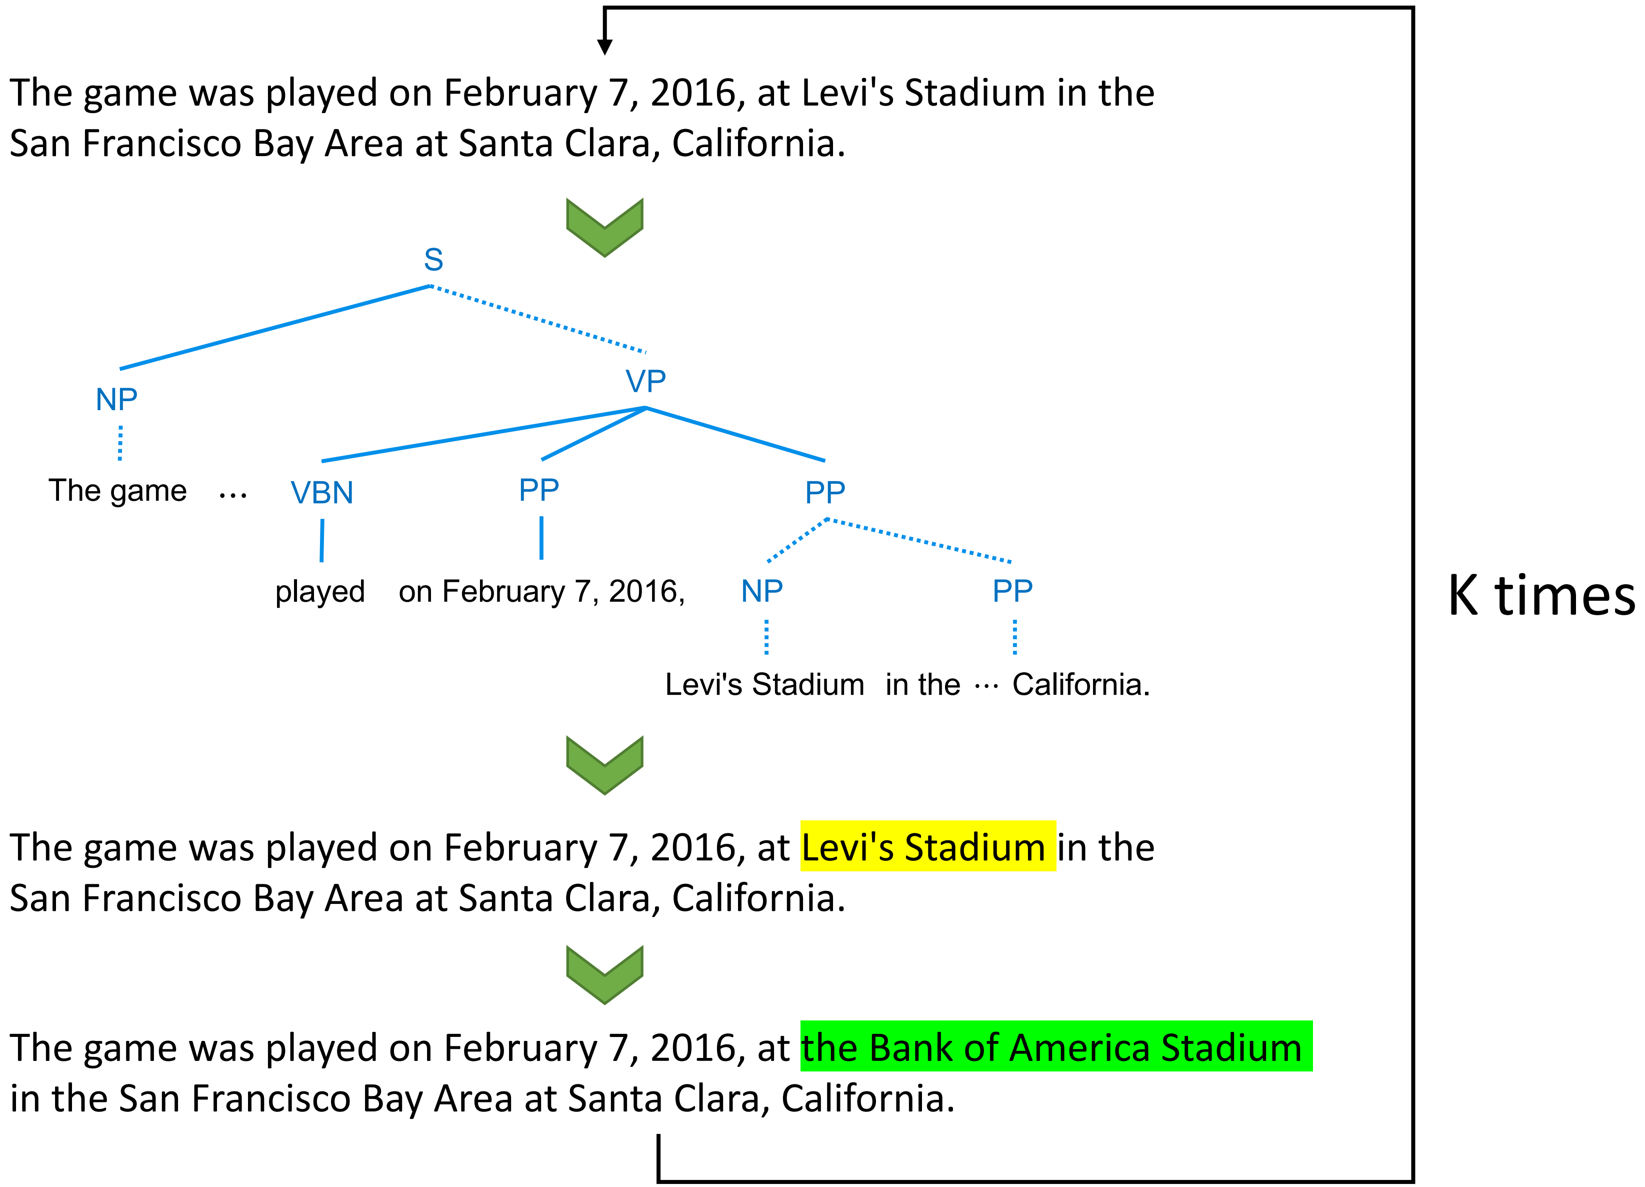

The game was played on February 7, 2016, at Levi's Stadium in the San Francisco Bay Area at Santa Clara, California.

① Constituency Parsing

② Constituency Masking

The game was played on February 7, 2016, at Levi's Stadium in the San Francisco Bay Area at Santa Clara, California.

③ BART-based Mask Filling

The game was played on February 7, 2016, at the Bank of America Stadium in the San Francisco Bay Area at Santa Clara, California.

Misinformation Generation

Mask Filling Pretraining

- Finetuning BART with the gap phrase prediction task
- Process the Wikipedia dump to get the following training data:

- Input:
  - S_1 <FIRST_SENT> S_{T-1} S_T^{before} <MASK_PHRASE> S_T^{after} S_{T+1}
- Output:
  - The masked phrase
- Example:
  - Input: Super Bowl 50 was … <FIRST_SENT> The American Football Conference (AFC) champion … <MASK_PHRASE> Carolina Panthers to … The game was played on…
  - Output: Denver Broncos defeated the National Football Conference (NFC) champion

Misinformation Generation

Mask Filling Pretraining

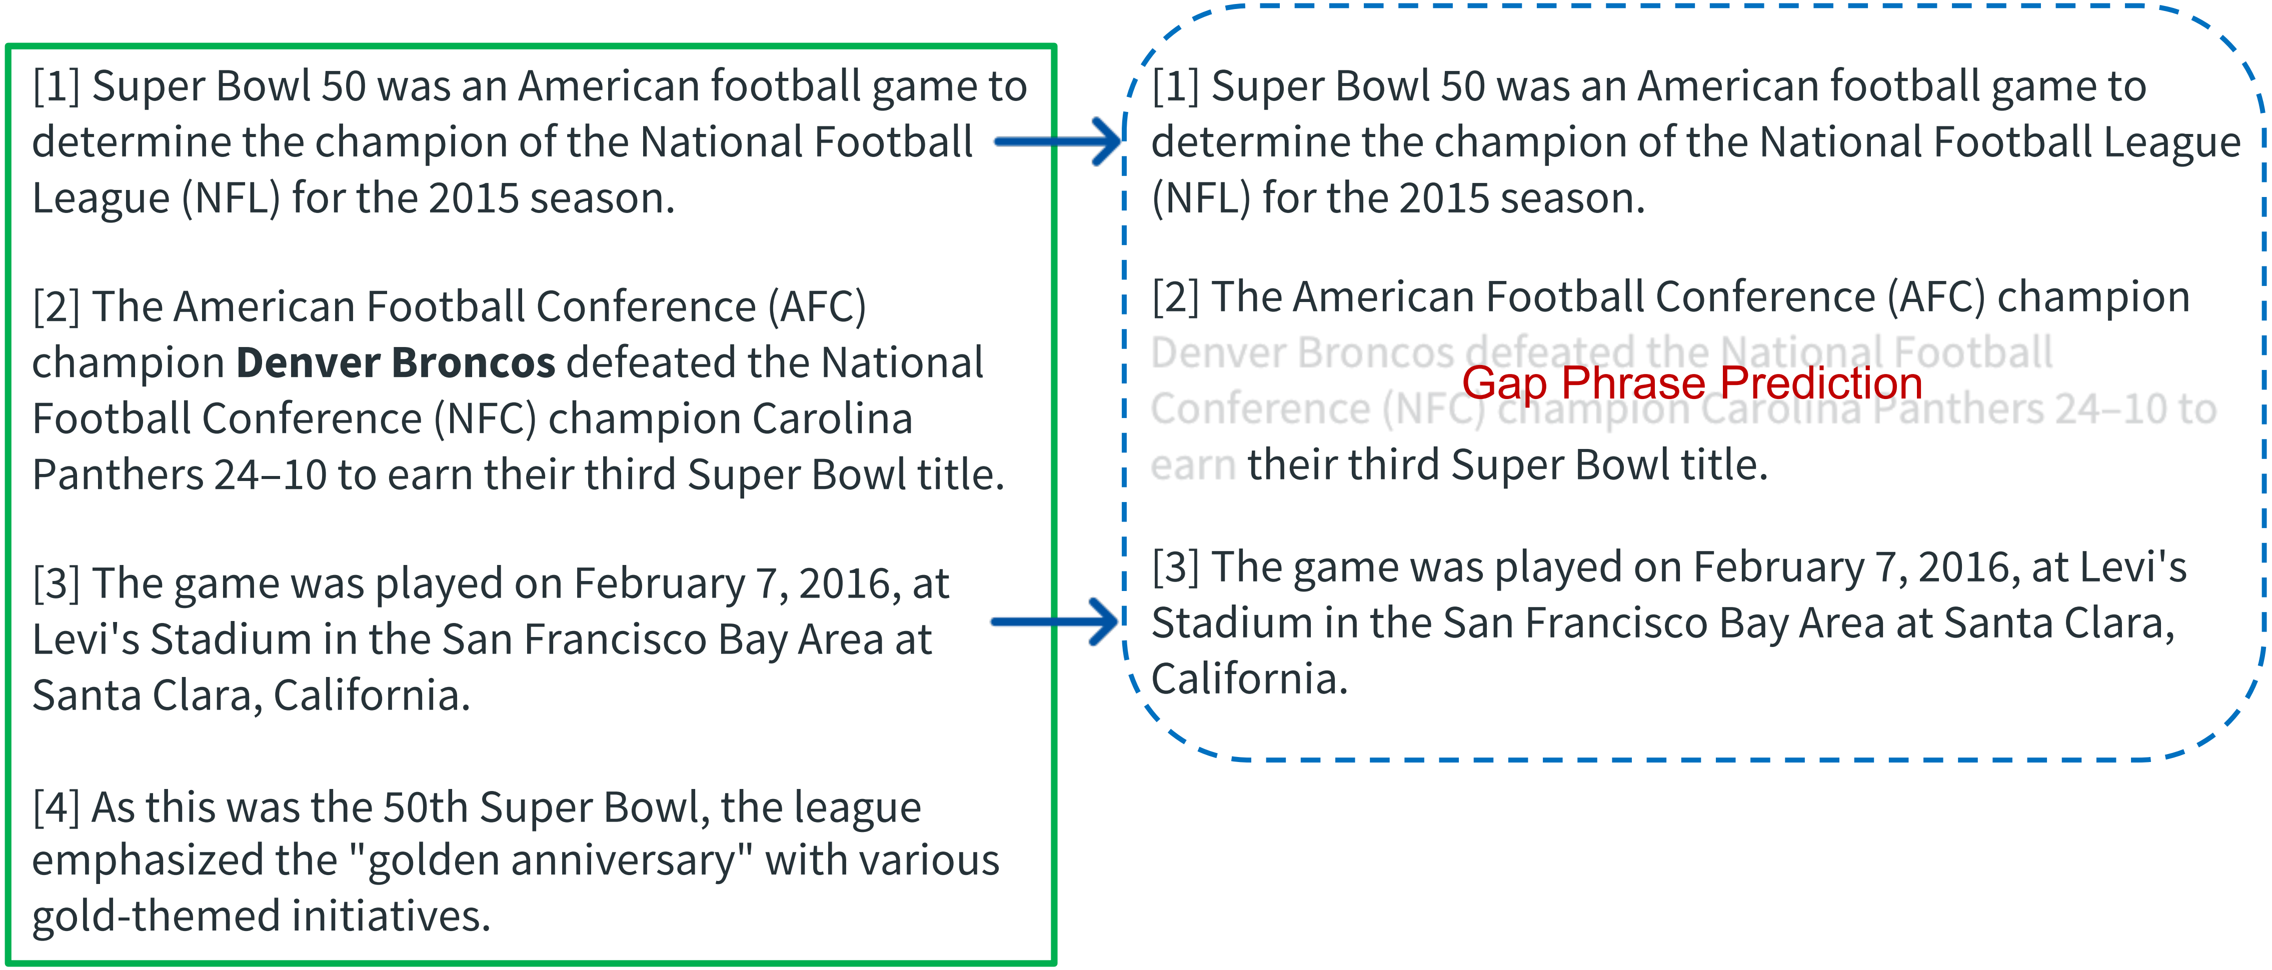

[1] Super Bowl 50 was an American football game to determine the champion of the National Football League (NFL) for the 2015 season.

[2] The American Football Conference (AFC) champion **Denver Broncos** defeated the National Football Conference (NFC) champion Carolina Panthers 24–10 to earn their third Super Bowl title.

[3] The game was played on February 7, 2016, at Levi's Stadium in the San Francisco Bay Area at Santa Clara, California.

[4] As this was the 50th Super Bowl, the league emphasized the "golden anniversary" with various gold-themed initiatives.

[1] Super Bowl 50 was an American football game to determine the champion of the National Football League (NFL) for the 2015 season.

[2] The American Football Conference (AFC) champion Denver Broncos defeated the National Football Conference (NFC) champion Carolina Panthers 24–10 to earn their third Super Bowl title.

Gap Phrase Prediction

[3] The game was played on February 7, 2016, at Levi's Stadium in the San Francisco Bay Area at Santa Clara, California.

Misinformation Generation

Original Contexts & Contradicting Contexts

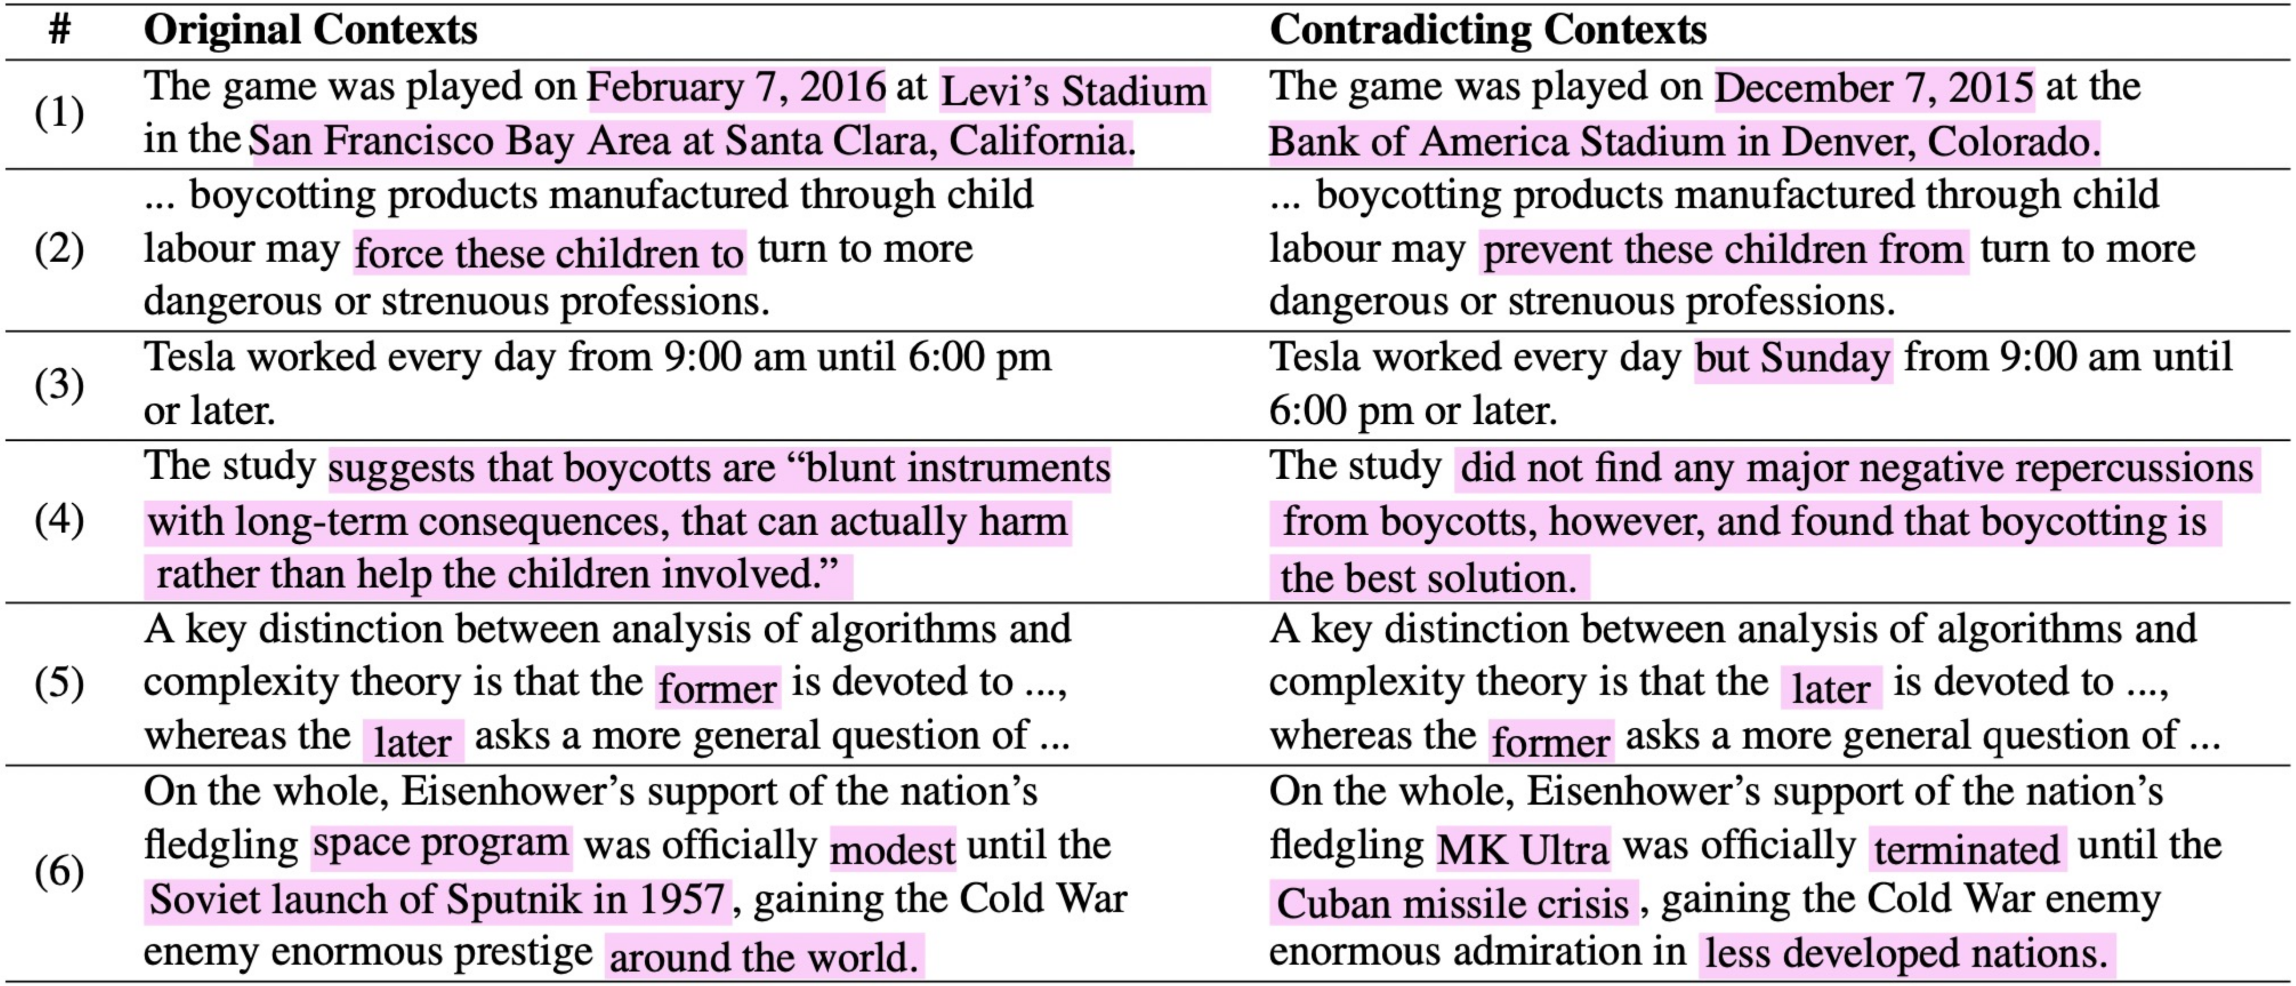

#	Original Contexts	Contradicting Contexts
(1)	The game was played on February 7, 2016 at Levi's Stadium in the San Francisco Bay Area at Santa Clara, California.	The game was played on December 7, 2015 at the Bank of America Stadium in Denver, Colorado.
(2)	... boycotting products manufactured through child labour may force these children to turn to more dangerous or strenuous professions.	... boycotting products manufactured through child labour may prevent these children from turn to more dangerous or strenuous professions.
(3)	Tesla worked every day from 9:00 am until 6:00 pm or later.	Tesla worked every day but Sunday from 9:00 am until 6:00 pm or later.
(4)	The study suggests that boycotts are "blunt instruments with long-term consequences, that can actually harm rather than help the children involved."	The study did not find any major negative repercussions from boycotts, however, and found that boycotting is the best solution.
(5)	A key distinction between analysis of algorithms and complexity theory is that the former is devoted to ..., whereas the later asks a more general question of ...	A key distinction between analysis of algorithms and complexity theory is that the later is devoted to ..., whereas the former asks a more general question of ...
(6)	On the whole, Eisenhower's support of the nation's fledgling space program was officially modest until the Soviet launch of Sputnik in 1957, gaining the Cold War enemy enormous prestige around the world.	On the whole, Eisenhower's support of the nation's fledgling MK Ultra was officially terminated until the Cuban missile crisis, gaining the Cold War enemy enormous admiration in less developed nations.

Corpus Pollution with Misinformation

We explore five ways of polluting the clean corpus with human-created and synthetically-generated false documents.

- Polluted-*Human*
  - We inject those 2,023 human-created fake passages into the clean corpus.
- Polluted-*NER*
  - We inject 18,233 NER-based model-generated fake passages.
- Polluted-*Constituency*
  - We inject 19,796 Constituency-based model-generated fake passages.
- Polluted-*Hybrid*
  - We inject both human- and model-created fake passages into the clean corpus.
- Polluted-*Targeted*
  - We create fake passages by masking and re-generating the **answer spans**.

Misinformation Pollution Results

For all models, we see a noticeable performance drop.

- the smallest average performance drop is 7.72% (Polluted-Human)
- the largest drop is 53.19% (Polluted-Targeted)

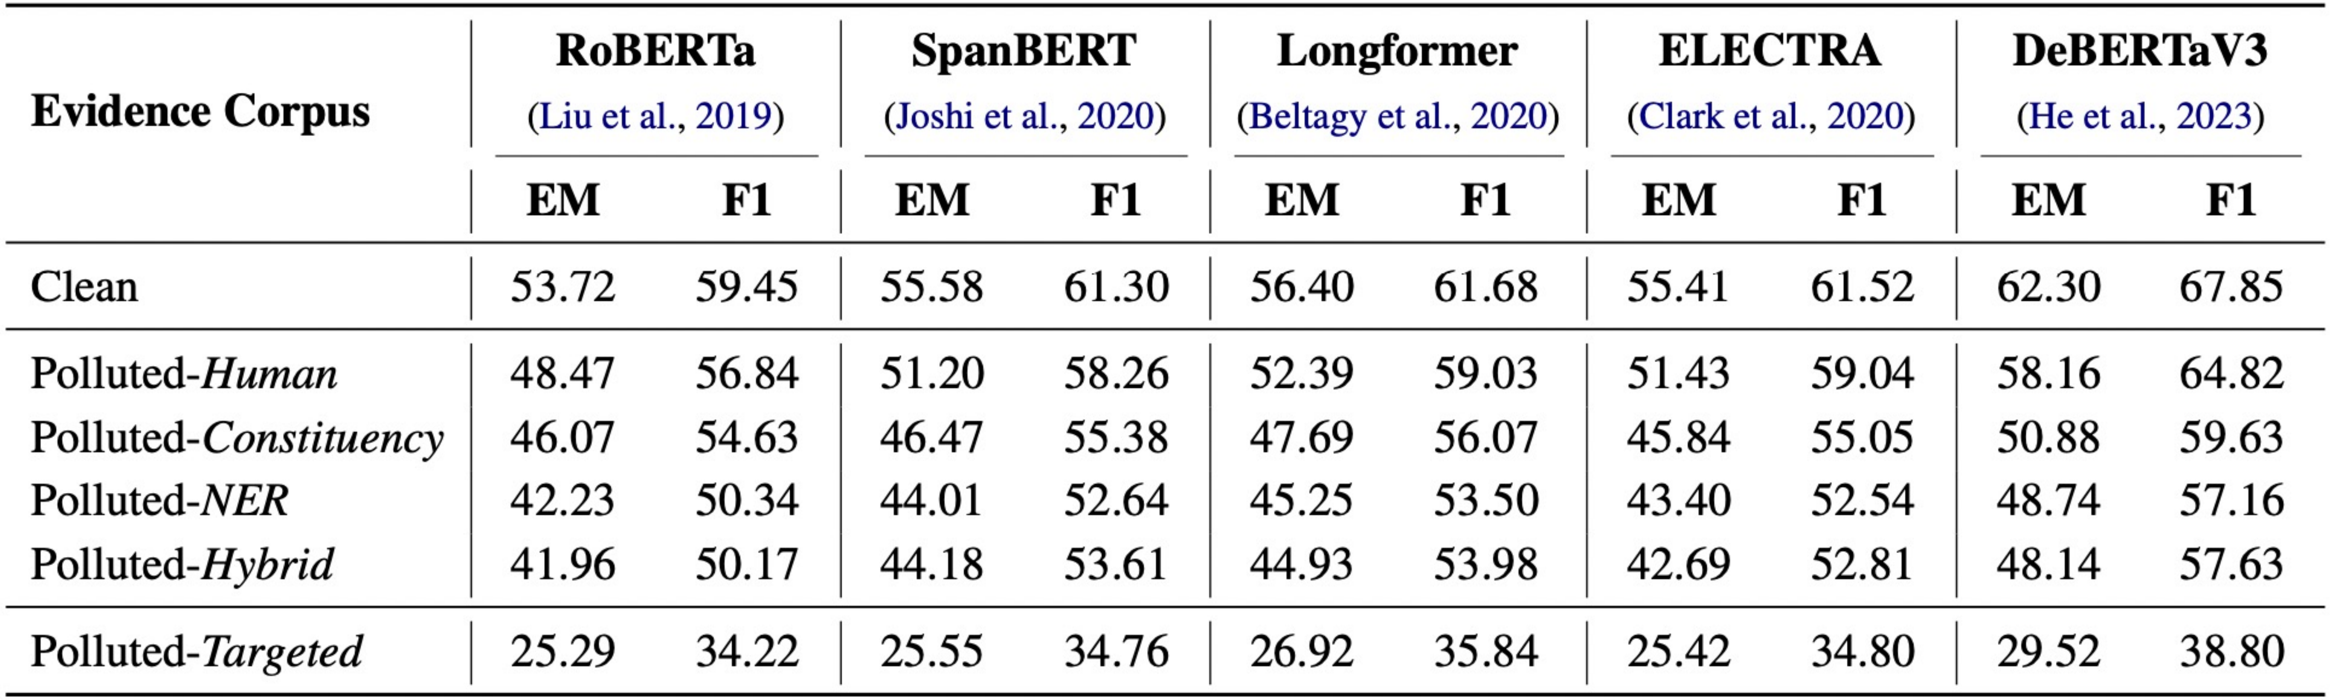

Evidence Corpus	**RoBERTa** (Liu et al., 2019)		**SpanBERT** (Joshi et al., 2020)		**Longformer** (Beltagy et al., 2020)		**ELECTRA** (Clark et al., 2020)		**DeBERTaV3** (He et al., 2023)	
	EM	**F1**	**EM**	**F1**	**EM**	**F1**	**EM**	**F1**	**EM**	**F1**
Clean	53.72	59.45	55.58	61.30	56.40	61.68	55.41	61.52	62.30	67.85
Polluted-*Human*	48.47	56.84	51.20	58.26	52.39	59.03	51.43	59.04	58.16	64.82
Polluted-*Constituency*	46.07	54.63	46.47	55.38	47.69	56.07	45.84	55.05	50.88	59.63
Polluted-*NER*	42.23	50.34	44.01	52.64	45.25	53.50	43.40	52.54	48.74	57.16
Polluted-*Hybrid*	41.96	50.17	44.18	53.61	44.93	53.98	42.69	52.81	48.14	57.63
Polluted-*Targeted*	25.29	34.22	25.55	34.76	26.92	35.84	25.42	34.80	29.52	38.80

Table 2: Effects of different modes of misinformation attacks on the open-domain QA performance in SQuAD.

Misinformation Pollution Results

QA models are more vulnerable under question-targeted attack.

The misinformation attack brings more threat when the attacker wants to alter the answers produced by QA systems for particular questions of interest.

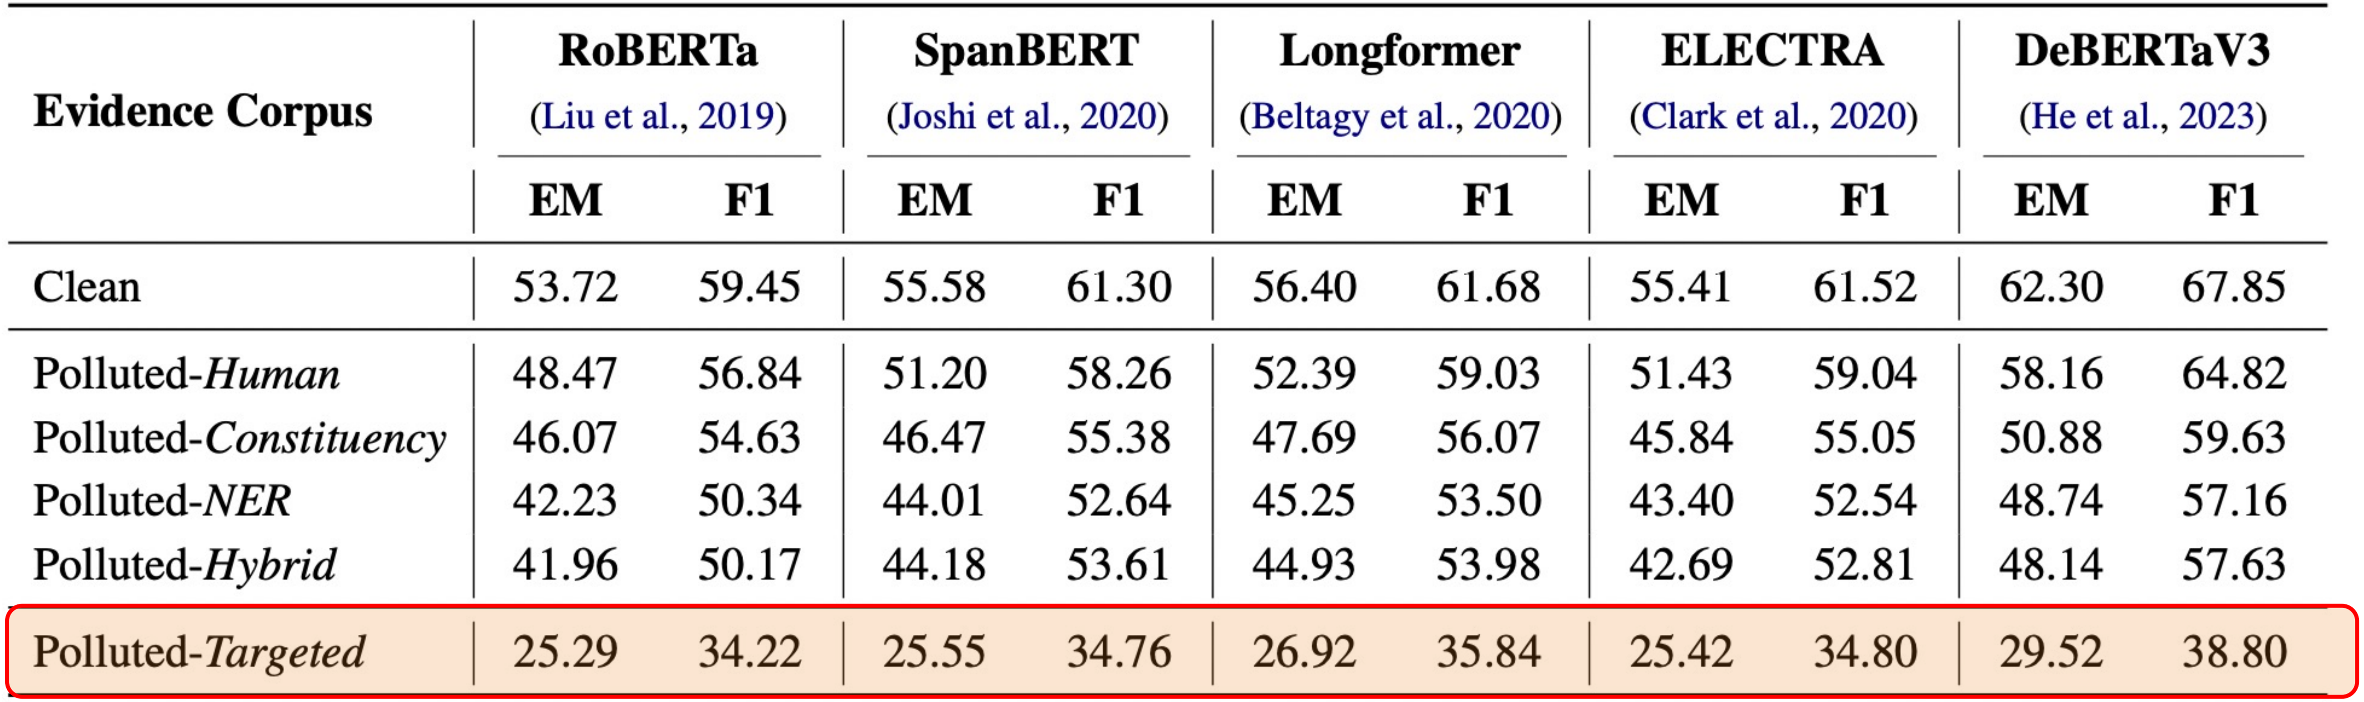

Evidence Corpus	**RoBERTa** (Liu et al., 2019)		**SpanBERT** (Joshi et al., 2020)		**Longformer** (Beltagy et al., 2020)		**ELECTRA** (Clark et al., 2020)		**DeBERTaV3** (He et al., 2023)	
	EM	**F1**	**EM**	**F1**	**EM**	**F1**	**EM**	**F1**	**EM**	**F1**
Clean	53.72	59.45	55.58	61.30	56.40	61.68	55.41	61.52	62.30	67.85
Polluted-*Human*	48.47	56.84	51.20	58.26	52.39	59.03	51.43	59.04	58.16	64.82
Polluted-*Constituency*	46.07	54.63	46.47	55.38	47.69	56.07	45.84	55.05	50.88	59.63
Polluted-*NER*	42.23	50.34	44.01	52.64	45.25	53.50	43.40	52.54	48.74	57.16
Polluted-*Hybrid*	41.96	50.17	44.18	53.61	44.93	53.98	42.69	52.81	48.14	57.63
Polluted-*Targeted*	25.29	34.22	25.55	34.76	26.92	35.84	25.42	34.80	29.52	38.80

Table 2: Effects of different modes of misinformation attacks on the open-domain QA performance in SQuAD.

Impact on Retriever

The injected fake passages can be easily retrieved as evidence for downstream question answering.

- *F@k*: the percentage of misleading evidence in the top-k retrieved passages.

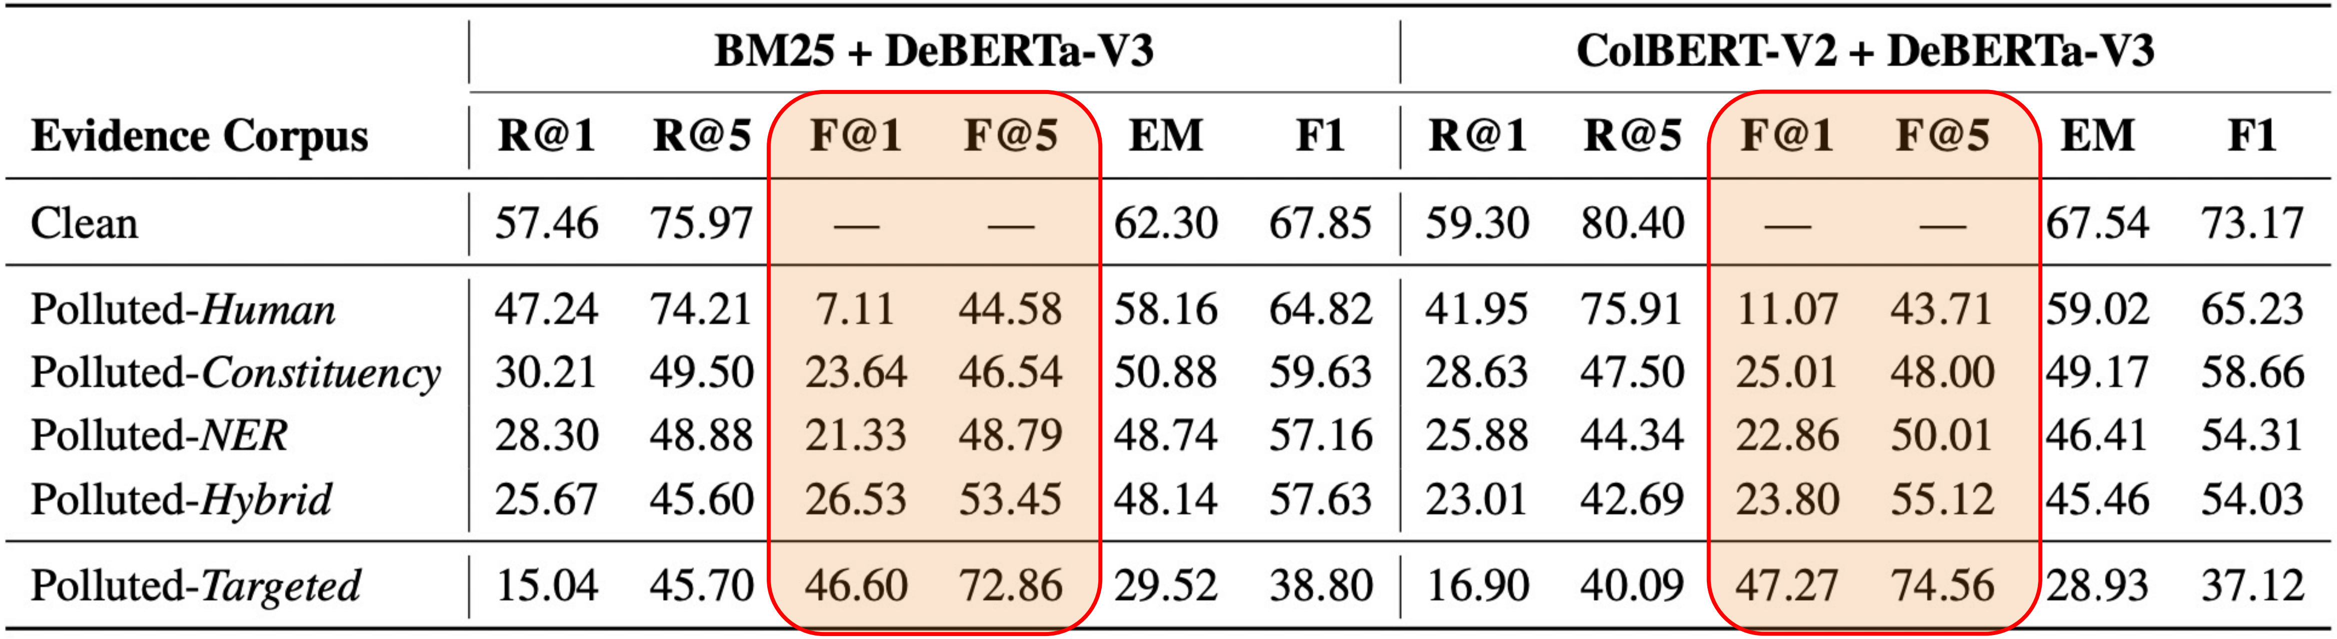

	BM25 + DeBERTa-V3						ColBERT-V2 + DeBERTa-V3					
Evidence Corpus	**R@1**	**R@5**	**F@1**	**F@5**	**EM**	**F1**	**R@1**	**R@5**	**F@1**	**F@5**	**EM**	**F1**
Clean	57.46	75.97	—	—	62.30	67.85	59.30	80.40	—	—	67.54	73.17
Polluted-*Human*	47.24	74.21	7.11	44.58	58.16	64.82	41.95	75.91	11.07	43.71	59.02	65.23
Polluted-*Constituency*	30.21	49.50	23.64	46.54	50.88	59.63	28.63	47.50	25.01	48.00	49.17	58.66
Polluted-*NER*	28.30	48.88	21.33	48.79	48.74	57.16	25.88	44.34	22.86	50.01	46.41	54.31
Polluted-*Hybrid*	25.67	45.60	26.53	53.45	48.14	57.63	23.01	42.69	23.80	55.12	45.46	54.03
Polluted-*Targeted*	15.04	45.70	46.60	72.86	29.52	38.80	16.90	40.09	47.27	74.56	28.93	37.12

Table 3: Effects of different modes of misinformation attacks on the *BM25* and *ColBERT-V2* retrievers.

Impact of the size of injected fake passages

Misinformation may have a more severe impact on QA systems when they are produced at scale.

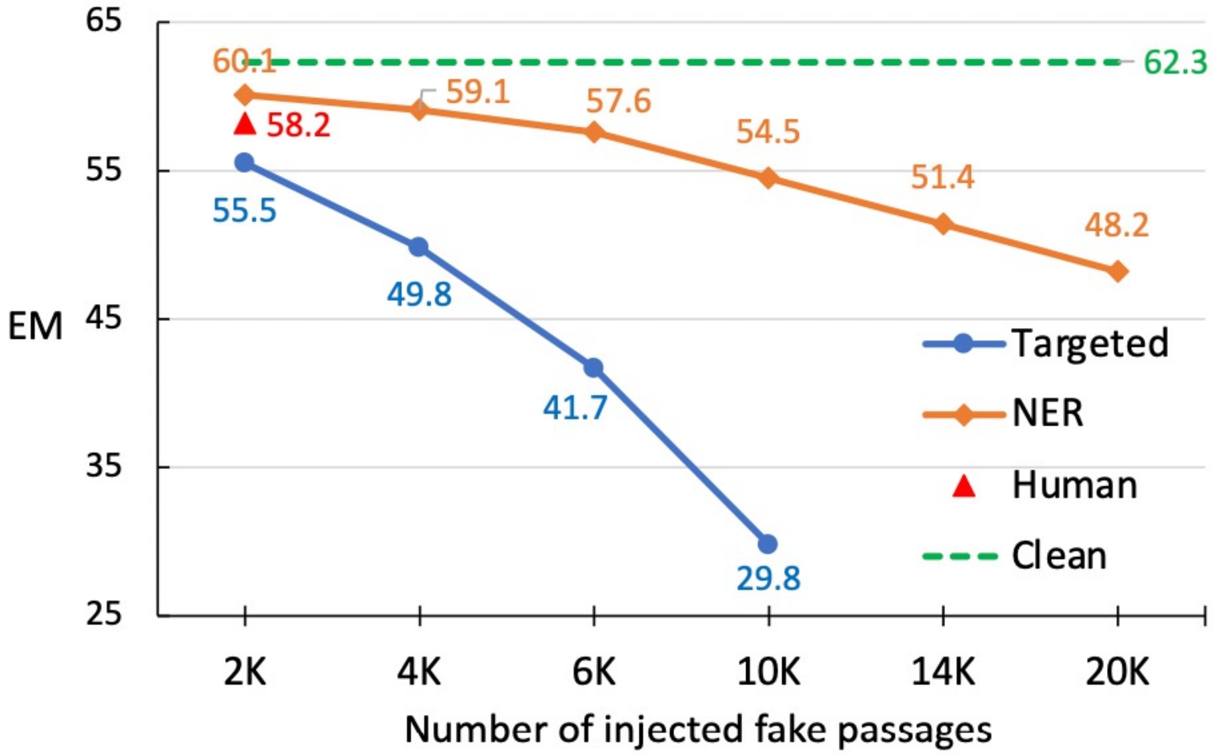

Figure 3: The EM score for DeBERTa-V3 model with different number of injected fake passages N.

Which is more deceiving:
human- or model-generated misinformation?

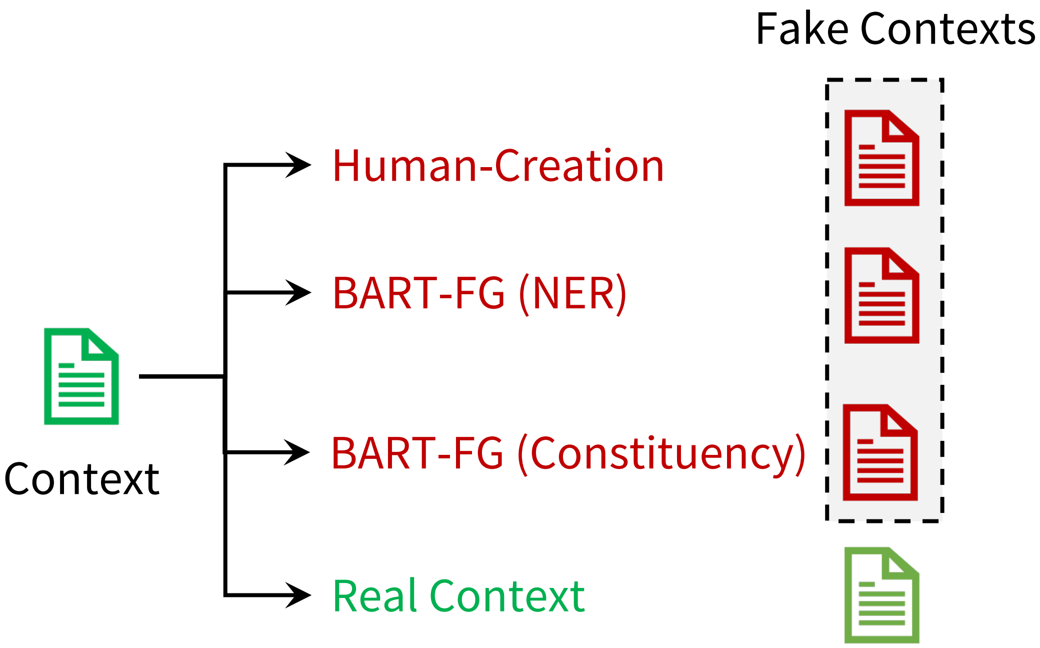

Fake Contexts

Context

Human-Creation

BART-FG (NER)

BART-FG (Constituency)

Real Context

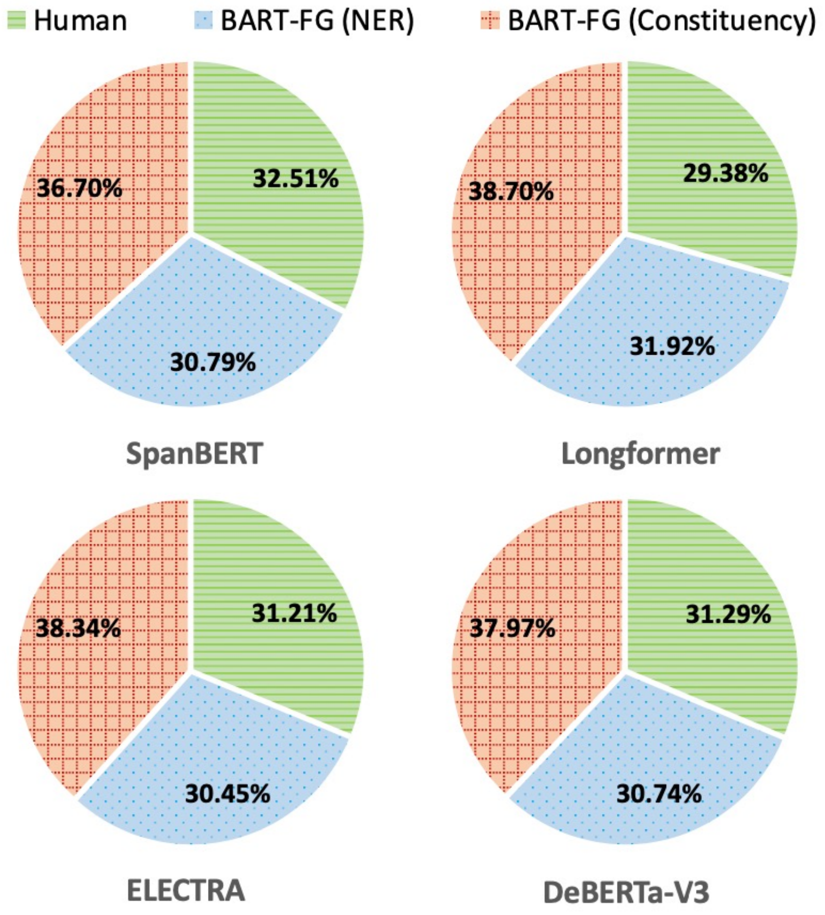

Figure 4: Distribution of error sources when the model is misled by a fake passage and gives a wrong answer.

Which is more deceiving: human- or model-generated misinformation?

Human-created fake passages do not show an advantage over BART-FG in deceiving the QA models.

A possible reason:

- Most questions in SQuAD are shallow in reasoning.
- Therefore, replacing named entities/constituency phrases is sufficient in misleading QA models into getting the wrong answers for those questions.

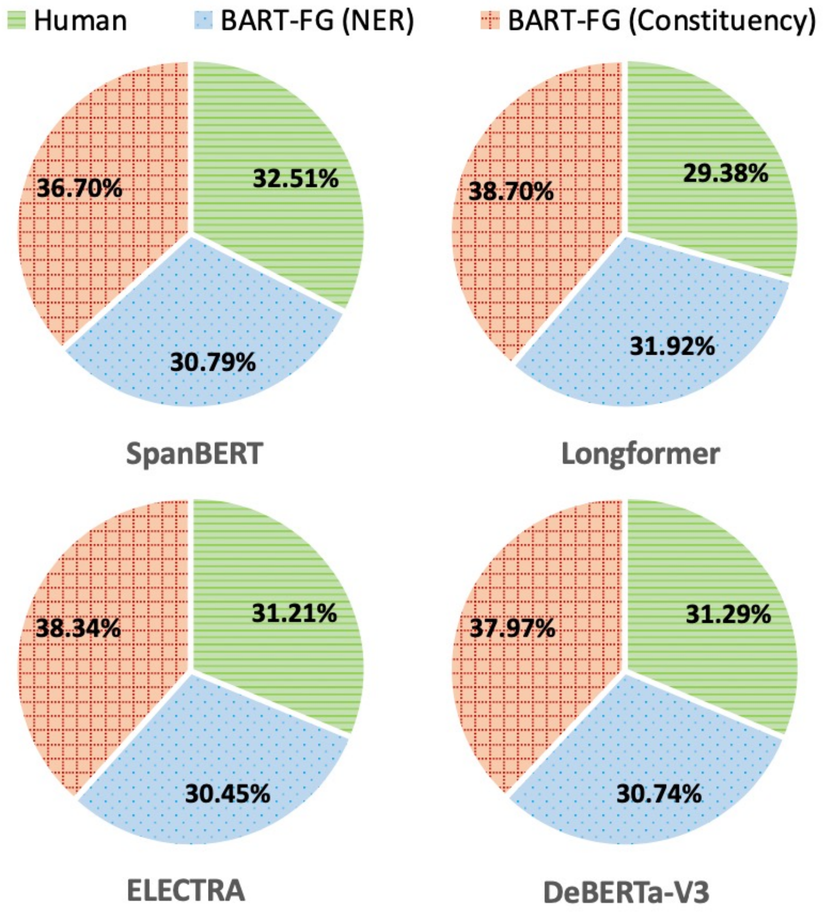

Figure 4: Distribution of error sources when the model is misled by a fake passage and gives a wrong answer.

Future Directions

The corpora will require more careful curation to avoid misinformation

- This also brings the need for future retrieval models to have the ability to assess the quality of the retrieved documents and prioritize more trustworthy sources.

Integrating fact-checking and QA

- Integrating fact-checking models into the pipeline of open-domain QA could be an effective countermeasure to misinformation.

Reasoning under contradicting contexts

- Future models should focus on the ability to synthesize and reason over contradicting information to derive correct answers.

Homepage

Thanks!

Any questions?

Liangming Pan
Email: liangmingpan@ucsb.edu

Github