

SPARSECL: Sparse Contrastive Learning for Contradiction Retrieval

Haike Xu^{1*} Zongyu Lin^{2*} Yizhou Sun² Kai-Wei Chang² Piotr Indyk¹

¹MIT ²University of California Los Angeles


<https://sparsecl.github.io/>

Abstract

Contradiction retrieval refers to identifying and extracting documents that explicitly disagree with or refute the content of a query, which is important to many downstream applications like fact checking and data cleaning. To retrieve contradiction argument to the query from large document corpora, existing methods such as similarity search and crossencoder models exhibit significant limitations. The former struggles to capture the essence of contradiction due to its inherent nature of favoring similarity, while the latter suffers from computational inefficiency, especially when the size of corpora is large. To address these challenges, we introduce a novel approach: SPARSECL that leverages specially trained sentence embeddings designed to preserve subtle, contradictory nuances between sentences. Our method utilizes a combined metric of cosine similarity and a sparsity function to efficiently identify and retrieve documents that contradict a given query. This approach dramatically enhances the speed of contradiction detection by reducing the need for exhaustive document comparisons to simple vector calculations. We validate our model using the Arguana dataset, a benchmark dataset specifically geared towards contradiction retrieval, as well as synthetic contradictions generated from the MSMARCO and HotpotQA datasets using GPT-4. Our experiments demonstrate the efficacy of our approach not only in contradiction retrieval with more than 30% accuracy improvements on MSMARCO and HotpotQA across different model architectures but also in applications such as cleaning corrupted corpora to restore high-quality QA retrieval. This paper outlines a promising direction for improving the accuracy and efficiency of contradiction retrieval in large-scale text corpora.

1 Introduction

Figure 1: Performance gains in NDCG@10 score across different sentence embedding models and datasets, showcasing the effectiveness and robustness of our SPARSECL compared with standard contrastive learning (CL)



We study the problem of contradiction retrieval. Given a large document corpus and a query passage, the goal is to retrieve document(s) in the corpus that contradict the query, assuming they exist. This

*Equal Contribution




Figure 2: Comparison of our SPARSECL with Cross-Encoder and Contrastive-Learning based Bi-Encoder for contradiction retrieval.

problem has a large number of applications, including counter-argument detection [23] and fact verification [21]. The standard approaches to retrieving contradictions are two-fold. One is to use a bi-encoder [26, 8, 9] that maps each document to a feature space such that two contradicting documents are mapped close to each other (e.g., according to the cosine metric) and use nearest neighbor search algorithms. The second approach is to train a cross-encoder model [26] that determines whether two documents contradict each other, and apply it to each document or passage in the corpus.

Unfortunately, both methods suffer from limitations. The first approach is inherently incapable of representing the “contradiction relation” between the documents, due to the fact that the cosine metric is transitive: if A is similar to B , and B is similar to C , then A is also similar to C . As an example, consider an original sentence and its paraphrase in Table 5. Both of them contradict the sentence in the third column but they are not contradicting each other. The second approach, which uses a cross-encoder model, can capture the contradiction between sentences to some extent, but it is much more computationally expensive. Our experiment in Appendix F shows that compared with standard vector computation, running a cross-encoder is at least 200 times slower.

In this paper, we propose to overcome these limitations by introducing SPARSECL for efficient contradiction retrieval using sparse-aware sentence embeddings. The key idea behind our approach is to train a sentence embedding model to preserve *sparsity of differences* between the contradicted sentence embeddings. When answering a query, we calculate a score between the query and each document in the corpus, based on *both* the cosine similarity and the sparsity of the difference between their embeddings, and retrieve the ones with the highest scores. Our specific measure of sparsity is defined by the Hoyer measure of sparsity [5], which uses the scaled ratio of the ℓ_1 norm and the ℓ_2 norm of a vector as a proxy of the number of non-zero entries in the vector. Unlike the cosine metric, the Hoyer measure is not transitive (please see Appendix C for an example), which avoids the limitations of the former. At the same time this method is much more efficient than a cross-encoder, as both the cosine metric and the Hoyer measure are easy to compute given the embeddings. The Hoyer sparsity histogram of our trained embeddings, is displayed in Figure 3

We first evaluate our method on the counter-argument detection dataset Arguana [23], which to the best of our knowledge, is the only publicly available dataset suitable for testing contradiction retrieval. In addition, we generate two synthetic data sets, where contradictions for documents in MSMARCO [13] and HotpotQA [29] datasets are synthetically generated using GPT-4 [1]. Our experiments demonstrate the efficacy of our approach in contradiction retrieval, as seen in Table 1. We also apply our method to corrupted corpus cleaning problem, where the goal is to filter out contradictory sentences in a corrupted corpus and preserve good QA retrieval accuracy.

To summarize. our contributions can be divided into three folds:

- We introduce a novel contradiction retrieval method that employs specially trained sentence embeddings combined with a metric that includes both cosine similarity and the Hoyer

measure of sparsity. This approach effectively captures the essence of contradiction while being computationally efficient.

- Our method demonstrates superior performance on both real and synthetic datasets, achieving significant improvements in contradiction retrieval metrics compared to existing methods. This underscores the effectiveness of our embedding and scoring approach.
- We apply our contradiction retrieval method to the problem of corpus cleaning, showcasing its utility in removing contradictions from corrupted datasets to maintain high-quality QA retrieval. This application highlights the practical benefits of our approach in real-world scenarios.




Figure 3: Histograms for the Hoyer sparsity of different pairs of sentence embedding differences on HotpotQA test set. The left figure is the histogram produced by a standard sentence embedding model (“bge-base-en-v1.5”), where the median Hoyer sparsity values for random pairs, paraphrases, and contradictions are 0.212, 0.211, 0.211. The right figure is the histogram produced by our sentence embedding model fine-tuned from “bge-base-en-v1.5” using our SPARSECL method, where the median Hoyer sparsity values for random pairs, paraphrases, and contradictions are 0.212, 0.281, 0.632.

2 Related Work

Counter Argument Retrieval A direct application of our contradiction retrieval task in “counter-argument retrieval”. Since the curation of Arguana dataset by [23], there has been a few previous work on retrieving the best counter-argument for a given argument [14, 18]. In terms of methods, [23] uses a weighted sum of different word and embedding similarities and [18] designs a “Bipolar-encoder” and a classification head. We believe that our method relying only on cosine similarity and sparsity is simpler than theirs and produces better results in the experiment. In addition, some analyses in the counter-argument retrieval papers are specific to the “debate” setting, e.g. they rely on topic, stance, premise/conclusion, and some other inherent structures in debates for help, which may prevent their methods from being generalized to broader scenarios.

Fact verification and LLM hallucination Addressing the hallucination problem in Large Language Models has been a subject of many research efforts in recent years. According to the three types of different hallucinations in [31], here we only focus on those so called “Fact-Conflicting Hallucination” where the outputs of LLM contradict real world knowledge. The most straightforward way to mitigate this hallucination issue is to assume an external groundtruth knowledge source and augment LLM’s outputs with an information retrieval system. There have been a few works on this line showing the success of this method [15, 11]. This practice is very similar to “Fact-Verification” [21, 16] where the task is to judge whether a claim is true or false based on a given knowledge base.

However, as pointed out by [31], in the era of LLM, the external knowledge base can encompass the whole internet. It is impossible to assume that all the information there are perfectly correct and there may exist conflicting information within the database. In the context of our paper, instead of using a groundtruth database to check an external claim, our goal is to check the internal contradictions between different documents in an unknown corpus.

Learning augmented LLM and retrieval corpus attack Augmenting large language models with retrieval has been shown to be useful for many purposes. Recently, there have been a few

works [32, 34] studying the vulnerability of retrieval system from adversarial attack. In specific, they show that adding a few corrupted data to the corpus will significantly drop the retrieval accuracy. This phenomenon bring our attention to the necessity of checking the factuality of the knowledge database. Note that the type of corrupted documents considered by their papers are different from ours. While they consider the injection of adversarially generated documents, we consider the existence of contradicted documents as a natural part of the corpus. Also their purpose is to show the effect of adversarial attack, while we provide a defense method for a certain kind of corrupted database.

3 Method

Problem Formulation We consider the contradiction retrieval problem: given a passage corpus $C = \{p_1, p_2, \dots, p_n\}$ and a query passage q , retrieve the “best” passage p^* that contradicts q . We assume that several similar passages supporting q might exist in the corpus C .

Embedding based method Judging whether two passages contradict each other is a standard Natural Language Inference task and can be easily tackled by many off-the-shelf language models [22, 28], . However, to retrieve the best candidate from the corpus, we have to iterate the whole corpus, or at least send the candidates retrieved by similarity search to the language model to determine if they constitute contradiction. This is time consuming, given that there are potentially many similar passages in the corpus. Therefore, in our paper, we mainly focus on those methods that only rely on their passage embeddings. Specifically, we want to design a simple scoring function F that given the embeddings of two passages, outputs a score between $[0, 1]$, indicating the likelihood that they are contradicting each other.

Sparsity Enhanced Embeddings Following the idea from counter-argument retrieval papers [23], such a score function should be a combination of similarity and dissimilarity functions. Observe that a dissimilarity function is basically a negation of a similarity function, so the authors of [23] design several different similarity functions and set the scoring function to maximize one of them and minimize another. Here, instead of enumerating different similarity functions, we consider another notion: the “sparsity” of their embedding differences. The basic intuition is as follows. Suppose that all sentences are represented as vectors in a “semantic” basis, where each coordinate represents one clearly identifiable semantic meaning. Then a contradiction between two passages should manifest itself as a difference in a few coordinates, while other coordinates should be quite close to each other. The issue, however, is that we do not know how to construct the appropriate basis, and the sparsity is defined with respect to a fixed coordinate system. Nevertheless, following this intuition, we fine-tune sentence embedding models using contrastive learning, by rewarding the sparsity of the difference vectors between embeddings of contradicting passages. Please see Figure 3 for the Hoyer sparsity histogram of our trained embeddings.

SPARSECL We use contrastive learning to fine-tune any pretrained sentence embedding model to generate the desired sparsity-enhanced embeddings. The choice of positive and negative examples are exactly the reverse of the choice we make when the training sets are Natural Language Inference datasets. The positive example for a passage is its contradiction passage in the training set. The hard negative example for a passage is its similar passage in the training set. There are also other random in-batch passages as soft negative examples. The sparsity function we choose here is Hoyer sparsity function from [5]. Let h_1 and h_2 be two sentence embeddings and their embeddings have dimension d . We define

$$\text{Hoyer}(h_1, h_2) = \left(\sqrt{d} - \frac{\|h_1 - h_2\|_1}{\|h_1 - h_2\|_2} \right) / \left(\sqrt{d} - 1 \right).$$

This is a transformed version of the ratio of the l_1 to the l_2 norm, with output normalized to $[0, 1]$.

Finally, for each training tuple (x_i, x_i^+, x_i^-) with their embeddings (h_i, h_i^+, h_i^-) , batch size N , and temperature τ , its loss function is defined as

$$l_i = -\log \frac{e^{\text{Hoyer}(h_i, h_i^+)/\tau}}{\sum_{j=1}^N \left(e^{\text{Hoyer}(h_i, h_j^+)/\tau} + e^{\text{Hoyer}(h_i, h_j^-)/\tau} \right)}.$$

Scoring function for contradiction retrieval For the score function for contradiction retrieval, we use a weighted sum of the standard cosine similarity and our sparsity function. Note that the cosine similarity is provided separately by any off-the-shelf sentence embedding model in a zeroshot manner. It can also be fine-tuned. Let $E()$ be the standard sentence embedding model and $E_s()$ be our sparsity-enhanced sentence embedding model trained by SPARSECL. Then the final score function for contradiction retrieval is

$$F(p_1, p_2) = \cos(E(p_1), E(p_2)) + \alpha \cdot \text{Hoyer}(E_s(p_1), E_s(p_2)).$$

where α is a scalar tuned using the validation set. Note that the criterion for contradiction is usually case-dependent, so it is necessary that we reserve a parameter to adapt to different notions of contradiction. To get the answer passages, we calculate the score function for all passages and report the top 10 of them².

4 Experiments

We test our contradiction retrieval method on a counterargument retrieval task Arguana [23] and two synthetic datasets adapted from HotpotQA[29] and MSMARCO[13]. Then, we apply our contradiction retrieval task to a new experimental setting: retrieval corpus cleaning. Finally, we perform ablation studies to explain the functionality of each component of our method.

4.1 Counter-argument Retrieval

Dataset Arguana is a dataset curated in [23], where the author provide a corpus of 6753 argument-counterargument pairs, taken from 1069 debates with 15 themes on idebate.org. For each debate, the arguments are further divided into two opposing stances (pro and con). For each stance, there are paired arguments and counter-arguments. The dataset is split into the training set (60% of the data), the validation set (20%), and the test set (20%). This ensures that data from each individual debate is included in only one set and that debates from every theme are represented in every set. The task goal is: given an argument, retrieve its best counter-argument.

Training We use Arguana’s training set to fine-tune our sparsity-enhanced sentence embedding model via SPARSECL. To construct our training data, for each argument and counter-argument pair (x_i, x_i^c) in the Arguana’s training set, we set x_i^c to be the positive example of x_i . We select all the other arguments and counter-arguments from the same debate and stance as x_i ’s hard negatives. We fine-tune three pretrained sentence embedding models of different sizes (“UAE-Large-V1” [8], “GTE-large-en-v1.5” [9], and “bge-base-en-v1.5” [26]) selected from the top 10 models on HuggingFace’s MTEB Retrieval Arguana leaderboard. Please refer to Table 7 for our training parameters.

Baselines We mainly compare our method to the similarity-based method. Since Arguana is one of the datasets in the MTEB Retrieval benchmark, directly searching for the similar passages in the corpus can already produce quite good test results. We report the performance of the state-of-the-art pretrained sentence embedding models including “GTE-large-en-v1.5”, “UAE-Large-V1”, “bge-base-en-v1.5”, “SFR-Embedding-Mistral” [10], “voyage-lite-02-instruct” when used to directly retrieve the most similar argument to each query (Zeroshot). For a fair comparison, we also report the results of fine-tuning these models using standard contrastive learning (CL) on the same dataset used for SPARSECL.

Test The Arguana test set consists of 1401 query arguments and counter-argument pairs. Following the standard test setting, we search for an answer of a query within the whole corpus (training set + validation set + test set) and report NDCG@10 scores. The α parameter we used in the score function varies across different datasets and models. We select α based on the best NDCG@10 score on the validation set. Please refer to Table 8 in Appendix E for our specific α choices and parameter searching details. When we directly use a model to provide cosine similarity scores in a zeroshot

²In the actual implementation, for time efficiency, we first use FAISS [4] to retrieve the top K candidates with cosine similarity and then rerank them using our cosine + sparsity score function. We set a very large K (e.g. $K = 1000$) so that empirically this is almost equivalent to searching for the maximal cosine + sparsity score in the whole corpus

Model	Method	Arguana	MSMARCO	HotpotQA
SFR-Mistral	Zeroshot (Cosine)	0.672	0.605	0.595
VOYAGE	Zeroshot (Cosine)	0.703	N/A	N/A
BGE	Zeroshot (Cosine)	0.658	0.600	0.595
	Zeroshot (Cosine) + SPARSECL(Hoyer)	0.704	0.909	0.967
BGE	CL (Cosine)	0.687	0.564	0.562
	CL (Cosine) + SPARSECL(Hoyer)	0.722	0.883	0.965
UAE	Zeroshot (Cosine)	0.683	0.597	0.587
	Zeroshot (Cosine) + SPARSECL(Hoyer)	0.732	0.902	0.955
UAE	CL (Cosine)	0.712	0.442	0.541
	CL (Cosine) + SPARSECL(Hoyer)	0.750	0.869	0.943
GTE	Zeroshot (Cosine)	0.725	0.603	0.597
	Zeroshot (Cosine) + SPARSECL(Hoyer)	0.797	0.953	0.977
GTE	CL (Cosine)	0.778	0.651	0.597
	CL (Cosine) + SPARSECL(Hoyer)	0.813	0.952	0.979

Table 1: Results for different models and methods on the contradiction retrieval task. Experiments are run on the Arguana dataset [23] and modified MSMARCO[13] and HotpotQA[29] datasets. We report NDCG@10 score here, the higher the better. “UAE” stands for “UAE-Large-V1”, “BGE” stands for “bge-base-en-v1.5”, “GTE” stands for “gte-large-en-v1.5”, “SFR-Mistral” stands for “SFR-Embedding-Mistral”, “VOYAGE” stands for “voyage-lite-02-instruct”. The “Method” column denotes the score function used to retrieve contradictions. We consider two score functions: cosine similarity and cosine similarity plus Hoyer sparsity. “Zeroshot” denotes the direct testing of the model without any fine-tuning. “CL” denotes fine-tuning using standard contrastive learning. “SPARSECL” denotes fine-tuning using Hoyer sparsity contrastive learning (our method).

manner, we use its default pooler (“cls”) for that model. When we use a fine-tuned model (via either CL or SPARSECL) to provide either cosine similarity scores or sparsity scores, we use the “avg” pooler.

Results The detailed results are presented in Table 1. Across all models—“GTE-large-en-v1.5”, “UAE-Large-V1”, and “bge-base-en-v1.5”—an average improvement of 4.5% in counter-argument retrieval were observed when incorporating our SPARSECL to either Zeroshot or CL. Furthermore, our combination of GTE with CL (Cosine) + SPARSECL (Hoyer) method achieving NDCG@10 score 0.813 is also the highest score on Arguana as far as we know (we also report the Zeroshot results of the other top 3 sentence embedding models “SFR-Mistral” and “VOYAGE” on MTEB for reference.)

This pattern of enhancement was consistently observed regardless of whether the embedding models were fine-tuned or not. Notably, standard cosine similarity fine-tuning alone also contributed to performance gains. For instance, fine-tuned GTE models showed an increase from 0.725 to 0.778 on the Arguana dataset using standard cosine similarity alone. This suggests that the Arguana dataset inherently favors scenarios where the counterargument is the most similar passage to the query, which may amplify the benefits of fine-tuning.

These findings highlight the robustness of our approach, particularly when traditional similarity metrics are augmented with sparsity measures to capture subtle nuances in contradiction. Further insights can be gleaned from our ablation study detailed in Section 4.4, where we analyze the impact of similar non-contradictory passages within the corpus.

4.2 Contradiction retrieval on synthetic datasets

The task of “contradiction retrieval” generalizes beyond the argument and counter-argument relationship in the debate area, e.g. passages with conflicting factual information should also be considered as “contradictions”. To test our method’s validity for these more general forms of contradictions, we construct two synthetic datasets to test our method’s performance.

Data set construction Given a QA retrieval dataset, e.g. MSMARCO [13], for each answer passage x_i of a query q_i , we use Large Language Models (specifically, GPT-4 [1]) to generate 3 synthetic answers paraphrasing x_i or contradicting x_i . Let the generated paraphrases be $\{x_{i1}^+, x_{i2}^+, x_{i3}^+\}$ and the generated contradictions be $\{x_{i1}^-, x_{i2}^-, x_{i3}^-\}$. We then delete x_i from the corpus and add the set of generated passages $\{x_{i1}^+, x_{i2}^+, x_{i3}^+, x_{i1}^-, x_{i2}^-, x_{i3}^-\}$ to the corpus. In the test phrase, the queries are $\{x_{i1}^+, x_{i2}^+, x_{i3}^+\}$, each of which has the same answers $\{x_{i1}^-, x_{i2}^-, x_{i3}^-\}$. We generate the paraphrases and contradictions for the validation set, test set, and a randomly sampled 10000 documents from the training set.

The reason why we only keep the generated text but not the original one is that all the GPT-4 generated passages are easily distinguishable from the human written ones, which makes language models vulnerable to shortcuts. We performed a random inspection to check that all the paraphrases and contradictions are generated as intended. Please refer to Table 5 to see two examples of the generated paraphrases and contradictions. We report the prompts and the temperature parameter we use to generate these data in Appendix A.

Training To prepare the training data for contrastive learning, for each paraphrase and contradiction set $\{x_{i1}^+, x_{i2}^+, x_{i3}^+, x_{i1}^-, x_{i2}^-, x_{i3}^-\}$ generated from the same original passage, we form 9 pieces of training data $(x_{ia}^+, x_{ib}^-, x_{ic}^+)$ for 9 different combinations of paraphrases, contradictions, and a randomly selected hard negative from the remaining two paraphrases. We then perform SPARSECL to fine-tune a sparsity-enhanced embedding.

Baseline We are not aware of any accurate methods for retrieving contradictions that only rely on sentence embeddings. Therefore, the only baseline we provide is a standard contrastive learning with cosine similarity (CL), using the same training data (contradictions as positive examples and paraphrases as negative examples) that we use for our SPARSECL.

Test Similar to the testing strategy for Arguana, we define our corpus to consist of all generated text (training set + validation set + test set). We query the paraphrases $\{x_{i1}^+, x_{i2}^+, x_{i3}^+\}$ of the original passage x_i and set the groundtruth answers to be the generated contradictions $\{x_{i1}^-, x_{i2}^-, x_{i3}^-\}$. We select the α parameter with the maximal NDCG@10 score on the validation set and report the NDCG@10 score obtained by applying that α to the test set.

The results are reported in Table 1. For both MSMARCO and HotpotQA data sets, incorporating our SPARSECL method achieves over 30 percentage points gain compared with the pure cosine-similarity-based method. The large improvement is due to the existence of paraphrases in the corpus, that are strong confounders for the pure similarity-based methods. We also observe that fine-tuning using standard contrastive learning with cosine similarity (CL) yields performance gains for Arguana but not for MSMARCO and HotpotQA. Our explanation is that, for MSMARCO and HotpotQA, the generated paraphrases are more similar to the query than the contradictions. Therefore fine-tuning with the standard cosine similarity is unlikely to work.

4.3 Retrieval Corpus Cleaning

As an application of contradiction retrieval, we test how well our method can be used to find inconsistencies within a corpus and clean the corpus for future training or QA retrieval. We first inject corrupted data contradicting existing documents into the corpus, and measure the retrieval accuracy degradation for retrieved answers. Then, we use our contradiction retrieval method to filter out corrupted data and measure the retrieval accuracy again.

Data Similarly to the data generation in Section 4.2, we construct a new corpus containing LLM-generated paraphrases and contradictions based on MSMARCO and HotpotQA data sets. We start with an original corpus C and its subset S . We then generate paraphrases and contradictions for S as in Section 4.2.

For HotpotQA, S contains all answer documents for the test set, 10000 answer documents sampled from the training set, and 1000 answer documents sampled from the development set. For MSMARCO, S contains all answer documents for the dev set, and 11000 answer documents sampled from the training set.

We then curate 3 different versions of the corpus based on the original corpus C and the subset S .

- The initial corpus C^+ : For each original answer document x in S , we remove x from C and instead add 3 LLM-generated paraphrases $\{x_1^+, x_2^+, x_3^+\}$ to C . The result forms the *initial* corpus C^+ .
- The corrupted corpus C^- : For each original answer document x in S , we generate 3 contradictions $\{x_1^-, x_2^-, x_3^-\}$ and add them to C^+ to get the *corrupted* corpus C^- .
- The cleaned corpus C^\natural : We apply our data cleaning procedure to the corrupted corpus C^- , obtaining the *cleaned* dataset C^\natural .

Test We test the retrieval accuracy (NDCG@10) and the corruption ratio (Recall@10) for answering the original queries in the test set. The goal of our experiment is to show how retrieval algorithms behave on these three constructed corpora C^+ , C^- , and C^\natural .

Data Cleaning Our sparsity-based method can only identify contradictions within the data set, but we do not know which element in a contradiction pair is correct. To perform data cleaning, we make the assumption that for each original passage $x \in S$, we are given one of its paraphrases as the groundtruth. Then, our task is reduced to searching for passages contradicting a given ground truth document and filtering them out.

Method We use the GTE-large-en-v1.5 model without fine-tuning to provide the cosine similarity score for this data cleaning experiment. We use the model from our contradiction retrieval experiment in section 4.2 trained on MSMARCO and HotpotQA to provide the sparsity score. The α parameter is also identical to the one used in section 4.2. For each ground truth document, we filter out the top 3 scored documents from the corpus.

Note that the optimal choice of α for contradiction retrieval may not be the optimal choice for data cleaning because of different test objectives. We apply the same α only for simplicity, as our goal is to demonstrate the validity of applying our method to the data cleaning problem.

Datasets	Original	Corrupted		Cleaned	
	Acc	Acc	Corrupt	Acc	Corrupt
HotpotQA	0.676	0.567	0.443	0.652	0.020
MSMARCO	0.435	0.381	0.413	0.414	0.040

Table 2: Experimental results for the impact of corrupted data on QA retrieval and contradiction retrieval for filtration. “Acc” represents the retrieval accuracy measured by the NDCG@10 score and “Corrupt” represents the fraction of returned passages that are corrupted, as measured by Recall@10.

Table 2 shows the results. We observe that the retrieval accuracy on the corrupted corpus drops significantly, as the generated contradictions cause the embedding model to retrieve them as query answers. The corruption ratio measures the average fraction of the top-10 retrieved documents that correspond to the generated contradicting passages. This performance is above 40% for both datasets. After performing our corpus cleaning procedure, which searches for the passages contradicting the given ground truth documents and removes the top-3 for each of them, we can recover more than 60% of the performance loss due to corruption and at the same time reduce the corruption ratio to less than 5%.

4.4 Ablation Studies

We perform the following two ablation studies to further understand sparsity-based retrieval method.

Arguana retrieval results analysis In the standard Arguana dataset, even though the task is to retrieve the counter-argument for the query, the retrieval based solely on similarity still gives reasonable results. This means that counter-arguments are also the most similar arguments to the query, which makes the data set an imperfect test bed for testing contradiction retrieval.

To further compare our sparsity-based method and the pure similarity-based method, we augment Arguana by adding arguments’ paraphrases to the corpus. Specifically, for any argument x and

its counter-argument x^- in the original corpus C , we use GPT-4 to generate three paraphrases $\{x_1, x_2, x_3\}$ of x . We then form three new corpora with an increasing number of paraphrases added to the corpus: C_1 contains all x_1 and x^- , C_2 contains all x_1, x_2 , and x^- , and C_3 contains all x_1, x_2, x_3 , and x^- .

In the testing phase, we query the counter-arguments for one of x 's paraphrases, the answer of which should still be x^- . We observe how the performance varies when the corpora we retrieve from are C_1, C_2, C_3 .

Models	Methods	C_1	C_2	C_3
BGE	Zeroshot (Cosine)	0.561	0.355	0.267
	Zeroshot (Cosine) + SPARSECL(Hoyer)	0.682	0.679	0.675
BGE	CL (Cosine)	0.471	0.303	0.228
	CL (Cosine) + SPARSECL(Hoyer)	0.619	0.618	0.615

Table 3: Counter-argument retrieval results on the augmented Arguana dataset with different numbers of similar arguments in the corpus. C_x denotes testing counter-argument retrieval on the corpus with x existing paraphrases (including itself) of the query argument.

We present our experimental results in Table 3. As the number of paraphrases in corpus increases from 1 to 3, the performance of the similarity-based method drops significantly. Thus it is reasonable to deduce that, as the number of similar arguments in the corpus increases further, the NDCG@10 scores for similarity-based methods will converge to 0. On the other hand, the performance of our sparsity-based method is stable with respect to the number of paraphrases in the corpus.

Different sparsity functions Our intuition in Section 3 does not give clear guidelines on which sparsity function to use in our SPARSECL. Thus, we also experiment with different choices of sparsity functions, selected from [5]. Specifically, we consider two other sparsity functions (l_2/l_1 and κ_4), which are scale invariant and differentiable (see Table III in [5]). Note that both of these two sparsity functions have ranges $[0, 1]$, and higher values of those functions correspond to sparser vectors.

$$\frac{l_2}{l_1} = \frac{\|h_1 - h_2\|_2}{\|h_1 - h_2\|_1} \quad \kappa_4 = \frac{\|h_1 - h_2\|_4^4}{\|h_1 - h_2\|_2^2}.$$

Model	Method	l_2/l_1	κ_4	Hoyer	Cosine (baseline)
BGE	Zeroshot (Cosine) + SPARSECL	0.675	0.684	0.704	0.657
BGE	CL (Cosine) + SPARSECL	0.702	0.707	0.722	0.687

Table 4: NDCG@10 scores for Arguana using SPARSECL with different sparsity functions. We also report two baselines that use only the cosine similarity (zeroshot and contrastive learning).

As per Table 4, compared to the cosine similarity method, the combination of the cosine similarity score with the sparsity score trained by SPARSECL, yields higher NDCG@10 scores for each sparsity function. However, Hoyer sparsity yields the highest accuracy. We believe that simple sparsity functions have a more benign optimization landscape and thus are easier for language models to optimize.

5 Conclusion

In this work, we introduced a novel approach to contradiction retrieval that leverages sparsity-enhanced sentence embeddings combined with cosine similarity to efficiently identify contradictions in large document corpora. This method addresses the limitations of the traditional similarity search as well as computational inefficiencies of the cross-encoder models, proving its effectiveness on benchmark datasets like Arguana and on synthetic contradictions retrieval from MSMARCO and HotpotQA. An interesting question that we do not study here but deserves further study is whether it is possible to perform nearest neighbor search w.r.t. the Hoyer sparsity measure in *sublinear time*, as opposed to performing linear scan that we use in this work.

References

- [1] Josh Achiam, Steven Adler, Sandhini Agarwal, Lama Ahmad, Ilge Akkaya, Florencia Leoni Aleman, Diogo Almeida, Janko Altenschmidt, Sam Altman, Shyamal Anadkat, et al. GPT-4 technical report. *arXiv preprint arXiv:2303.08774*, 2023.
- [2] Akari Asai, Timo Schick, Patrick Lewis, Xilun Chen, Gautier Izacard, Sebastian Riedel, Hanneh Hajishirzi, and Wen-tau Yih. Task-aware retrieval with instructions. In Anna Rogers, Jordan Boyd-Graber, and Naoaki Okazaki, editors, *Findings of the Association for Computational Linguistics: ACL 2023*, pages 3650–3675, Toronto, Canada, July 2023. Association for Computational Linguistics.
- [3] Limeng Cui and Dongwon Lee. Coaid: Covid-19 healthcare misinformation dataset, 2020.
- [4] Matthijs Douze, Alexandre Guzhva, Chengqi Deng, Jeff Johnson, Gergely Szilvassy, Pierre-Emmanuel Mazaré, Maria Lomeli, Lucas Hosseini, and Hervé Jégou. The faiss library. 2024.
- [5] Niall Hurley and Scott Rickard. Comparing measures of sparsity. *IEEE Transactions on Information Theory*, 55(10):4723–4741, 2009.
- [6] Rohan Jha, Charles Lovering, and Ellie Pavlick. Does data augmentation improve generalization in nlp? *arXiv preprint arXiv:2004.15012*, 2020.
- [7] Philippe Laban, Tobias Schnabel, Paul N. Bennett, and Marti A. Hearst. SummaC: Re-visiting NLI-based models for inconsistency detection in summarization. *Transactions of the Association for Computational Linguistics*, 10:163–177, 2022.
- [8] Xianming Li and Jing Li. Angle-optimized text embeddings. *arXiv preprint arXiv:2309.12871*, 2023.
- [9] Zehan Li, Xin Zhang, Yanzhao Zhang, Dingkun Long, Pengjun Xie, and Meishan Zhang. Towards general text embeddings with multi-stage contrastive learning. *arXiv preprint arXiv:2308.03281*, 2023.
- [10] Rui Meng, Ye Liu, Shafiq Rayhan Joty, Caiming Xiong, Yingbo Zhou, and Semih Yavuz. Sfr-embedding-mistral:enhance text retrieval with transfer learning. Salesforce AI Research Blog, 2024.
- [11] Grégoire Mialon, Roberto Dessì, Maria Lomeli, Christoforos Nalmpantis, Ram Pasunuru, Roberta Raileanu, Baptiste Rozière, Timo Schick, Jane Dwivedi-Yu, Asli Celikyilmaz, Edouard Grave, Yann LeCun, and Thomas Scialom. Augmented language models: a survey, 2023.
- [12] Niklas Muennighoff, Nouamane Tazi, Loic Magne, and Nils Reimers. MTEB: Massive text embedding benchmark. In Andreas Vlachos and Isabelle Augenstein, editors, *Proceedings of the 17th Conference of the European Chapter of the Association for Computational Linguistics*, pages 2014–2037, Dubrovnik, Croatia, May 2023. Association for Computational Linguistics.
- [13] Tri Nguyen, Mir Rosenberg, Xia Song, Jianfeng Gao, Saurabh Tiwary, Rangan Majumder, and Li Deng. MS MARCO: A human generated machine reading comprehension dataset. *CoRR*, abs/1611.09268, 2016.
- [14] Matan Orbach, Yonatan Bilu, Assaf Toledo, Dan Lahav, Michal Jacovi, Ranit Aharonov, and Noam Slonim. Out of the echo chamber: Detecting countering debate speeches. *arXiv preprint arXiv:2005.01157*, 2020.
- [15] Ruiyang Ren, Yuhao Wang, Yingqi Qu, Wayne Xin Zhao, Jing Liu, Hao Tian, Hua Wu, Ji-Rong Wen, and Haifeng Wang. Investigating the factual knowledge boundary of large language models with retrieval augmentation, 2023.
- [16] Tal Schuster, Adam Fisch, and Regina Barzilay. Get your vitamin C! robust fact verification with contrastive evidence. In Kristina Toutanova, Anna Rumshisky, Luke Zettlemoyer, Dilek Hakkani-Tur, Iz Beltagy, Steven Bethard, Ryan Cotterell, Tamoy Chakraborty, and Yichao Zhou, editors, *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 624–643, Online, June 2021. Association for Computational Linguistics.

[17] Gautam Kishore Shahi and Durgesh Nandini. *FakeCovid- A Multilingual Cross-domain Fact Check News Dataset for COVID-19*. ICWSM, Jun 2020.

[18] Hongguang Shi, Shuirong Cao, and Cam-Tu Nguyen. Revisiting the role of similarity and dissimilarity in best counter argument retrieval. *arXiv preprint arXiv:2304.08807*, 2023.

[19] Amit Singhal et al. Modern information retrieval: A brief overview. *IEEE Data Eng. Bull.*, 24(4):35–43, 2001.

[20] Nandan Thakur, Nils Reimers, Andreas Rücklé, Abhishek Srivastava, and Iryna Gurevych. BEIR: A heterogeneous benchmark for zero-shot evaluation of information retrieval models. In *Thirty-fifth Conference on Neural Information Processing Systems Datasets and Benchmarks Track (Round 2)*, 2021.

[21] James Thorne, Andreas Vlachos, Christos Christodoulopoulos, and Arpit Mittal. FEVER: a large-scale dataset for fact extraction and VERification. In Marilyn Walker, Heng Ji, and Amanda Stent, editors, *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, June 2018. Association for Computational Linguistics.

[22] Hugo Touvron, Louis Martin, Kevin Stone, Peter Albert, Amjad Almahairi, Yasmine Babaei, Nikolay Bashlykov, Soumya Batra, Prajwala Bhargava, Shruti Bhosale, et al. Llama 2: Open foundation and fine-tuned chat models. *arXiv preprint arXiv:2307.09288*, 2023.

[23] Henning Wachsmuth, Shahbaz Syed, and Benno Stein. 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, 2018.

[24] Xiaoyue Wang, Jianyou Wang, Weili Cao, Kaicheng Wang, Ramamohan Paturi, and Leon Bergen. Birco: A benchmark of information retrieval tasks with complex objectives, 2024.

[25] Zhihao Wang, Zongyu Lin, Peiqi Liu, Guidong ZHeng, Junjie Wen, Xianxin Chen, Yujun Chen, and Zhilin Yang. Learning to detect noisy labels using model-based features. *arXiv preprint arXiv:2212.13767*, 2022.

[26] Shitao Xiao, Zheng Liu, Peitian Zhang, and Niklas Muennighoff. C-pack: Packaged resources to advance general chinese embedding, 2023.

[27] Qizhe Xie, Minh-Thang Luong, Eduard Hovy, and Quoc V Le. Self-training with noisy student improves imagenet classification. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 10687–10698, 2020.

[28] Haike Xu, Zongyu Lin, Jing Zhou, Yanan Zheng, and Zhilin Yang. A universal discriminator for zero-shot generalization. *arXiv preprint arXiv:2211.08099*, 2022.

[29] Zhilin Yang, Peng Qi, Saizheng Zhang, Yoshua Bengio, William Cohen, Ruslan Salakhutdinov, and Christopher D. Manning. HotpotQA: A dataset for diverse, explainable multi-hop question answering. In Ellen Riloff, David Chiang, Julia Hockenmaier, and Jun’ichi Tsujii, editors, *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 2369–2380, Brussels, Belgium, October–November 2018. Association for Computational Linguistics.

[30] Xinyu Zhang, Nandan Thakur, Odunayo Ogundepo, Ehsan Kamalloo, David Alfonso-Hermelo, Xiaoguang Li, Qun Liu, Mehdi Rezagholizadeh, and Jimmy Lin. MIRACL: A Multilingual Retrieval Dataset Covering 18 Diverse Languages. *Transactions of the Association for Computational Linguistics*, 11:1114–1131, 09 2023.

[31] Yue Zhang, Yafu Li, Leyang Cui, Deng Cai, Lemao Liu, Tingchen Fu, Xinting Huang, Enbo Zhao, Yu Zhang, Yulong Chen, Longyue Wang, Anh Tuan Luu, Wei Bi, Freda Shi, and Shuming Shi. Siren’s song in the ai ocean: A survey on hallucination in large language models, 2023.

[32] Zexuan Zhong, Ziqing Huang, Alexander Wettig, and Danqi Chen. Poisoning retrieval corpora by injecting adversarial passages. *arXiv preprint arXiv:2310.19156*, 2023.

- [33] Jing Zhou, Zongyu Lin, Yanan Zheng, Jian Li, and Zhilin Yang. Not all tasks are born equal: Understanding zero-shot generalization. In *The Eleventh International Conference on Learning Representations*, 2022.
- [34] Wei Zou, Runpeng Geng, Binghui Wang, and Jinyuan Jia. Poisonedrag: Knowledge poisoning attacks to retrieval-augmented generation of large language models, 2024.

A Data generation details for MSMARCO and HotpotQA experiments in Section 4.2

We use “gpt-4-turbo” to generate paraphrases and contradictions for our experiment in Section 4.2. The prompts we use are in Table 6. We set $temperature = 1$ and $n = 3$ (to generate 3 outputs). Please see Table 5 for some examples of generated paraphrases and contradictions.

Datasets	Orginal	Paraphrase	Contradiction
MSMARCO	In addition to the high financial value of higher education, higher education also makes individuals much more intelligent than what they would be with just a high school education...	Beyond its significant monetary worth , higher education substantially enhances a person's intelligence compared to merely completing high school...	Besides the low financial significance of higher education, higher education often renders individuals no more intelligent than they would be with just a high school education...
HotpotQA	Ice hockey is a contact team sport played on ice , usually in a rink , in which two teams of skaters use their sticks to shoot a vulcanized rubber puck into their opponent's net to score points...	Ice hockey is a contact sport where two teams compete on an ice surface , typically in a rink , using sticks to hit a vulcanized rubber puck into the opposing team's net to earn points...	Ice hockey is a non-contact team sport played on grass , often in an open field , where two teams of players use their feet to kick a soft leather ball into their opponent's goal to score points...

Table 5: Examples of passages from MSMARCO and HotpotQA datasets, with their generated paraphrases, and generated contradictions. Highlighted key-words represent exact matchings or contradictions

Task	Prompt
Generating paraphrases	Paraphrase the given paragraph keeping its original meaning. Do not add information that is not present in the original paragraph. Your response should be as indistinguishable to the original paragraph as possible in terms of length, language style, and format. Begin your answer directly without any introductory words.
Generating contradictions	Rewrite the given paragraph to contradict the original content. Ensure the revised paragraph changes the factuality of the original. Your response should be as indistinguishable to the original paragraph as possible in terms of length, language style, and format. Begin your answer directly without any introductory words.

Table 6: Prompts used to generate paraphrases and contradictions for MSMARCO and HotpotQA documents.

B Additional related work

Complex retrieval tasks Information retrieval is a well-studied area [19] and there have been many benchmarks for testing retrieval performance such as BEIR [20], MTEB [12], and MIRACL [30]. However, most of the datasets, through varying in some degrees, focus only on “retrieving the most similar document”. People have noted that there exist some more complex retrieval tasks (e.g. Arguana [23] retrieves counter-arguments that refute a query argument), and build retrieval benchmark focusing on complex retrieval goals, e.g. BIRCO [24] and BERRI [2].

To retrieve according to different instructions, [2] trains TART, a multi-task retrieval system with task instructions attached as prompts in front of the query content. However, when answering queries,

they are still searching for the most similar sentence embedding, though the prompt is different for different tasks. As far as we know, our paper studies the first non-similarity-based search problem.

Data inconsistency and misinformation detection Data inconsistency, refers to the factually incorrectness in the content, might come from different sources, including their natural existence in the corpus [17, 3], data augmentations [6, 33], and pseudo labeling [27, 25], which might lead to negative influence if serving as training dataset. There have been a few datasets on detecting the factually wrong information. For example, [7] detects whether a given summary is consistent with the input document, [17, 3] detects whether a given COVID-19 related news is true or false. Most of these datasets lie in a specific domain and require external knowledge to judge the correctness of each piece of data. On the contrary, the “data inconsistency” notion we consider in our paper doesn’t depend on any external knowledge, but is a relationship between different pieces of data in the same corpus. The goal of our method is to find such “contradiction pairs” in corpus efficiently, but not to judge which one is consistent with the real world knowledge.

C An example demonstrating the non-transitivity of Hoyer sparsity

Here, we provide a simple example to demonstrate that using Hoyer sparsity to measure “contradiction” can bypass the challenging scenario for similarity metrics where “A contradicts C, B contradicts C, but A doesn’t contradict B”.

Specifically, we will construct three embeddings with dimension d for sentences A, B, and C, such that $\text{Hoyer}(A, C) > 1 - O\left(\frac{1}{\sqrt{d}}\right)$, $\text{Hoyer}(B, C) > 1 - O\left(\frac{1}{\sqrt{d}}\right)$, and $\text{Hoyer}(A, B) < O\left(\frac{1}{\sqrt{d}}\right)$, exactly matching their contradiction relationship.

We construct the following d dimensional embeddings where $\epsilon < \frac{1}{d}$ can be any parameter.

$$\begin{aligned} A &= (1, 0, 0, \dots, 0) \\ B &= (1, 0, \epsilon, \dots, \epsilon) \\ C &= (0, 1, 0, \dots, 0) \end{aligned}$$

Then, we calculate their l_1 over l_2 ratios:

$$\begin{aligned} \frac{\|A - B\|_1}{\|A - B\|_2} &= \sqrt{d - 2} \\ \frac{\|A - C\|_1}{\|A - C\|_2} &= \sqrt{2} \\ \frac{\|B - C\|_1}{\|B - C\|_2} &= \frac{2 + (d - 2)\epsilon}{\sqrt{2 + (d - 2)\epsilon^2}} < \frac{3}{\sqrt{2}} \end{aligned}$$

Applying their l_1 over l_2 ratio bounds to the Hoyer sparsity formula will give us the desired relationship.

D Compute Resources

Most of our experiments are not so computationally extensive, which can be run by one single A6000 GPU. We run our major experiments on A6000 and A100 GPUs.

E Hyper-parameters for Training and Inference

Here we present the training details (Table 7) for our experiments on Arguana and synthetic HotpotQA and MSMARCO. We report the α parameters tuned on the validation set in Table 8. We search the α parameters from the range $[0, 10]$ by first dividing the range into 10 intervals, calculating the NDCG@10 score on the validation set for each interval’s midpoint, and then diving into that interval for a finer search. We stop when the interval range is smaller than 0.01

Models	Model Size	Backbone	CL		PARSECL		temp	bz
			ep	lr	ep	lr		
GTE-large-en-v1.5	434M	BERT + RoPE + GLU	1	1e-5	3	2e-5	0.01	64
UAE-Large-V1	335M	BERT	1	2e-5	3	2e-5	0.05	64
bge-base-en-v1.5	109M	BERT	1	2e-5	3	2e-5	0.02	64

Table 7: Training parameters for Arguana. We set max sequence length to be 512 for Arguana dataset and 256 for HotpotQA and MSMARCO datasets.

Models	Methods	Arguana	MSMARCO	HotpotQA
GTE	Zeroshot (Cosine) + SPARSECL(Hoyer)	0.88	2.65	2.36
	CL (Cosine) + SPARSECL(Hoyer)	0.20	0.35	5.44
UAE	Zeroshot (Cosine) + SPARSECL(Hoyer)	0.15	1.00	1.06
	CL (Cosine) + SPARSECL(Hoyer)	0.12	1.01	1.22
BGE	Zeroshot (Cosine) + SPARSECL(Hoyer)	0.18	2.19	4.82
	CL (Cosine) + SPARSECL(Hoyer)	0.12	2.53	3.72

Table 8: α choices for different methods and datasets

F Efficiency test of cross-encoder and vector calculation

To further compare the efficiency of cross-encoders and Hoyer sparsity calculations, we perform the following experiments:

- We choose “bge-reranker-base” and “bge-reranker-large” to be our cross-encoders. We use them to calculate the similarity between one query from Arguana’s test set and 100 documents from Arguana’s corpus. We report the average running time of this method for 100 queries.
- We choose “bge-base-en-v1.5” and “bge-large-en-v1.5” to be our bi-encoders. Suppose we have preprocessed all the sentence embeddings. We use it to calculate the Hoyer sparsity between one query embedding from Arguana’s test set and 100 document embeddings from Arguana’s corpus. We report the average running time of this method for 100 queries.

Please see Table 9 for the running time of different methods. We can see that the calculation of Hoyer sparsity is at least 200 times faster than running a cross-encoder.

Cross-encoder	Model size	Time
bge-reranker-base	278M	0.8832s
bge-reranker-large	560M	1.6022s
Bi-encoder	Embedding dimension	Time
bge-base-en-v1.5	768	0.0029s
bge-large-en-v1.5	1024	0.0036s

Table 9: Average running time for calculating the score functions between one Arguana query and 100 Arguana documents