

Knowledge-Enhanced Evidence Retrieval for Counterargument Generation

Yohan Jo¹, Haneul Yoo², JinYeong Bak³, Alice Oh², Chris Reed⁴, Eduard Hovy¹

¹Carnegie Mellon University, USA, ²KAIST, South Korea

³Sungkyunkwan University, South Korea, ⁴University of Dundee, UK

yohanj@cs.cmu.edu, haneul.yoo@kaist.ac.kr, jy.bak@skku.edu,
alice.oh@kaist.edu, c.a.reed@dundee.ac.uk, hovy@cmu.edu

Abstract

Finding counterevidence to statements is key to many tasks, including counterargument generation. We build a system that, given a statement, retrieves counterevidence from diverse sources on the Web. At the core of this system is a natural language inference (NLI) model that determines whether a candidate sentence is valid counterevidence or not. Most NLI models to date, however, lack proper reasoning abilities necessary to find counterevidence that involves complex inference. Thus, we present a knowledge-enhanced NLI model that aims to handle causality- and example-based inference by incorporating knowledge graphs. Our NLI model outperforms baselines for NLI tasks, especially for instances that require the targeted inference. In addition, this NLI model further improves the counterevidence retrieval system, notably finding complex counterevidence better.¹

1 Introduction

Generating counterarguments is key to many applications, such as debating systems (Slonim, 2018), essay feedback generation (Woods et al., 2017), and legal decision making (Feteris et al., 2017). In NLP, many prior studies have focused on generating counterarguments to the main conclusions of long arguments, usually motions. Although such counterarguments are useful, argumentative dialogue is usually interactive and synchronous, and one often needs to address specific statements in developing argument. For instance, in the Change-MyView (CMV) subreddit, challengers often quote and counter specific statements in the refuted argument, where 41% of these attacks are about factual falsehood, such as exceptions, feasibility, and lack of evidence (Jo et al., 2020). Hence, the scope of our work is narrower than most prior work. Instead

¹Source code and data are available at <https://github.com/yohanjo/kenli>.

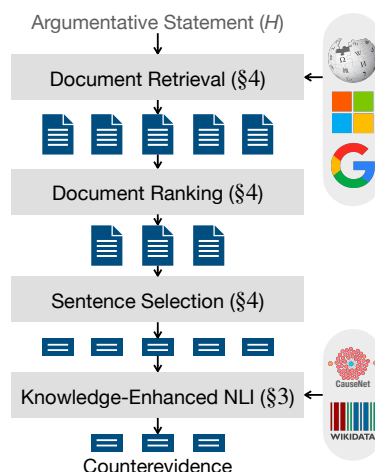


Figure 1: Architecture overview.

of generating a counterargument to a *complete argument*, we aim to find counterevidence to *specific statements* in an argument. This counterevidence may serve as essential building blocks for developing a larger counterargument and also allow for more interactive development of argumentation.

We adopt a popular fact-verification framework (Thorne et al., 2018): given a statement to refute, we retrieve relevant documents from the Web and select counterevidence (Figure 1). At the core of this framework is a module that determines whether a candidate sentence is valid counterevidence to the given statement. A natural choice for this module is a natural language inference (NLI) model. But NLI models to date have shown a lack of reasoning abilities (Williams et al., 2020), which is problematic because counterarguments often involve complex inference. To overcome this limitation, we enhance NLI by focusing on two types of inference informed by argumentation theory (Walton et al., 2008). The first is *argument from examples*, as in:

Claim: *Vegan food reduces the risk of diseases.*

Counterevidence: *Legume protein sources can result in phytohemagglutinin poisoning.*

The inference is based on the fact that “legume

protein sources” and “phytohemagglutinin poisoning” are **examples** of “vegan food” and “diseases”, respectively. The second type of inference is *argument from cause-to-effect*, as in:

Claim: *Veganism reduces the risk of diabetes.*

Counterevidence: *Vegan diets suffer from poor nutrition.*

The inference is based on the fact that poor nutrition can **cause** diabetes.

In order to handle causality- and example-based inference, we develop a knowledge-enhanced NLI model (§3). By incorporating two knowledge graphs—CauseNet and Wikidata—into the model while training on public NLI datasets, the accuracy of the NLI model improves across NLI datasets, especially for challenging instances that require the targeted inference.

We integrate this NLI model into the entire retrieval system to find counterevidence to argumentative statements from two online argument platforms, ChangeMyView (CMV) and Kialo (§4). We demonstrate that our knowledge-enhanced NLI model improves the system, finding more complex counterevidence. We also conduct in-depth analyses of the utility of different types of source documents and document search methods (Wikipedia, Bing, and Google).

Our contributions are as follows:

- A knowledge-enhanced NLI model to handle causality- and example-based inference.
- A counterevidence retrieval system improved by the NLI model, along with analyses of different document types and search methods.
- A new challenging dataset of counterevidence retrieval, along with all search results and retrieved documents from Bing and Google.

2 Related Work

2.1 Counterargument Generation

In NLP, there are two main approaches to counterargument generation. **Retrieval-based** approaches retrieve existing arguments from debates that best serve as counterarguments, based on how similar a claim is to the target argument (Le et al., 2018) and how dissimilar a premise is (Wachsmuth et al., 2018). Some studies retrieve texts that contain negative consequences of the target argument (Reisert et al., 2015; Sato et al., 2015). Recently, a human-curated corpus was developed (Orbach et al., 2020).

Neural language generation approaches take

the target argument as input and generate a counterargument using a neural network (Hua and Wang, 2018; Hua et al., 2019). These approaches still retrieve evidence sentences from Wikipedia or news articles that are similar to the target argument, which are fed to a neural network to decode a counterargument. Our work is complementary to these studies, as high-quality counterevidence is essential to decoding high-quality counterarguments.

Most of these studies build a counterargument against an *entire* argument. Thus, generated counterarguments might counter the main conclusion of the target argument without addressing specific points in it. In contrast, our work aims to find counterevidence that directly addresses specific statements in the target argument.

2.2 Fact Verification

Since we want to find counterevidence to specific statements in the target argument, our work is closely related to fact verification (Li and Zhou, 2020). Recently, this research area has garnered much attention, especially with the emergence of the FEVER (Fact Extraction and VERification) task (Thorne et al., 2018). The FEVER task aims to predict the veracity of statements, and most approaches follow three steps: document retrieval, sentence selection, and claim verification. Recent studies examined homogeneous model architectures across different steps (Tokala et al., 2019; Nie et al., 2019, 2020a). Especially BERT has been shown to be effective in both retrieval and verification (Soleimani et al., 2020), and a joint model of BERT and pointer net achieved state-of-the-art performance in this task (Hidey et al., 2020). Our work builds on this model (§4).

2.3 Knowledge-Enhanced Language Models

The last step of fact verification, i.e., claim verification, relies heavily on natural language inference (NLI) between an evidence text and a statement to verify. Recently, transformer-based language models (LMs) have been prevailing in NLI (Nie et al., 2020b), but they still show a lack of reasoning abilities (Williams et al., 2020). Hence, researchers have tried to improve LMs by integrating knowledge, mainly via two approaches.

The first is to exploit knowledge graphs (KGs) mainly to learn better embeddings of tokens and entities (Wang et al., 2021; Peters et al., 2019; Zhang et al., 2019; Lauscher et al., 2020). Once learning is done, the model does not require external

knowledge during inference. The second type of models use the triple information of entities linked to the input text during inference, either by encoding knowledge using a separate network (He et al., 2020; Chen et al., 2018) or by converting the KG to input text tokens (Liu et al., 2020). Our work adopts the second approach, as it depends less on the snapshot of the KG used for pretraining. But our model design has clear distinctions from previous work in the way that KGs are integrated and entity paths are taken.

3 Knowledge-Enhanced NLI

A natural language inference (NLI) model is the core of our entire system (Figure 1). Given a statement to refute, the system retrieves and ranks relevant documents, and then obtains a set of candidate sentences for counterevidence. For each candidate, the NLI model decides whether it entails, contradicts, or neither the statement. In this section, we first motivate the model design and explain our Knowledge-Enhanced NLI model (KENLI), followed by evaluation settings and results.

3.1 Motivation

Many NLI models have difficulty in capturing the relation between statements when their words are semantically far apart. For instance, if a statement refutes another based on example- or causality-based inference using technical terms (e.g., legume protein sources as an example of vegan food), the semantic gap between the words can make it hard to capture the relation between the two statements without explicit knowledge.

To reduce semantic gaps between words, our method aims to bridge entities in the two statements using a knowledge graph (KG) so that the information of an entity in one statement flows to an entity in the other statement, along with the information of the intermediate entities and relations on the KG. This information updates the embeddings of the tokens linked to the entities.

3.2 Model

KENLI (Figure 2 left) is based on RoBERTa-base (Liu et al., 2019), which takes a pair of premise P and hypothesis H as input and computes the probability of whether their relation is entailment, contradiction, or neutral. To bridge entities between P and H , the **Knowledge Enhancement (KE) Net** is inserted between two layers (e.g., 10th and 11th

layers), splitting RoBERTa into **Encoder1** and **Encoder2**. It updates intermediate token embeddings from Encoder1 and feeds them to Encoder2. The final prediction is made through a fully-connected layer on top of the CLS embedding.

The KE Net (Figure 2 middle) exploits a knowledge graph (KG) where nodes are entities and edges are directed relations between entities (e.g., ‘instance_of’, ‘cause’). Its main goal is to let information flow between entities in P and H through the KG. Suppose the KG has a set of relations $R = \{r_i\}_{i=1}^{|R|}$. For each input text pair, $T = \{t_i\}_{i=1}^{|T|}$ is the tokens in P that are linked to entities. Their initial embeddings $\{t_i^0\}_{i=1}^{|T|}$ are the intermediate token embeddings from Encoder1. $E = \{e_i\}_{i=1}^{|E|}$ denotes entities under consideration, with initial embeddings $\{e_i^0\}_{i=1}^{|E|}$. Considering all entities in the KG for every input pair is computationally too expensive. Recall that our motivation is to bridge entities between P and H . Hence, for each input pair, we first include entity paths whose source is in P and destination is in H . We add more destinations with the constraint that the total number of considered entities is no greater than λ and the length of each path is no greater than ν (λ and ν are hyperparameters). To obtain e_j^0 , we simply encode the name of each entity with RoBERTa Encoder1 and sum all the token embeddings.

The KE Net is a stack of KE cells. Each KE cell handles one-hop inference on the KG using two transformers TR1 and TR2. TR1 updates each entity embedding based on its neighboring entities, and TR2 updates token embeddings based on the embeddings of linked entities. More specifically, in the l -th KE cell, TR1 takes $\{e_i^{l-1}\}_{i=1}^{|E|}$ as input and updates their embeddings using self-attention. Each attention head corresponds to each relation, and the attention mask for the k -th head $M^k \in \mathbb{R}^{|E| \times |E|}$ allows information flow between entities that have the k -th relation:

$$M_{ij}^k = \begin{cases} 1 & \text{if } i = j \text{ or } (e_i, r_k, e_j) \in \text{KG} \\ 0 & \text{otherwise.} \end{cases}$$

TR2 takes the concatenation of $\{t_i^{l-1}\}_{i=1}^{|T|}$ and $\{e_i^l\}_{i=1}^{|E|}$ as input and updates the token embeddings using one attention head with attention mask $M \in \mathbb{R}^{|T+E| \times |T+E|}$:

$$M_{ij} = \begin{cases} 1 & \text{if } i \leq |T| \text{ and} \\ & (i = j \text{ or } t_i \text{ is linked to } e_{j-|T|}) \\ 0 & \text{otherwise.} \end{cases}$$

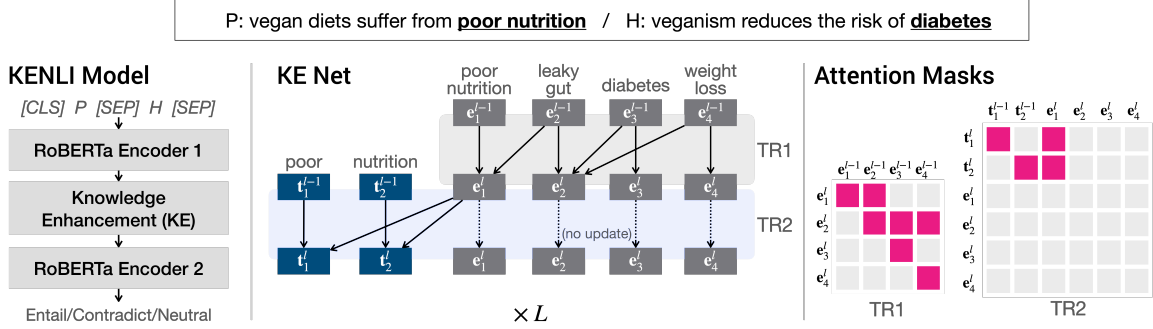


Figure 2: KENLI Model. This example illustrates using two KG paths: “poor nutrition $\xrightarrow{\text{cause}}$ leaky gut $\xrightarrow{\text{cause}}$ diabetes” and “poor nutrition $\xrightarrow{\text{cause}}$ leaky gut $\xrightarrow{\text{cause}}$ weight loss”.

Entity embeddings are not updated in TR2.

After token embeddings are updated by L KE cells (i.e., L -hop inference), the token embedding of t_i is updated as $t_i \leftarrow t_i^0 + t_i^L$ and fed to Encoder2 along with the other token embeddings in the input.

3.3 Knowledge Graphs

Our work uses two knowledge graphs: CauseNet and Wikidata. **CauseNet** (Heindorf et al., 2020) specifies *claimed* causal relations between entities, extracted from Wikipedia and ClueWeb12 based on linguistic markers of causality (e.g., “cause”, “lead”) and infoboxes. We discard entity pairs that were identified by less than 5 unique patterns, since many of them are unreliable. This results in total 10,710 triples, all having the ‘cause’ relation.

Wikidata (Vrandečić and Krötzsch, 2014) is a database that specifies a wide range of relations between entities. We use the October 2020 dump and retain triples that have 8 example-related relations: *instance_of*, *subclass_of*, *part_of*, *has_part*, *part_of_the_series*, *located_in_the_administrative_territorial_entity*, *contains_administrative_territorial_entity*, and *location*. The importance of information about physical and temporal containment in NLI was discussed recently (Williams et al., 2020). This filtering results in 95M triples, which we call **WikidataEx**.

3.4 Data

Our data mainly come from public NLI datasets: MNLI (Williams et al., 2018), ANLI (Nie et al., 2020b), SNLI (Bowman et al., 2015), and FEVER-NLI (Nie et al., 2019). We split the data into train, validation, and test sets as originally or conventionally set up for each dataset (Table 1). Due to limited computational resources, our training set includes only MNLI and ANLI.

Dataset	Train	Val	Test
MNLI	392,702	–	9,815
MNLI-MM	–	–	9,832
ANLI	162,865	3,200	3,200
SNLI	–	9,842	9,824
SNLI-Hard	–	–	3,261
FEVER-NLI	–	9,999	9,999
Example-NLI	30,133	2,867	3,468
ANLI-Contain	–	–	277
ANLI-Cause	–	–	1,078
BECauSE	–	–	2,814

Table 1: Number of NLI pairs by dataset.

The public NLI datasets alone may not include enough instances that require example- and causality-based inference. As a result, the NLI model may not learn to exploit the KGs well. To alleviate this issue, we generate synthetic NLI pairs that are built on example-based inference as follows (details are in Appendix A.2). Given a pair of P and H in the public datasets, we modify P to P' by replacing an entity that occurs in both P and H with an incoming entity on WikidataEx (e.g., “England” with “Yorkshire”). This achieves two goals. First, P' includes an entity that is an example of another entity in H so that the (P', H) pair requires example-based inference, with the same expected relation as the (P, H) pair. Second, this example relation comes from our KG so that the NLI model learns how to use the KG. Generating similar NLI pairs for causality-based inference is more challenging, and we leave it to future work.

Inference Evaluation: We use additional datasets to evaluate NLI models’ inference abilities. For example-based inference, we first use a diagnostic subset of ANLI that has been annotated with various categories of required inference, such as counting, negation, and coreference (Williams

et al., 2020). We choose the instances of the ‘Containment’ category, which requires inference on part-whole and temporal containment between entities (**ANLI-Contain**). In addition, we use the test set of our **Example-NLI** data after manually inspecting their labels.

For causality-based inference, we use the instances in the diagnostic ANLI set that belong to the ‘CauseEffect’ and ‘Plausibility’ categories (**ANLI-Cause**). They require inference on logical conclusions and the plausibility of events. In addition, we use **BECAUSE 2.0** (Dunietz et al., 2017), which specifies the ‘Cause’ and ‘Obstruct’ relations between text spans based on linguistic markers of causality. Since it has only two classes, we randomly pair up text spans to generate ‘neutral’ pairs. For reliability, we discard pairs where at least one text comprises only one word. Although this data is not for NLI, we expect that the better NLI models handle the causality between events, the better they may distinguish between the cause, obstruct, and neutral relations. See Table 1 for statistics.

3.5 Experiment Settings

For KENLI, the KE Net is inserted between the 10th and 11th layers of RoBERTa, although the location of insertion has little effect on NLI performance. The KE Net has a stack of two KE cells, allowing for 2-hop inference on a KG. We test KENLI with CauseNet (**KENLI+C**) and with WikidataEx (**KENLI+E**); we do not combine them so we can understand the utility of each KG more clearly. The maximum number of entities for each input (λ) and the maximum length of each KG path (ν) are set to 20 and 2, respectively. To see the benefit of pretraining the KE Net (as opposed to random initialization) prior to downstream tasks, we also explore pretraining it with masked language modeling on the training pairs while the original RoBERTa weights are fixed (**KENLI+E+Pt** and **KENLI+C+Pt**). The Adam optimizer is used with a learning rate of $1e-5$. See Appendix F.1 for a reproducibility checklist.

We compare KENLI with three baselines. The first two are state-of-the-art language models enhanced with knowledge graphs. **K-BERT** (Liu et al., 2020) exploits a KG during both training and inference, by verbalizing subgraphs around the entities linked to the input and combining the verbalized text into the input. **AdaptBERT** (Lauscher et al., 2020) uses a KG to enhance BERT using

bottleneck adapters (Houlsby et al., 2019); after that, it is fine-tuned for downstream tasks like normal BERT. We pretrain AdaptBERT for masked language modeling on sentences that verbalize CauseNet (10K) and a subset of WikidataEx (10M) for four epochs. We use the hyperparameter values as suggested in the papers. The last baseline is **RoBERTa**-base fine-tuned on the NLI datasets. RoBERTa trained with the ANLI dataset recently achieved a state-of-the-art performance for NLI (Nie et al., 2020b).

Input texts are linked to WikidataEx entities by the Spacy Entity Linker². CauseNet has no public entity linker, so we first stem all entities and input words using Porter Stemmer and then use exact stem matching for entity linking. The stemming allows verbs in input texts to be linked to entities (e.g., “infected–infection”, “smokes–smoking”).

3.6 Results

Table 2 shows the F1-scores of each model averaged over 5 runs with random initialization.

In the NLI evaluation, KENLI (rows 6–9) generally outperforms the baseline models (rows 1–5) across datasets. Especially KENLI with WikidataEx (rows 8–9) performs best overall and notably well for difficult datasets (SNLI-Hard, FEVER-NLI, and ANLI). This suggests that KENLI effectively incorporates example-related knowledge, which benefits prediction of nontrivial relations between statements. KENLI with CauseNet (rows 6–7) slightly underperforms KENLI+E, and its average F1-score across datasets is comparable to RoBERTa (row 5). Without pretraining (row 6), it performs slightly better than RoBERTa overall except for two difficult datasets ANLI and SNLI-Hard. With pretraining (row 7), its performance is best for the most difficult dataset ANLI, but slightly lower than or comparable to RoBERTa for the other datasets. This variance in performance across datasets makes it hard to conclude the benefit of CauseNet in general cases.

However, according to the inference evaluation, KENLI’s strength is clearer compared to other models. For example-based inference, KENLI+E (row 8) significantly outperforms the other models (ANLI-Contain) or performs comparably well (Example-NLI). Its performance is best or second-best for causality-based inference as well (ANLI-

²<https://pypi.org/project/spacy-entity-linker/>

	NLI Evaluation							Inference Evaluation			
	MNLI	MNLI-MM	ANLI	SNLI	SNLI-Hard	FEVER-NLI	Micro Avg	Example-NLI	ANLI-Contain	ANLI-Cause	BECause
AdaptBERT+C	83.0	83.6	44.7	78.8	68.2	68.2	75.4	58.4	42.3	35.0	27.6
AdaptBERT+E	83.2	83.5	44.7	78.7	68.4	67.8	75.4	58.8	43.4	34.7	27.5
K-BERT+C	83.7	83.9	45.2	80.0	70.3	68.9	76.2	58.9	42.4	34.7	27.2
K-BERT+E	83.4	83.7	46.0	79.3	69.5	69.2	76.0	59.0	44.2	35.8	26.9
RoBERTa	87.3	87.0	<u>48.6</u>	84.2	74.6	71.9	79.7	61.8	47.7	35.0	27.6
KENLI+C	<u>87.5</u>	87.1	48.2	<u>84.3</u>	74.8	71.4	79.7	<u>62.0</u>	48.2	35.1	<u>27.9</u>
KENLI+C+Pt	87.3	86.9	48.8	84.2	74.2	71.9	79.7	61.7	<u>48.4</u>	35.2	27.8
KENLI+E	87.3	87.2	48.5	84.2	75.1	72.5*	<u>79.9*</u>	61.9	49.2	<u>35.5</u>	28.0
KENLI+E+Pt	87.6	<u>87.1</u>	48.4	84.6	<u>75.1</u>	<u>72.5*</u>	80.0†	62.0	46.9	35.2	27.6

Table 2: F1-scores of NLI models by dataset. Statistical significance was measured by the paired bootstrap against the best baseline ($p < 0.05^*$, 0.01^\dagger). Bold and underline each indicate top1 and top2 results, respectively.

Cause and BECause). This suggests that the benefit of example-related knowledge is not limited to example-based inference only. Although KENLI+C (rows 6–7) shows comparable performance to RoBERTa for the general NLI tasks, it consistently outperforms RoBERTa when example- and causality-based inference is required. Example NLI pairs that are classified by only one model are shown in Table 8 in Appendix B.

Pretraining KENLI (rows 7 & 9) does not show a conclusive benefit compared to no pretraining (rows 6 & 8). Particularly for difficult datasets and inference evaluation, KENLI+E without pretraining (row 8) performs better than pretraining (row 9). The benefit of pretraining for KENLI+C varies depending on the dataset and inference task, making no substantial difference overall.

4 Retrieval of Counterevidence

Our system for counterevidence retrieval builds on DeSePtion (Hidey et al., 2020), a state-of-the-art system for the fact extraction and verification (FEVER) task (Thorne et al., 2018). As Figure 1 shows, given a statement to verify, it retrieves and ranks relevant documents, ranks candidate evidence sentences, and predicts whether the statement is supported, refuted, or neither. We adapt DeSePtion to suit our task, where the main contribution is to strengthen the last stage via our knowledge-enhanced NLI model (a detailed comparison between our system and DeSePtion is in Appendix C). We first explain individual stages and then describe evaluation settings and results.

4.1 Stages

Document Retrieval: Documents that may contain counterevidence are retrieved. Given a state-

ment to verify, we retrieve candidate documents from Wikipedia, Bing, and Google. For Wikipedia, we use the Spacy Entity Linker to retrieve the articles of Wikidata entities linked to the statement. And for each linked entity, we additionally sample at most five of their instance entities and the corresponding articles, which potentially include counterexamples to the statement. We retrieve additional Wikipedia pages by using named entities in the statement as queries for the wikipedia library³. We also conduct TF-IDF search using DrQA (Chen et al., 2017) indexed for the FEVER task. For Bing and Google, we use their search APIs. Wikipedia pages are excluded from their search results, and PDF files are processed using the pdfminer library.

Document Ranking: Retrieved documents are ranked via DeSePtion with some adaptation. First, RoBERTa is trained to predict whether each document is relevant or not, on the FEVER data. It takes the concatenation of a document snippet and the statement to verify as input. For documents from Bing and Google, we use search result snippets; for Wikipedia, we obtain snippets by concatenating the title of each Wikipedia page with its sentence that is most similar to the statement based on RoBERTa. The last embedding of the CLS token is used as the embedding of the input document, and a pointer net takes these embeddings of all documents and sequentially outputs pointers to relevant documents.

The number of retrieved documents varies depending on the search method, much fewer for the Google API than Wikipedia and Bing in general. Since this imbalance makes it difficult to compare the utility of the different search methods, we keep the number of candidate documents the same across

³<https://pypi.org/project/wikipedia/>

the methods, by ranking documents from different search methods separately and pruning low-ranked documents of Wikipedia and Bing. As a result, the three methods have the same average number of candidate documents per statement (~ 8).

Sentence Selection: For each statement to verify, we select the top 200 candidate sentences (in all ranked documents) whose RoBERTa embeddings have the highest cosine similarity to the statement.

Relation Prediction: We classify whether each candidate sentence is valid counterevidence to the statement to verify. We simply use an NLI model to compute the probability of contradiction and rank sentences by this score. The reason DeSePtion is not used here is described in Appendix C.

See Appendix F.2 for a reproducibility checklist.

4.2 Data

Input statements to our system come from two public argument datasets (Jo et al., 2020) collected from the ChangeMyView (CMV) subreddit and Kialo. On CMV, the user posts an argument, and other users attempt to refute it often by attacking specific sentences. Each sentence in an argument becomes our target of counterevidence. On Kialo, the user participates in a discussion for a specific topic and makes a statement (1–3 sentences) that either supports or attacks an existing statement in the discussion. We find counterevidence to each statement. To use our resources more efficiently, we discard CMV sentences or Kialo statements that have no named entities or Wikidata entities, since they often do not have much content to refute. We also run coreference resolution for third-person singular personal pronouns using the neuralcoref 4.0 library⁴. We randomly select 94 posts (1,599 sentences) for CMV and 1,161 statements for Kialo for evaluation.

4.3 Evaluation

We evaluate four NLI models. The first three models are directly from §3. That is, **RoBERTa** is fine-tuned on the NLI data. **KENLI+C** and **KENLI+E** are trained with CauseNet and WikidataEx, respectively, without pretraining. The last baseline is **LogBERT**, a state-of-the-art model for argumentative relation classification (Jo et al., 2021). Given a

⁴<https://github.com/huggingface/neuralcoref>

pair of statements, it predicts whether the first statement supports, attacks, or neither the second statement based on four logical relations between them, namely, textual entailment, sentiment, causal relation, and normative relation⁵. Since LogBERT captures the support and attack relations beyond textual entailment, this baseline would show whether NLI is sufficient for finding counterevidence.

We collect a ground-truth set of labeled data using MTurk. First, for each statement to refute, we include in the ground-truth set the top candidate sentence from each model if the probability of ‘contradiction’ is ≥ 0.5 (i.e., max four sentences). As a result, the ground-truth set consists of 4,783 (CMV) and 3,479 (Kialo) candidate sentences; they are challenging candidates because at least one model believes they are valid counterevidence.

Each candidate sentence is scored by two Turkers with regard to how strongly it refutes the statement (very weak=0, weak=1, strong=2, and very strong=3). Each candidate sentence is displayed with the surrounding sentences in the original source document as context, as well as a link to the source document. If a candidate is scored as both very weak and very strong, these scores are considered unreliable, and the candidate is scored by a third Turker. For each candidate, the mean score s is taken as the ground-truth validity as counterevidence: ‘valid’ if $s \geq 1.5$ and ‘invalid’ otherwise. More details are described in Appendix A.3.

According to the additional question of whether reading the source document is necessary to make a decision for each candidate, about 40% of answers and 65% of candidates required reading source documents. This might indicate that three sentences are insufficient for making robust decisions about counterevidence, but it could also be the case that, since our system checks all documents and filter them by relevance in earlier stages, it would not benefit much from more than three sentences.

We use four evaluation metrics on the ground-truth set. **Precision**, **recall**, and **F1-score** are computed based on whether the model-predicted probability of contradiction for each candidate is ≥ 0.5 . These metrics, however, make the problem binary classification, missing the nuanced degree of validity for each candidate. Thus, we measure **Kendall’s τ** between mean validity scores from

⁵For implementation, BERT-base is fine-tuned for the four classification tasks and then for argumentative relation classification on the Kialo arguments (both normative and non-normative) in the original paper.

	CMV				Kialo			
	Prec	RecL	F1	τ	Prec	RecL	F1	τ
RoBERTa	48.3	63.6	54.9	0.002	58.0	57.0	57.5	0.022
KENLI+C	48.8	<u>65.0</u>	55.8	0.014	<u>59.0</u>	62.2 [‡]	60.6 [‡]	0.038 [†]
KENLI+E	<u>48.9</u>	71.3[‡]	58.0[‡]	0.015	58.0	65.2 [‡]	61.4 [‡]	0.038 [†]
LogBERT	51.4[†]	61.8	<u>56.1</u>	0.031*	60.0	66.2[‡]	62.9[‡]	0.045[†]

Table 3: Accuracy of evidence retrieval. For precision, recall, and F1-score, statistical significance was calculated using the paired bootstrap against RoBERTa; for Kendall’s τ , the statistical significance of each correlation value was calculated ($p < 0.05^*$, 0.01^\dagger , 0.001^\ddagger).

human judgments and each model’s probability scores. High τ indicates a good alignment between the human judgment and the model judgment about the strength of validity of each candidate.

4.4 Results

Table 3 summarizes the accuracy of evidence retrieval. Both KENLI+C (row 2) and KENLI+E (row 3) outperform RoBERTa (row 1) for both CMV and Kialo. The motivation behind KENLI was to capture statement pairs that require complex inference, by bridging entities with KGs. As expected, KENLI identifies more instances of contradiction that are missed by RoBERTa, as indicated by its high recall. The recall of KENLI+E is substantially higher than RoBERTa’s by 7.7 and 8.3 points for CMV and Kialo, respectively, while its improvement of precision is relatively moderate. KENLI+C has a similar pattern but with a smaller performance gap with RoBERTa.

To see if KENLI+E indeed effectively captures counterevidence that requires example-based inference, we broke down its F1-score into one measured on candidate sentences for which KG paths exist between their tokens and the statement’s tokens and one measured on the other candidate sentences with no connecting KG paths (Figure 3). The F1-score gap between KENLI+E and RoBERTa is substantially higher for the candidate sentences where KG paths exist. The gap of recall is even higher, indicating that KENLI+E indeed captures complex counterevidence more effectively than RoBERTa. KG paths that benefit KENLI+E the most include “player PART_OF game”, “Tel Aviv District LOCATED_IN Israel”, and “neurovascular system HAS_PART brain”.

LogBERT slightly underperforms KENLI+E for CMV, but it outperforms KENLI+E for Kialo, possibly because LogBERT is trained on arguments

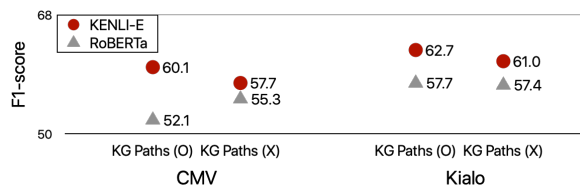


Figure 3: F1-scores of KENLI+E and RoBERTa by the existence of KG paths in candidate sentences.

from Kialo and may learn useful linguistic information in Kialo. While KENLI+E has relatively high recall, LogBERT is notable for high precision compared to the other models. This is somewhat counterintuitive because LogBERT uses four logical relations between two statements, which one might expect to improve recall by capturing a broad range of mechanisms for contradiction. In reality, however, LogBERT seems to make conservative predictions based on whether strong signals exist for the logical relations. Combining the two models may result in high recall by incorporating different KGs and, at the same time, improve precision by incorporating different types of signals (e.g., sentiment). For example, KENLI could be pretrained on the four logical mechanisms in the same way that LogBERT is. Alternatively, we could incorporate the KE Net in the middle of LogBERT. We leave this direction to future work.

We conducted a further analysis on how differently KENLI+E and LogBERT behave. We find that LogBERT excels when a pair of statement and candidate sentence has strong signals for four logical relations—textual contradiction, negative sentiment, obstructive causal relation, and refuting normative relation—that have been found to have high correlations with LogBERT’s decision of ‘contradiction’ (Jo et al., 2021). For instance, for CMV, when we focus only on the pairs whose candidate sentences express negative sentiment toward their target statements with probability greater than 0.5, LogBERT’s recall and F1-score substantially increase to 96.5 and 65.3, respectively, whereas KENLI+E’s recall and F1-score drop to 63.1 and 56.9. This means that negative sentiment toward the target statement is a useful cue for counterevidence, and LogBERT uses this signal better than KENLI. The caveat, however, is LogBERT’s F1-score drops significantly to 49.9 for candidate sentences not expressing negative sentiment, while KENLI+E’s F1-score remains stable at 58.4. A similar pattern occurs for obstructive causal rela-

tion and textual contradiction, and for Kialo as well. More details are in Appendix D.

According to Kendall’s τ , LogBERT shows the best alignment with human judgments on the validity scores of candidate sentences among the four models. However, correlations overall are rather weak, ranging between 0.002 and 0.045. One possible reason is that the models are trained for binary classification, not to predict the degree of validity. As a result, 75% of KENLI+E’s probability scores are skewed toward > 0.9 or < 0.1 versus only 10% of human scores (after normalized between 0 and 1). Therefore, KENLI+E’s probability scores seem to capture something different than human scores by design. KENLI may have a better alignment with human scores if it is trained explicitly on the degree of validity using regression, though collecting such data would be expensive. Another reason for the low correlation is that sometimes a candidate sentence’s validity is vague due to limited context (e.g., Statement: *However, the similarities end there.* / Candidate Sentence: *However, the similarities do not end there.*). In such cases, human scores tend to lie in the middle area (0.33 for this example), whereas KENLI still makes a confident decision (0.99 for this example) within the limited context. We also find that KENLI tends to overpredict ‘valid’ as a candidate sentence and the target statement share more words. Such an overreliance on overlapping words could exacerbate the misalignment with human judgments.

We conducted further analyses on the utility of document types. Previous work on counterargument generation and fact extraction/verification relies heavily on Wikipedia and sometimes news articles. However, we find that valid counterevidence resides in more various sources. While knowledge archives (e.g., Wikipedia, lectures) take the highest proportion (27–37%), counterevidence resides in mainstream news, personal blogs, research journals, etc. as well (Figure 8 in Appendix E.2). Moreover, if we break down model accuracy into different document types, the models are more reliable (i.e., achieve higher F1-scores) for specialized magazines and Q&A forums than knowledge archives (Figure 9 in Appendix E.2). If we break down model accuracy by search methods, Wikipedia achieves lower scores than Bing and Google across all metrics (Table 10 in Appendix E.1). These results suggest that counterargument generation and fact extraction/verification should consider more

diverse sources of evidence beyond Wikipedia. See Appendices E.2 and E.1 for more details.

In §1, we assumed the scenario where one makes a counterargument by first detecting attackable points in the target argument and then retrieving counterevidence to those points. To see the feasibility of automating this pipeline, we took an existing model that aims to detect attackable sentences in arguments (Jo et al., 2020) and analyzed whether this model can identify sentences that have counterevidence according to our collected data. We find that attackability scores predicted by this model tend to be higher for argument sentences for which we were able to find counterevidence (Figure 10 in Appendix E). This result suggests that computational models for attackability detection and our counterevidence retrieval system could create synergy to fully automate counterargument generation. See Appendix E.3 for more details.

5 Conclusion

In this paper, we built a counterevidence retrieval system. To retrieve counterevidence that involves complex inference, we presented a knowledge-enhanced NLI model with specific focus on causality- and example-based inference. The NLI model demonstrates improved performance for NLI tasks, especially for instances that require the targeted inference. Integrating the NLI model into the retrieval system further improves counterevidence retrieval performance, especially recall, showing the effectiveness and utility of our method of incorporating knowledge graphs in NLI.

Acknowledgments

This work was partly supported by CMU’s GuSH funding. JinYeong Bak is partly supported by Institute of Information & communications Technology Planning & Evaluation (IITP) grant funded by the Korea government (MSIT) (No.2019-0-00421, AI Graduate School Support Program (Sungkyunkwan University)) and partly by the MSIT, Korea, under the High-Potential Individuals Global Training Program (2020-0-01550) supervised by the IITP. Alice Oh and Haneul Yoo are supported by Institute of Information & communications Technology Planning & Evaluation (IITP) grant funded by the Korea government (MSIT) (No.2017-0-01779, A machine learning and statistical inference framework for explainable artificial intelligence).

References

- Samuel R. Bowman, Gabor Angeli, Christopher Potts, and Christopher D. Manning. 2015. [A large annotated corpus for learning natural language inference](#). In *Proceedings of the 2015 Conference on Empirical Methods in Natural Language Processing*, pages 632–642, Lisbon, Portugal. Association for Computational Linguistics.
- Danqi Chen, Adam Fisch, Jason Weston, and Antoine Bordes. 2017. [Reading Wikipedia to Answer Open-Domain Questions](#). In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1870–1879, Vancouver, Canada. Association for Computational Linguistics.
- Qian Chen, Xiaodan Zhu, Zhen-Hua Ling, Diana Inkpen, and Si Wei. 2018. [Neural Natural Language Inference Models Enhanced with External Knowledge](#). In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 2406–2417.
- Zhuyun Dai and Jamie Callan. 2019. [Deeper Text Understanding for IR with Contextual Neural Language Modeling](#). In *Proceedings of the 42nd International ACM SIGIR Conference on Research and Development in Information Retrieval*, page 985–988, New York, NY, USA. Association for Computing Machinery.
- Jesse Dunietz, Lori Levin, and Jaime Carbonell. 2017. [The BECause Corpus 2.0: Annotating Causality and Overlapping Relations](#). In *Proceedings of the 11th Linguistic Annotation Workshop*, pages 95–104.
- Eveline T Feteris, Feteris, and Olivier. 2017. *Fundamentals of legal argumentation*, volume 1. Springer.
- Bin He, Di Zhou, Jinghui Xiao, Xin Jiang, Qun Liu, Nicholas Jing Yuan, and Tong Xu. 2020. [BERT-MK: Integrating Graph Contextualized Knowledge into Pre-trained Language Models](#). In *Findings of the Association for Computational Linguistics: EMNLP 2020*, pages 2281–2290, Online. Association for Computational Linguistics.
- Stefan Heindorf, Yan Scholten, Hening Wachsmuth, Axel-Cyrille Ngonga Ngomo, and Martin Potthast. 2020. [CauseNet: Towards a Causality Graph Extracted from the Web](#). In *Proceedings of the 29th ACM International Conference on Information and Knowledge Management*.
- Christopher Hidey, Tuhin Chakrabarty, Tariq Alhindi, Siddharth Varia, Kriste Krstovski, Mona Diab, and Smaranda Muresan. 2020. [DeSePtion: Dual Sequence Prediction and Adversarial Examples for Improved Fact-Checking](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 8593–8606, Online. Association for Computational Linguistics.
- Neil Houlsby, Andrei Giurgiu, Stanislaw Jastrzebski, Bruna Morrone, Quentin De Laroussilhe, Andrea Gesmundo, Mona Attariyan, and Sylvain Gelly. 2019. [Parameter-Efficient Transfer Learning for NLP](#). In *Proceedings of the 36th International Conference on Machine Learning*, volume 97, pages 2790–2799, Long Beach, California, USA. PMLR.
- Xinyu Hua, Zhe Hu, and Lu Wang. 2019. [Argument Generation with Retrieval, Planning, and Realization](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 2661–2672. Association for Computational Linguistics.
- Xinyu Hua and Lu Wang. 2018. [Neural Argument Generation Augmented with Externally Retrieved Evidence](#). In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 219–230.
- Yohan Jo, Seojin Bang, Eduard Hovy, and Chris Reed. 2020. [Detecting Attackable Sentences in Arguments](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 1–23. Association for Computational Linguistics.
- Yohan Jo, Seojin Bang, Chris Reed, and Eduard Hovy. 2021. [Classifying Argumentative Relations Using Logical Mechanisms and Argumentation Schemes](#). *Transactions of the Association for Computational Linguistics*.
- Anne Lauscher, Olga Majewska, Leonardo F R Ribeiro, Iryna Gurevych, Nikolai Rozanov, and Goran Glavaš. 2020. [Common Sense or World Knowledge? Investigating Adapter-Based Knowledge Injection into Pretrained Transformers](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing*.
- Dieu Thu Le, Cam-Tu Nguyen, and Kim Anh Nguyen. 2018. [Dave the debater: a retrieval-based and generative argumentative dialogue agent](#). In *Proceedings of the 5th Workshop on Argument Mining*, pages 121–130.
- Qifei Li and Wangchunshu Zhou. 2020. [Connecting the Dots Between Fact Verification and Fake News Detection](#). In *Proceedings of the 28th International Conference on Computational Linguistics*, pages 1820–1825, Barcelona, Spain (Online). International Committee on Computational Linguistics.
- Weijie Liu, Peng Zhou, Zhe Zhao, Zhiruo Wang, Qi Ju, Haotang Deng, and Ping Wang. 2020. [K-BERT: Enabling Language Representation with Knowledge Graph](#). In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 34, pages 2901–2908.
- Yinhan Liu, Myle Ott, Naman Goyal, Jingfei Du, Mandar Joshi, Danqi Chen, Omer Levy, Mike Lewis, Luke Zettlemoyer, and Veselin Stoyanov. 2019. [Roberta: A robustly optimized bert pretraining approach](#).

- Yixin Nie, Lisa Bauer, and Mohit Bansal. 2020a. [Simple Compounded-Label Training for Fact Extraction and Verification](#). In *Proceedings of the Third Workshop on Fact Extraction and VERification (FEVER)*, pages 1–7.
- Yixin Nie, Haonan Chen, and Mohit Bansal. 2019. [Combining Fact Extraction and Verification with Neural Semantic Matching Networks](#). In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 33, pages 6859–6866.
- Yixin Nie, Adina Williams, Emily Dinan, Mohit Bansal, Jason Weston, and Douwe Kiela. 2020b. [Adversarial NLI: A New Benchmark for Natural Language Understanding](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 4885–4901.
- Matan Orbach, Yonatan Bilu, Assaf Toledo, Dan Lahav, Michal Jacovi, Ranit Aharonov, and Noam Slonim. 2020. [Out of the Echo Chamber: Detecting Countering Debate Speeches](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 7073–7086.
- Matthew E Peters, Mark Neumann, Robert Logan, Roy Schwartz, Vidur Joshi, Sameer Singh, and Noah A Smith. 2019. [Knowledge Enhanced Contextual Word Representations](#). 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 43–54.
- Paul Reiser, Naoya Inoue, Naoaki Okazaki, and Kentaro Inui. 2015. [A Computational Approach for Generating Toulmin Model Argumentation](#). In *Proceedings of the 2nd Workshop on Argumentation Mining*, pages 45–55.
- Misa Sato, Kohsuke Yanai, Toshinori Miyoshi, Toshihiko Yanase, Makoto Iwayama, Qinghua Sun, and Yoshiki Niwa. 2015. [End-to-end Argument Generation System in Debating](#). In *Proceedings of ACL-IJCNLP 2015 System Demonstrations*, pages 109–114.
- Noam Slonim. 2018. [Project Debater](#). In *Computational Models of Argument*, volume 305, pages 4–4.
- Amir Soleimani, Christof Monz, and Marcel Worring. 2020. [BERT for Evidence Retrieval and Claim Verification](#). In *Proceedings of the 42nd European Conference on IR Research*, pages 359–366.
- 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.
- Santosh Tokala, Vishal G, Avirup Saha, and Niloy Ganguly. 2019. [AttentiveChecker: A Bi-Directional Attention Flow Mechanism for Fact Verification](#). 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 2218–2222, Minneapolis, Minnesota. Association for Computational Linguistics.
- Denny Vrandečić and Markus Krötzsch. 2014. [Wiki-data: a free collaborative knowledgebase](#). *Communications of the ACM*, 57(10):78–85.
- Henning Wachsmuth, Shahbaz Syed, and Benno Stein. 2018. [Retrieval of the Best Counterargument without Prior Topic Knowledge](#). In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 241–251.
- Douglas Walton, Chris Reed, and Fabrizio Macagno. 2008. *Argumentation Schemes*. Cambridge University Press.
- Xiaozhi Wang, Tianyu Gao, Zhaocheng Zhu, Zhengyan Zhang, Zhiyuan Liu, Juanzi Li, and Jian Tang. 2021. [KEPLER: A Unified Model for Knowledge Embedding and Pre-trained Language Representation](#). *Transactions of the Association for Computational Linguistics*, 9:176–194.
- 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.
- Adina Williams, Tristan Thrush, and Douwe Kiela. 2020. [ANLIzing the Adversarial Natural Language Inference Dataset](#). *arXiv*.
- Bronwyn Woods, David Adamson, Shayne Miel, and Elijah Mayfield. 2017. [Formative Essay Feedback Using Predictive Scoring Models](#). In *Proceedings of the 23rd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, pages 2071–2080.
- Zhengyan Zhang, Xu Han, Zhiyuan Liu, Xin Jiang, Maosong Sun, and Qun Liu. 2019. [ERNIE: Enhanced Language Representation with Informative Entities](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 1441–1451.

Appendix

A Annotation Tasks

A.1 Annotation Principle

For all annotation tasks, we recruited annotators on Amazon Mechanical Turk (MTurk). Participants should meet the following qualifications: (1) residents of the U.S., (2) at least 500 HITs approved, and (3) HIT approval rate greater than 97%. Each HIT includes several questions and one attention question. The attention question asks the annotator to select a specific option, and we rejected and discarded a HIT if the annotator failed the attention question. For all annotation tasks, the annotation manuals are publicly available.

A.2 Annotation of Example-Based NLI data

This section describes our method for synthetic building of example-based NLI data that was augmented with the public NLI datasets in our experiments. The entire process consists of two steps. First, we generate synthetic example-based pairs using a pretrained language model (§A.2.1). Next, we annotate their labels using MTurk (§A.2.2).

A.2.1 Generating Synthetic NLI Pairs

We synthetically generate example-based NLI pairs as follows. Given a pair of P and H in the public datasets in Table 1, we modify P to P' by replacing an entity that occurs in both P and H with an incoming entity on WikidataEx. For example, in the following pair

P : a breakdancer man is performing for the kids at school

H : a man is break dancing at a school

“school” occurs in both P and H , so we may generate P' by replacing “school” with an instance of the school (e.g., “preschool”) based on WikidataEx. To avoid broken or implausible sentences, we retain P' only if its perplexity is lower than or equal to that of P based on GPT2. Table 4 shows examples of synthetically generated P' and their perplexity. P is the original statement from the SNLI dataset, and P'_1 – P'_4 are generated statements after the entity “school” is replaced. The perplexity of P'_1 and P'_2 is lower than that of the original statement P , so we pair each of them with H and add the pairs to our synthetic NLI data. However, P'_3 and P'_4 are discarded because their perplexity is higher than that of P .

ID	Statement	Perplexity
P	a breakdancer man is performing for the kids at school	3.08
P'_1	a breakdancer man is performing for the kids at licensed victuallers’ school	2.67
P'_2	a breakdancer man is performing for the kids at preschool	2.95
P'_3	a breakdancer man is performing for the kids at boys republic	3.09
P'_4	a breakdancer man is performing for the kids at language teaching	3.10

Table 4: Examples of generated example-based statement and its perplexity measured by GPT2.

		Original Label		
		Entail	Neutral	Contradict
New Label	Entail	1,698	548	373
	Neutral	228	543	139
	Contradict	151	302	1,030
	Broken	20	15	24
	No Majority	336	333	291

Table 5: Confusion matrix of example-based NLI data labels.

A.2.2 Label Annotation

For each of the generated NLI pairs, we ask annotators whether H is correct or wrong given the context P' . They can choose from the four options: definitely correct (entail), definitely wrong (contradict), neither definitely correct nor definitely wrong (neutral), and broken English (Figure 4). Each HIT consists of 10 pairs and one attention question. Each pair is labeled by three annotators and is discarded if the three annotators all choose different labels.

A.2.3 Analysis

To see how the labels of the generated pairs (P' , H) differ from the labels of their original pairs (P , H), we manually analyzed 6,031 pairs (Table 5). Only 59 sentences were labeled as broken, meaning that our GPT2-based generation method effectively generates sensible statements P' . Most original pairs of entailment and contradiction keep their labels, but many of originally neutral pairs turn to either entailment or contradiction after entity replacement.

A.3 Annotation of Evidence Validity

In this task, annotators were asked to mark how strongly each counterevidence candidate sentence refutes the statement it attempts to refute (i.e., statements from CMV or Kialo). The four options of

Context

a breakdancer man is performing for the kids at licensed victuallers' school

Statement

A man is break dancing at a school

Question

Given the context, the statement is ?

Definitely Correct
Definitely Wrong
Neither Definitely Correct nor Definitely Wrong
Broken

Figure 4: Example-based NLI labeling page.

Argument

Cultural appropriation is actually a good thing. I personally have learned and been inspired by many people of different cultures. My honest opinion is "**Cultural appropriation**" is a product of progress, or progress is the product of cultural appropriation. Whichever comes first, I find it hard to believe they are not connected. I come from an insular town called Vidor, Texas. We are unfortunately known for being super white and racist. In my friend group, I am now the minority. I wouldn't have been able to get from a family of confederate sympathizers, to the amazing relationships I have without several forms of cultural appropriation. I can't imagine how a world where more people remain segregated is better for anyone.

Refuting Evidence 1

This can be controversial when members of a dominant culture appropriate from disadvantaged minority cultures. **According to critics of the practice, cultural appropriation differs from acculturation, assimilation, or equal cultural exchange in that this appropriation is a form of colonialism.** When cultural elements are copied from a minority culture by members of a dominant culture, these elements are used outside of their original cultural context—sometimes even against the expressly stated wishes of members of the originating culture. [See the full reference \(en.wikipedia.org\)](#)

Questions

- How strongly does this evidence **REFUTE** the highlighted sentence in the argument?

Very Weak	Weak	Strong	Very Strong
-----------	------	--------	-------------
- Was it necessary to follow the reference to make your decision?

Not Necessary	Necessary
---------------	-----------

Refuting Evidence 2

Dear worker, read this instruction carefully. This item has been inserted to check if you really pay attention to the content of evidence. Answer **Strong** and **Necessary** for this question. If you miss this question, we cannot trust your answers in this HIT and will not pay you for this HIT. [See the full reference \(en.wikipedia.org\)](#)

Questions

- How strongly does this evidence **REFUTE** the highlighted elements in the argument?

Very Weak	Weak	Strong	Very Strong
-----------	------	--------	-------------
- Was it necessary to follow the reference to make your decision?

Not Necessary	Necessary
---------------	-----------

Figure 5: CMV evaluation page. Refuting Evidence 2 is an attention question.

strength are ‘very weak’, ‘weak’, ‘strong’, and ‘very strong’, with corresponding scores 0, 1, 2, and 3 (Figures 5 and 6). For each statement from CMV, the entire post is displayed with the target statement highlighted so the annotator can consider the context of the statement when making a decision. For each candidate sentence, the annotators should also answer whether reading the source document is necessary to make a judgment.

Each HIT includes four statements to refute, along with at most four candidate counterevidence sentences for each statement, and one attention question. Each candidate sentence was labeled by two annotators. If a candidate sentence was labeled as both ‘very weak’ and ‘very strong’, we treated the labels as unreliable (146 candidates in 131 sen-

tences from CMV, 71 candidates in 65 statements from Kialo) and allocated a third annotator. We average their scores, which becomes the candidate sentence’s final strength. The average variance of scores for each candidate sentence is 0.48, meaning that annotators on average have a score difference less than 1 point.

A.4 Annotation of Document Types

In this task, we annotate the type of source document for each candidate sentence. Each annotator was shown the network location identifier of a URL (e.g., “www.cnn.com”, “docs.python.org”) and asked to choose the type of the site from 14 categories (Table 6 and Figure 7). Total 1,987 unique location identifiers were annotated. Each HIT con-

Argument

Slavery was implicitly permitted in the original Constitution but removed by the 13th Amendment.

Refuting Evidence 1

W. E. B. Du Bois wrote in 1935: **Slavery was not abolished even after the Thirteenth Amendment.** There were four million freedmen and most of them on the same plantation, doing the same work they did before emancipation, except as their work had been interrupted and changed by the upheaval of war.

[See the full reference \(en.wikipedia.org\)](#)

Question

- How strongly does this evidence **REFUTE** the argument?

Very Weak Weak Strong Very Strong

- Was it necessary to follow the reference to make your decision?

Not Necessary Necessary

Refuting Evidence 2

Dear worker, read this instruction carefully. This item has been inserted to check if you really pay attention to the content of evidence. Answer **Strong** and **Not Necessary** for this question. If you miss this question, we cannot trust your answers in this HIT and will not pay you for this HIT.

[See the full reference \(en.wikipedia.org\)](#)

Question

- How strongly does this evidence **REFUTE** the argument?

Very Weak Weak Strong Very Strong

- Was it necessary to follow the reference to make your decision?

Not Necessary Necessary

Figure 6: Kialo evaluation page. Refuting Evidence 2 is an attention question.

Question

Click the following URL and choose its category. Refer to the example pages if necessary.

www.edge.org

Example Pages

- <https://www.edge.org/responses/what-do-you-think-about-machines-that-think>
- <https://www.edge.org/responses/q2013>
- <https://www.edge.org/responses/how-is-the-internet-changing-the-way-you-think>
- <https://www.edge.org/responses/what-scientific-idea-is-ready-for-retirement>
- <https://www.edge.org/responses/what-is-your-favorite-deep-elegant-or-beautiful-explanation>

Mainstream News

Research Journal

Report

Personal Blog

Magazine-psychology

Magazine-society

Magazine-finance

Magazine-culture

Magazine-scitech

Magazine-general

Knowledge Archives

Q&A

Forum

Broken

Figure 7: Document type annotation page.

sists of 10 identifiers and one attention question. Each identifier was annotated until two annotators chose the same category. If there was no such category for five annotators, we selected the decision of the most “trustworthy” annotator, who had the highest rate of decisions selected for other identifiers.

A.5 Ethical Considerations on Human Annotation

We consider ethical issues on our annotation tasks. The first consideration is fair wages. We compute the average time per HIT based on a small pilot study, and set the wage per HIT to be above the federal minimum wage in the U.S. (\$7.25⁶). Table 7

⁶<https://www.dol.gov/general/topic/wages/minimumwage>

shows that the expected hourly wage is higher than the federal minimum wage for all the annotation tasks.

We also preserve the privacy of crowdworkers. We do not ask for their personal information, such as names and gender. We collect Worker IDs to map each HIT result with the annotator and to accept or reject their work on MTurk. But the Worker IDs are discarded afterward to preserve their privacy.

Our annotation tasks are upfront and transparent with annotators. We provide the instruction manual of each task at the starting page, which informs the annotators of various task information, such as an estimated time needed for the task. Some annotators complained when their work was rejected. We generally responded within a business day with

Type	Description	Examples
Mainstream News	Mainstream news about daily issues and general topics.	www.cnn.com, www.bbc.com
Research Journal	Peer-reviewed papers or dissertations.	link.springer.com, www.nature.com
Report	Surveys, statistics, and reports. Should be a source of substantial data rather than a summary of reports.	www.whitehouse.gov, www.irs.gov, www.cdc.gov
Personal Blog	Personal blogs.	medium.com, jamesclear.com
Magazine–Psychology	Magazines about psychology, mental health, relationships, family.	www.psychologytoday.com
Magazine–Society	Magazines about social and political issues.	www.hrw.org, www.pewresearch.org
Magazine–Finance	Magazines about finance, business, management.	www.hbr.org
Magazine–Culture	Magazines about culture, education, entertainment, fashion, art.	www.vulture.com
Magazine–Scitech	Magazines about science, medicine, technology.	www.techdirt.com, www.webmd.com
Magazine–General	Magazines about multiple domains.	thedickinsonian.com
Knowledge Archives	Information archives for knowledge transfer, such as encyclopedias, books, dictionaries, lectures.	plato.stanford.edu, quizlet.com, www.wikihow.com
Q&A	Question and answering platforms.	stackoverflow.com, www.quora.com
Forum	Forums for opinion sharing and reviews.	www.reddit.com, www.debate.org
Broken	URLs are not accessible.	

Table 6: Evidence document types.

Task	# Questions/HIT	Time/HIT (secs)	Wage/HIT	Expected Hourly Wage
Example-based NLI	10	324	\$0.7	\$7.78
Evidence Validity – CMV	4	247	\$0.5	\$7.28
Evidence Validity – Kialo	4	240	\$0.5	\$7.50
Document Type	10	351	\$0.5	\$7.79

Table 7: Expected hourly wage of each annotation task. All wages are over the federal minimum wage in the U.S. (\$7.25). The number of questions per HIT does not include attention questions.

evidence of our decision (i.e., their failure at the attention question).

B Example NLI Pairs

Table 8 shows example NLI pairs that are classified correctly by only one model.

C Counterevidence Retrieval System

Document Retrieval: In this stage, documents that may contain counterevidence are retrieved. Given a statement to verify, DeSePtion retrieves candidate documents from Wikipedia in four ways: (1) using named entities in the statement as queries for the wikipedia library⁷, (2) using the statement as a query for Google, (3) TF-IDF search using DrQA (Chen et al., 2017), and (4) some heuristics. Note that all documents are from Wikipedia, in accordance with the FEVER task.

We make several adaptations that better suit our task. First, in addition to Wikipedia articles, we also retrieve web documents using Microsoft Bing and Google (wikipedia pages are excluded from

their search results). The three sources provide documents with somewhat different characteristics, and we compare their utility in Appendix E. Second, we use the Spacy Entity Linker to retrieve the articles of Wikidata entities linked to the statement. And for each linked entity, we additionally sample at most five of their instance entities and the corresponding articles. These expanded articles potentially include counterexamples to the statement⁸. Lastly, we do not use the heuristics.

Document Ranking: Given a set of candidate documents, DeSePtion ranks them using a pointer net combined with fine-tuned BERT. First, BERT is trained to predict whether each document is relevant or not, using the FEVER dataset; it takes the concatenation of the page title and the state-

⁸We considered retrieving web documents in a similar way, using query expansion, but ended up not doing it. One reason is that search engines already include example-related documents to some extent. For instance, for the query “Vegan diets can cause cancer”, Bing returns a document with the title “Can the Keto and Paleo Diets Cause Breast Cancer?”. Another practical reason is that query expansion requires arbitrarily many search transactions that are beyond the capacity of our resources.

⁷<https://pypi.org/project/wikipedia/>

Correct Model	NLI Pair
AdaptBERT+C	<p>P: How to blow a bubble with bubblegum
Buy some bubblegum. You can buy gum at pretty much every corner store. Chewing gums can be used to make bubbles, but they won't be as big, and they'll usually pop too easily. / H: Kids can learn to blow bubbles.</p> <p>P: Letters in support or condemnation of the QES program (though one may assume they will insist on programme) should be addressed to Mrs Anne Shelley, Secretary, Queen's English Society, 3 Manor Crescent, Guildford GU2 6NF, England. / H: Mrs. Anne Shelley is in charge of the QES program.</p> <p>P: A woman in black walks down the street in front of a graffitied wall. / H: A young woman is standing and staring at a painted mural.</p>
AdaptBERT+E	<p>P: Donna Noble is a fictional character in the long-running British science fiction television series Doctor Who . Portrayed by British actress and comedian Catherine Tate , she is a companion of the Tenth Doctor (David Tennant) . / H: Donna Noble is the therapist of the Doctor.</p> <p>P: Blond boy in striped shirt on the swing set. / H: A boy in a blue striped shirt is on the swing set.</p> <p>P: The diocese of Vannida (in Latin: Dioecesis Vannidensis) is a suppressed and titular See of the Roman Catholic Church. It was centered on the ancient Roman Town of Vannida, in what is today Algeria, is an ancient episcopal seat of the Roman province of Mauritania Cesariense. / H: The diocese of Vannida is located in Europe</p>
K-BERT+C	<p>P: Second, a clue may name a class of objects which includes the answer, like bird for COCK. / H: Clues may not be categorical.</p> <p>P: Asia has the highest number of child workers, but Sub-Saharan Africa has the highest proportion of working children relative to population. / H: There are more children in Asia than there are in Sub-Saharan Africa.</p> <p>P: Duende meant 'hobgoblin,' 'sprite,' or 'ghost' in Spanish for a long time, but it is not known when it acquired its artistic coloration. / H: Duende has had many meanings, but all were similar.</p>
K-BERT+E	<p>P: Well, the issue before this Court, I hasten to say, as I said before, is only whether, once the Congress makes that judgment, it can ever change it retrospectively. The issue before this Court is not whether, in the future, a certain length of time would be appropriate. / H: Congress will change the decision of a judgement without considering when the decision was made.</p> <p>P: John Goodman . His other film performances include lead roles in The Babe (1992) , The Flintstones (1994) and 10 Cloverfield Lane (2016) and supporting roles in Coyote Ugly (2000) , The Artist (2011) , Extremely Loud and Incredibly Close (2011) , Argo (2012) , Flight (2012) , The Hangover Part III (2013) , and Patriots Day (2016) . / H: John Goodman played Babe Ruth in The Babe.</p> <p>P: Man chopping wood with an axe. / H: The man is outside.</p>
RoBERTa	<p>P: Two men wearing dirty clothing are sitting on a sidewalk with their dog and begging for money using a cardboard sign. / H: These two illiterate men convinced their dog to write a panhandling sign for them.</p> <p>P: Sure enough, there was the chest, a fine old piece, all studded with brass nails, and full to overflowing with every imaginable type of garment. / H: The chest wasn't big enough to completely contain all of the garments.</p> <p>P: A man is spinning a little girl in the air above his head. / H: A man is carrying a little girl off the ground.</p>
KENLI+C	<p>P: "Nine Lives" is a song by American hard rock band Aerosmith. It was released in 1997 as the lead single and title track from the album "Nine Lives". The song was written by lead singer Steven Tyler, guitarist Joe Perry, and songwriter Marti Frederiksen. The song is four minutes, one second long. All the high-caliber guitar solos are played by Brad Whitford. / H: The song "Nine Lives" was released in 1997 and performed only by Brad Whitford.</p> <p>P: The last stages of uploading are like a mental dry-heave. / H: Uploading your consciousness feels like a mental dry-heave in the final stages.</p> <p>P: Cheerleaders in blue performing on a football field underneath a yellow football goal post. / H: A group of girls shake pom poms.</p>
KENLI+E	<p>P: The Ron Clark Story is a 2006 television film starring Matthew Perry. The film is based on the real-life educator Ron Clark. It follows the inspiring tale of an idealistic teacher who leaves his small hometown to teach in a New York City public school, where he faces trouble with the students. The film was directed by Randa Haines, and was released directly on television. / H: Ron Clark's hometown was not New York City.</p> <p>P: These social encounters oer children many opportunities to hear people refer to their own mental states and those of others and, therefore, to observe diereent points of view. / H: The social encounters give kids a lot of chances to hear people talk about their mental states and how things make them feel.</p> <p>P: An astute mother I observed in the grocery store had her 3-year-old son, Ricky, reach for items on the shelf and put them in the cart. / H: Ricky was a good boy and followed his mother doing nothing else.</p>

Table 8: NLI pairs that are correctly classified by only one model. 'Correct Model' is the only model that correctly classifies the NLI pairs to the right.

ment to verify as input. The output is used as the embedding of the document. Next, a pointer net takes these embeddings of all documents and sequentially outputs pointers to relevant documents.

In our adaptation, we use RoBERTa in place of BERT. More importantly, we use search snippets in place of page titles to take advantage of the relevant content in each document provided by search engines. The ranker is still trained on the FEVER dataset, but since it does not include search snippets, we heuristically generate snippets by concatenating the title of each Wikipedia page with its sentence that is most similar to the statement⁹. This technique substantially improves document relevance prediction on the FEVER dataset by 7.4% F1-score points. For web documents, we use search snippets provided by Bing and Google.

The number of retrieved documents varies a lot depending on the search method; the Google API retrieves much fewer documents than Wikipedia and Bing in general. Since this imbalance makes it difficult to compare the utility of the different search methods, we make the number of candidate documents the same across the methods, by ranking documents from different search methods separately and then pruning low-ranked documents of Wikipedia and Bing. This process lets the three methods have the same average number of candidate documents per statement (~ 8).

Sentence Selection: DeSePtion considers all sentences of the ranked documents. However, web documents are substantially longer than Wikipedia articles in general, so it is computationally too expensive and introduces a lot of noise to process all sentences. Therefore, for each statement to verify, we reduce the number of candidate sentences by selecting the top 200 sentences (among all ranked documents) whose RoBERTa embeddings have the highest cosine similarity to the statement.

Relation Prediction: In this stage, we classify whether each candidate sentence is valid counterevidence to the statement to verify. Here, instead of DeSePtion, we simply use an NLI model as-is. The reason is that the main goal of DeSePtion is to predict the veracity of a statement, rather than whether each sentence supports or refutes the statement. Thus, it assumes that once a statement is found to be supported or refuted, considering more

⁹We combine all token embeddings in the last layer of RoBERTa and measure the cosine similarity between these vectors.

sentences results in the same prediction. This assumption is justified for the FEVER task, where a statement cannot be both supported and refuted. In real-world arguments, however, a statement can be both supported and refuted, and our goal is to find refuting sentences. We compute the probability score that each sentence contradicts the statement and rank the sentences by these scores. This simple approach has been found to be effective in information retrieval (Dai and Callan, 2019).

D Comparison between KENLI+E and LogBERT

In order to better understand how differently KENLI+E and LogBERT behave, we broke down the performance of the models based on how strongly a pair of statement and candidate sentence signals the four logical relations—textual contradiction, negative sentiment, obstructive causal relation, and refuting normative relation—that have high correlations with LogBERT’s decision of ‘contradiction’ (Jo et al., 2021). More specifically, when LogBERT takes a pair of statement and candidate sentence, it can compute the probabilities of these four logical relations using its pretrained classification layers. For each of these relations, we split input pairs into two groups: pairs with probability greater than 0.5 and the other. Our hypothesis is that LogBERT performs well on the pairs in the first group, because they have a strong signal associated with ‘contradiction’ that LogBERT learns during pretraining.

Table 9 shows the breakdown of model accuracy. As expected, LogBERT usually achieves the highest recall and F1-score among all the models for input pairs that have strong signals of the logical mechanisms ($P > 0.5$). This pattern is pronounced for negative sentiment and obstructive causal relation. In contrast, LogBERT’s F1-score drops substantially when such signals are missing ($P \leq 0.5$), whereas KENLI+E’s performance is more stable regardless of those signals. This result implies the overreliance of LogBERT on the four logical relations, which is helpful when such relations exist in input pairs but rather harmful otherwise.

E In-Depth Analyses of Evidence Retrieval

E.1 Utility of Search Methods

One difference between our system and prior work is that we retrieved web documents using Bing

		CMV						Kialo					
		$P > 0.5$			$P \leq 0.5$			$P > 0.5$			$P \leq 0.5$		
		Prec	Rec1	F1	Prec	Rec1	F1	Prec	Rec1	F1	Prec	Rec1	F1
Textual Contradict	RoBERTa	50.2	60.9	55.0	45.6	68.6	54.8	60.4	53.8	56.9	53.9	64.0	58.5
	KENLI+C	51.2	65.8	57.6	44.7	63.7	52.5	61.6	59.2	60.3	54.7	69.0	61.1
	KENLI+E	51.5	69.7	59.2	44.9	74.3	56.0	60.8	62.1	61.4	53.2	72.0	61.2
	LogBERT	51.5	83.4	63.7	50.7	21.4	30.1	60.8	86.8	71.5	53.0	20.2	29.2
Negative Sentiment	RoBERTa	51.3	57.0	54.0	47.3	66.4	55.3	61.6	44.6	51.7	56.4	65.7	60.7
	KENLI+C	51.0	60.3	55.3	48.0	67.1	56.0	62.8	49.7	55.5	57.4	71.0	63.4
	KENLI+E	51.8	63.1	56.9	47.8	74.9	58.4	61.5	52.3	56.5	56.4	74.2	64.1
	LogBERT	49.4	96.5	65.3	53.3	46.9	49.9	60.8	96.9	74.7	58.8	44.7	50.8
Obstructive Causal	RoBERTa	52.5	60.5	56.2	47.8	64.0	54.7	61.5	53.6	57.3	57.3	57.7	57.5
	KENLI+C	52.5	66.9	58.8	48.3	64.8	55.4	61.9	59.8	60.9	58.5	62.7	60.5
	KENLI+E	54.5	69.8	61.2	48.2	71.5	57.6	61.9	64.0	63.0	57.2	65.4	61.0
	LogBERT	52.6	100.0	68.9	51.1	56.5	53.7	58.4	100.0	73.8	60.5	59.4	59.9
Refuting Normative	RoBERTa	46.6	62.9	53.5	49.6	64.0	55.9	59.3	61.6	60.4	56.9	53.8	55.3
	KENLI+C	47.8	64.1	54.8	49.5	65.7	56.5	60.8	66.3	63.4	57.8	59.4	58.6
	KENLI+E	48.5	70.5	57.4	49.1	71.9	58.4	59.0	69.6	63.9	57.2	62.1	59.6
	LogBERT	53.7	71.0	61.1	49.5	55.4	52.3	61.7	67.0	64.3	58.8	65.7	62.1

Table 9: Breakdown of model accuracy by the strength of logical relations in input pairs. The first column indicates a logical mechanism that is associated with LogBERT’s decision of ‘contradiction’. ‘ $P > 0.5$ (≤ 0.5)’ indicates input pairs whose probability of the logical mechanism is greater than (less than or equal to) 0.5. For each mechanism, bold numbers indicate the highest score for each metric.

	CMV			Kialo		
	P	R	F	P	R	F
Wikipedia	42.4	64.5	51.1	55.1	60.8	57.8
Bing	53.1	66.2	58.9	59.5	63.5	61.4
Google	47.0	64.8	54.5	59.9	62.6	61.2

Table 10: Accuracy of counterevidence retrieval by search methods.

and Google, whereas no prior work did that to our knowledge. Hence, comparing candidate sentences from Wikipedia, Bing, and Google will shed light on the usefulness of the search engines and inform future system designs. Table 10 shows candidate sentences retrieved from Bing and Google generally achieve higher F1-scores than those from Wikipedia. While Wikipedia provides comparably good recall, its precision is substantially lower than the other methods. This suggests that Wikipedia is a great source of a vast amount of relevant information, but the other search methods and more diverse types of documents should not be ignored if one needs more precise and nuanced counterevidence.

E.2 Utility of Document Types

One question we want to answer is: what types of documents are useful sources of counterevidence to argumentative statements? Prior work focuses mostly on Wikipedia articles (Thorne et al., 2018;

Hua and Wang, 2018), debates (Orbach et al., 2020; Wachsmuth et al., 2018), and occasionally news articles (Hua et al., 2019). In contrast, our candidate sentences come from more diverse types of documents, such as academic papers and government reports. To analyze the utility of different document types, we first annotated each candidate sentence with 13 different types using MTurk (Table 6). See Appendix A.4 for annotation details.

First of all, Figure 8 shows the distribution of document types for valid counterevidence. A lot of counterevidence exists in knowledge archives (27–37%), followed by mainstream news (8–13%), magazines about social issues (7–12%), personal blogs (5–10%), and research journals (6–8%). This suggests the benefit of using broader types of documents in counterevidence and fact verification than conventionally used Wikipedia and debates.

Table 9 summarizes the F1-score of counterevidence retrieval by document types (averaged across all models). For both CMV and Kialo, financial magazines and Q&A platforms are useful document types providing high F1-scores. For CMV, magazines about culture and research journals are beneficial, while in Kialo, general-domain magazines and forums are useful types. As we also observed in the earlier analysis of search methods, Wikipedia, which is conventionally used in fact verification, and mainstream news are relatively less

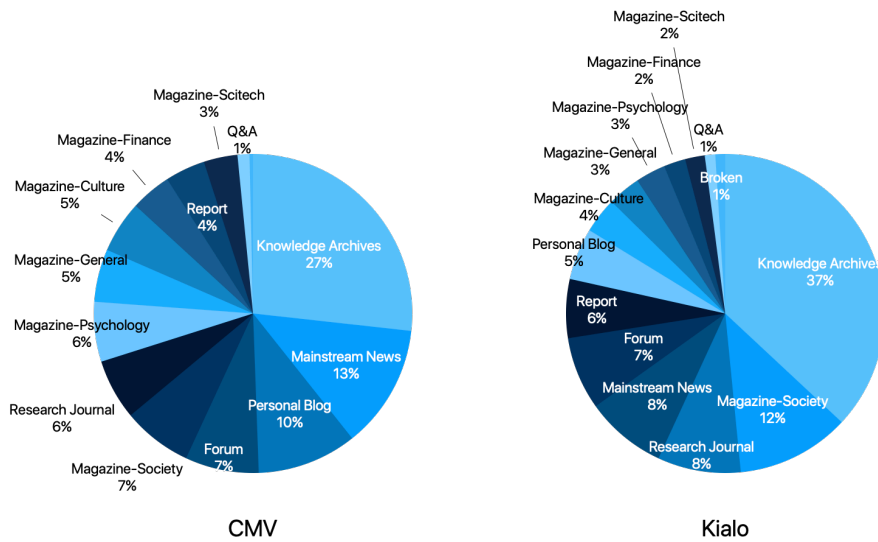


Figure 8: Distribution of document types for valid counter-evidence.

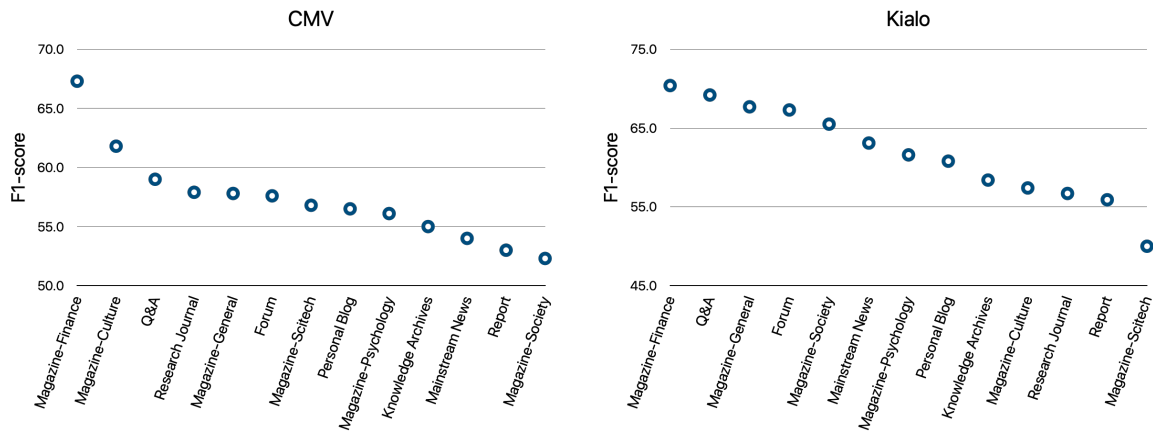


Figure 9: F1-scores of counter-evidence retrieval by evidence document types.

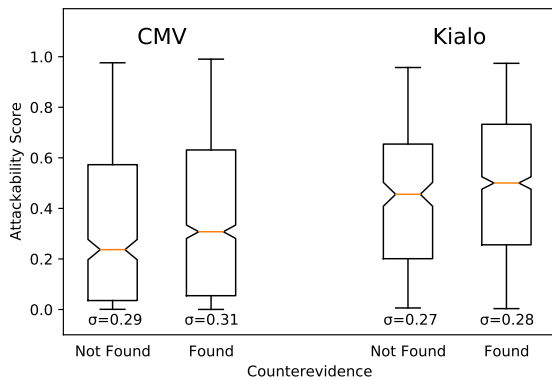


Figure 10: Attackability.

reliable. So are reports that contain a lot of detailed information.

E.3 Attackability

Our system was originally designed in consideration of the scenario where we counter an argument

by first detecting attackable sentences and then finding proper counter-evidence to them. Detecting attackable sentences in arguments has recently been studied for CMV based on persuasion outcomes (Jo et al., 2020). Here, we test if this method can help us find statements for which counter-evidence exists.

We assume that statements in our dataset are attackable if they have at least one candidate sentence that is valid counter-evidence. Figure 10 shows the distribution of the attackability scores of statements for which counter-evidence was found (Found) and statements for which no counter-evidence was found (Not Found). As expected, the attackability scores of statements that have counter-evidence are higher than the other statements for both CMV ($p = 0.001$) and Kialo ($p = 0.003$ by the Wilcoxon rank-sum test).

The attackability score of each statement also

has a small positive correlation with the number of candidate sentences that are valid counterevidence, resulting in Kendall’s τ of 0.057 (CMV, $p = 0.002$) and 0.085 (Kialo, $p = 0.006$). These results suggest that synergy can be made by integrating our system with attackability detection to build a complete counterargument generation system. We leave this direction to future work.

F Reproducibility Checklist

F.1 Knowledge-Enhanced NLI

- **A clear description of the mathematical setting, algorithm, and/or model:** Explained in the main text.
- **Submission of a zip file containing source code, with specification of all dependencies, including external libraries, or a link to such resources:** This information will be made available at a git repository upon publication.
- **Description of computing infrastructure used:** Intel(R) Xeon(R) Gold 5215 CPU @ 2.50GHz (20 CPUs), 128GiB System memory, Quadro RTX 8000 (4 GPUs).
- **The average runtime for each model or algorithm (training + inference):** KENLI+C: 311.9 mins / KENLI+E: 562.8 mins
- **Number of parameters in each model:** KENLI+C: 135,676,421 / KENLI+E: 174,330,206.
- **Corresponding validation F1-score (across all datasets) for each reported test result:** KENLI+C: 74.5 / KENLI+E: 75.0.
- **Explanation of evaluation metrics used, with links to code:** Explained in the main text.
- **The exact number of training and evaluation runs:** 5 runs.
- **Bounds for each hyperparameter:** None.
- **Number of hyperparameter search trials:** No hyperparameter search.
- **Hyperparameter configurations for best-performing models:** Explained in the main text.
- **Relevant details such as languages, and number of examples and label distributions:** Explained in the main text.
- **Details of train/validation/test splits:** Explained in the main text.
- **Explanation of any data that were excluded, and all pre-processing steps:** Explained in the main text.
- **A zip file containing data or link to a downloadable version of the data:** Example-NLI

will be made available in a git repository upon publication.

- **For new data collected, a complete description of the data collection process, such as instructions to annotators and methods for quality control:** Explained in the main text.

F.2 Evidence Retrieval

- **A clear description of the mathematical setting, algorithm, and/or model:** Explained in the main text.
- **Submission of a zip file containing source code, with specification of all dependencies, including external libraries, or a link to such resources:** This information will be made available at a git repository upon publication.
- **Description of computing infrastructure used:** Intel(R) Xeon(R) Gold 5215 CPU @ 2.50GHz (20 CPUs), 128GiB System memory, Quadro RTX 8000 (4 GPUs).
- **The average runtime for each model:**
 - Document retrieval: 36.1 mins.
 - Relation classification: KENLI+C: 44.3 mins / KENLI+E: 71.2 mins.
- **Number of parameters in each model:**
 - Document retrieval: 125,853,255.
 - Relation classification: KENLI+C: 135,676,421 / KENLI+E: 174,330,206.
- **Corresponding validation performance for each reported test result:** No validation.
- **Explanation of evaluation metrics used, with links to code:** Explained in the main text.
- **The exact number of evaluation runs:** 1 run.
- **Bounds for each hyperparameter:** None.
- **Number of hyperparameter search trials:** No hyperparameter search.
- **Relevant details such as languages, and number of examples and label distributions:** Explained in the main text.
- **Details of train/validation/test splits:** Explained in the main text.
- **Explanation of any data that were excluded, and all pre-processing steps:** Explained in the main text.
- **A zip file containing data or link to a downloadable version of the data:** The annotated data will be made available at a git repository upon publication.
- **For new data collected, a complete description of the data collection process, such as instructions to annotators and methods for**

quality control: Explained in the main text.