

# A Counterfactual Explanation Framework for Retrieval Models

Anonymous ACL submission

## Abstract

Explainability has become a crucial concern in today's world, aiming to enhance transparency in machine learning and deep learning models. Information retrieval is no exception to this trend. In existing literature on explainability of information retrieval, the emphasis has predominantly been on illustrating the concept of relevance concerning a retrieval model. The questions addressed include why a document is relevant to a query, why one document exhibits higher relevance than another, or why a specific set of documents is deemed relevant for a query. However, limited attention has been given to understanding why a particular document is not favored (e.g., not within top-K) with respect to a query and a retrieval model. In an effort to address this gap, our work focuses on the question of what terms need to be added within a document to improve its ranking. This, in turn, answers the question of which words in the document played a role in not being favored by a retrieval model for a particular query. We use a counterfactual framework to solve the above-mentioned research problem. To the best of our knowledge, we mark the first attempt to tackle this specific counterfactual problem (i.e. examining the absence of which words can affect the ranking of a document). Our experiments show the effectiveness of our proposed approach in predicting counterfactuals for both statistical (e.g. BM25) and deep-learning-based models (e.g. DRMM, DSSM, ColBERT, MonoT5). The code implementation of our proposed approach is available in <https://anonymous.4open.science/r/CfIR-v2>.

## 1 Introduction

The requirement of transparency of Artificial Intelligence (AI) models has made explainability crucial, and this applies to Information Retrieval (IR) models as well (Anand et al., 2022). The target audience plays a significant role in achieving

explainability for an IR model, as the units of explanation or questions may differ based on the end user. For instance, a healthcare specialist, who is a domain expert but not necessarily an IR specialist, might want to understand the reasons behind a ranked suggestion produced by a retrieval model in terms of words used (Singh and Anand, 2019). On the other hand, an IR practitioner may be more interested in understanding whether different IR axioms are followed by a retrieval model or not (Bondarenko et al., 2022).

This study focuses on the perspective of IR practitioners. To be more specific, we introduce a counterfactual framework designed for retrieval models, catering to the needs of IR practitioners. Existing literature in explainable IR (ExIR) addressed questions like why a particular document is relevant with respect to a query (Singh and Anand, 2019), between a pair of documents why one document is more relevant to the query (Penha et al., 2022) compared to the other and why a list of documents relevant to a query (Lyu and Anand, 2023). Broadly speaking, all the above-mentioned questions mainly focus on explaining the relevance of a document or a list of documents from different perspectives.

However, there is limited attention to explain the question like the absence of which words renders a document unfavorable to a retrieval model (i.e. not within top-K) remains unexplored. The above-mentioned explanation can give an idea to an IR practitioner about how to modify a retrieval model. For example, if it is observed that a retrieval model (e.g. especially neural IR models (Rekabsaz and Schedl, 2020)) does not favor documents because of not having certain gender specific words then the setting of the retrieval model needs to be debiased.

With the motivation described above, the fundamental research question which we address in this research work is described as follows.

- **RQ1:** What are the terms that should be added to a document which can push the document to a higher rank with respect to a particular retrieval model?

We would like to note that we have framed **RQ1** as a counterfactual setup in our research scope. Similar to existing research in counterfactual explanations in AI (Kanamori et al., 2021; Van Loov-  
eren and Klaise, 2021), we also attempt to change the output of model with the provided explanations (i.e. change the rank of a document in IR models). Our experimental results show that on an average in 70% cases the solution provided by the counterfactual setup improves the ranking of a document with respect to a query and a ranking model.

**Our Contributions** The main contributions of this paper are as follows.

- Propose a model-agnostic novel counterfactual framework for retrieval models.
- Estimated a set of terms that can explain why a document is not within top-K with respect to a query and a retrieval model.
- Provide a comprehensive analysis with existing state-of-the-art IR models.

The rest of the paper is organized as follows. Section 2 describes Related work. Section 3 describes the counterfactual framework used in our work, Section 4 describes the experimental setup and Section 5 discuss about results and ablation study. Section 6 concludes with this paper.

## 2 Related Work

### 2.1 Counterfactual Explanations

The xAI field gained significant momentum with the development of the Local Interpretable Model-agnostic Explanations (LIME) method (Ribeiro et al., 2016), which offers a way to explain any classification model. While models like LIME explain why a model predicts a particular output, counterfactual explainers address the question of what changes in input features would be needed to alter the output. Counterfactual xAI was first brought into the limelight in early 2010s with seminal work of Pearl (2018). The study in Karimi et al. (2020) provided a practical framework named Model-Agnostic Counterfactual Ex-

planations (MACE) for any model. Later series of models (Kanamori et al., 2021; Van Loov-  
eren and Klaise, 2021; Parmentier and Vidal, 2021; Carreira-Perpiñán and Hada, 2021; Pawel-  
czyk et al., 2022; Hamman et al., 2023) were proposed for counterfactual explanation based on different optimization frameworks. In our re-  
search scope, we use Counterfactual Explanation framework proposed in (Mothilal et al., 2020) (ex-  
plained in detail in Section 3).

### 2.2 Explainability in IR

**Pointwise Explanations** shows the important features responsible for the relevance score pre-  
dicted by a retrieval model for a query-document pair. Popular techniques include locally approx-  
imating the relevance scores predicted by the re-  
trieval model using a regression model (Singh and Anand, 2019).

**Pairwise Explanations** predict why a particular document was favored by a ranking model com-  
pared to others. The work in (Xu et al., 2024) proposed a counterfactual explanation method to compare the ranking of a pair of documents with respect to a particular query.

**Listwise Explanations** focus on explaining the key features for a ranked list of documents and a query. Listwise explanations (Yu et al., 2022; Lyu and Anand, 2023) aim to capture a more global perspective compared to pointwise and pairwise explanations. The study in (Lyu and Anand, 2023) proposed an approach which combines the output of different explainers to capture the different aspects of relevance. The study in (Yu et al., 2022) trained a transformer model to generate explanation terms for a query and a ranked list of documents.

**Generative Explanations** (Singh and Anand, 2020; Lyu and Anand, 2023) generally leverage natural language processing to create new text content, like summaries or justifications, that directly address the user’s query and information needs. Model-agnostic approaches (Singh and Anand, 2020) have been proposed to interpret the intent of the query as understood by a black box ranker.

From the above mentioned category of explanations in IR, we focus on pointwise explanation in our research scope. In pointwise explanation, rather than explaining what are the words which are relevant in a document for a particular query we address the research question what are

the words which are required to improve the ranking of the document with respect to a query.

**Search Engine Optimization** techniques (Egri and Bayrak, 2014; Erdmann et al., 2022) generally uses different features like commercial cost, links to optimize the performance of the search engine. A major difference of the work in (Egri and Bayrak, 2014; Erdmann et al., 2022) with our work is we only consider the words present in a document as a feature. Our objective is to improve the ranking of a particular document concerning a specific query and a retrieval model rather than improving the ranking of a document concerning any query belonging to a particular topic.

### 3 Counterfactual Framework for Information Retrieval (CFIR)

**Problem Statement** Let  $d$  represents a target document that does not appear in the top- $K$  retrieved results of a query  $q$  and retrieval model  $M$ . The objective in CFIR is to identify a set of terms  $w_i$  which, when added to  $d$ , improve its ranking with respect to  $q$  and model  $M$ .

The above mentioned setup for CFIR is formally defined in Equation 1 where *CFIR*, employs a counterfactual document generator  $c_k(f_{\{M,q\}}, d)$  which takes as input a classifier  $f_{M,q}$  and the document  $d$  to construct an counterfactual document  $d'$  such that  $d'$  is likely to get a higher rank (within top- $K$ ) than  $d$  for model  $M$  and query  $q$ . The objective of  $f_{\{M,q\}} : R^{|V|} \rightarrow \{0, 1\}$  (where  $V$  is the vocabulary, described in detail in Section 3.1) is to predict given a query  $q$  and a retrieval model  $M$  if a particular document  $d$  will be within top- $K$  or not. The counterfactual explanation is defined as the set of words present in  $d'$  but not in  $d$  (i.e. output of Equation 1).

$$\begin{aligned} CFIR(q, M, d) &= c_k(f_{\{M,q\}}, d) - d \\ &= d' - d = \cup_{i=1}^m \{w_i\} \end{aligned} \quad (1)$$

#### 3.1 Building Classifier ( $f_{\{M,q\}}$ )

Similar to existing xAI (Ribeiro et al., 2016) approaches, the classifier  $f_{\{M,q\}}$  in our research scope essentially locally approximates the behavior of a retrieval model  $M$ , for a query  $q$  and a subset of documents retrieved for the query  $q$ . However, in contrast to the regression model in (Ribeiro et al., 2016), we build a binary classification model to predict whether a document  $d$  will

be ranked within the top- $K$  results or not for a specific query  $q$  and retrieval model  $M$ .

For each document  $d$  for which we want to generate counterfactuals, we train a separate classifier. In the classifier setup, the top- $K$  documents for a query  $q$  and retrieval model  $M$  represent class 1 and any other document not belonging to this class represents class 0. Theoretically speaking, if a corpus had  $N$  number of documents, then there will be  $N - K$  documents which should have class label 0 and  $N - K$  is a very large number in general which can cause class imbalance issue. To avoid this issue, for the class 0, for each document  $d$  for which we want to generate a counterfactual, we choose a set of closest neighbors in the set of  $N - K$  documents and the size of the neighborhood should be similar to  $K$ .  $K$  serves as a pre-defined threshold, typically set to values such as 10, 20, or 30. For  $f_{\{M,q\}}$ , each document  $d$  is represented as a word term frequency based feature vector, denoted as  $d_{vec}$ .

Formally, **Feature Vector for Classifier**  $f_{\{M,q\}}$  is represented as  $d_{vec} = \{t.f_1^d, t.f_2^d \dots, t.f_{|V|}^d\}$  where  $t.f_i^d$  represents the term frequency of the word  $w_i$  in  $d$ . Using all the words from all the documents retrieved for a query to construct the vocabulary set  $V$  can pose challenges. Consequently, we take the union of the most significant  $n$  words from each document  $d$  using a function named  $Imp(d)$  (explained in detail in Section 4) to construct  $V$ .  $V = \cup_{i=1}^K \{\cup_{j=1, w_j \in Imp(d_i)}^n w_j\}$ . Appendix D depicts a step-by-step algorithm to construct the feature vector for the classifier and Figure 6 in Appendix D shows one sample feature vector for the classifier.

**Counterfactual Document Generator**  $c_k(f_{M,q}, d)$  in Equation 1 follows an approach similar to that of Mothilal et al. (2020). Specifically,  $c_k(f_{M,q}, d)$  generates  $k$  candidate counterfactuals  $c_1^{maxIter}, c_2^{maxIter}, \dots, c_k^{maxIter}$  (where  $maxIter$  is the maximum number of iterations upto which loss function is optimized) for each document  $d$ , from which we randomly select a single counterfactual ( $d'$  in Equation 1) that involves only insertion of new words without modifying or deleting existing ones in  $d$  (step 5 in Algorithm 1). We fix  $k$  to a sufficiently large constant in our experiments. Similar to (Mothilal et al., 2020), the objective of  $c_k(f_{M,q}, d)$  is to minimize three different criteria described as follows.




Figure 1: Schematic Diagram for Counterfactual Explanation Framework (CFIR)

- **Criteria 1:** Minimizing the distance between the desired outcome  $y'$  (within top- $K$ ) and the prediction of the classifier model  $f_{\{M,q\}}$  for a counterfactual example ( $c_i$ ).
- **Criteria 2:** Minimizing the distance between any generated counterfactual ( $c_i$ ) and the original document  $d$ . Broadly speaking, a counterfactual example closer to the original input should be more useful for a user.
- **Criteria 3:** Increasing diversity between generated counterfactuals.

Based on the above-mentioned criteria the loss function to generate  $c_1^{maxIter}, \dots, c_k^{maxIter}$  is described as follows.

$$\arg \min_{c_1, \dots, c_k} \left( \frac{1}{k} \sum_{i=1}^k \text{yloss}(f_{M,q}(c_i), y') + \frac{\lambda_1}{k} \sum_{i=1}^k \text{dist}(c_i, d) - \lambda_2 \text{div}(c_1, \dots, c_k) \right) \quad (2)$$

In Equation 2,  $\text{yloss}(\cdot)$  takes care of **Criteria 1**,  $\text{dist}(c_i, d)$  takes care of the **Criteria 2** and  $\text{div}$  takes care of the **Criteria 3** as discussed above.  $\lambda_1$  and  $\lambda_2$  in Equation 2 are hyperparameters that balance the contribution of second and third parts of loss function (i.e. controlling diversity and similarity). The detailed description of the computation of  $\text{yloss}$ ,  $\text{dist}$  and  $\text{div}$  function in Equation 2 is given in Equations 4, 5 and 6 respectively in Appendix F. The loss function in Equation 2 is optimized using the gradient descent method.

Algorithm 1 shows step by step execution of the counterfactual document generator  $c_k(f_{\{M,q\}}, d)$ . In Algorithm 1 we show how the counterfactual examples ( $c_1, \dots, c_k$ ) are randomly initialized. The generated counterfactual examples (i.e.  $c_i^{maxIter}$ s)

should change the prediction of classifier  $f_{\{M,q\}}$  from 0 to 1 (i.e. modified document should be within top  $K$ ). The set of words corresponding to the counterfactual explanation of  $d$  are the new words that have been added to  $d'_{vec}$  (i.e. feature vector representation of  $d'$  in Equation 1) compared to  $d_{vec}$ . Figure 1 shows the schematic diagram for counterfactual setup with the workflow between the different components (i.e. classifier and counterfactual document generator) within it.

#### Algorithm 1: CF Document Generator $c_k(f_{\{M,q\}}, d)$

**Input** : Classifier function:  $f_{\{M,q\}}$ , Feature Vector:  $d_{vec} = \{tf_1, tf_2, \dots, tf_{|V|}\}$ , Number of Counterfactuals:  $k$   
**Output** :  $\{d'_{vec} \in R^{|V|}\}$   
**Initialization:**  
for  $i \leftarrow 1$  to  $k$  do  
for  $j \leftarrow 1$  to  $|V|$  do  
 $c_{i,j}^0 = r \sim \text{Random}(\cdot)$   
/\*  $c_{i,j}^0$  is the  $j^{th}$  coordinate of  $c_i$  at  $0^{th}$  iteration \*/  
end for  
end for  
1 for  $t \leftarrow 0$  to  $maxIter$  do  
2 Compute the loss  $\frac{1}{k} \sum_{i=1}^k \text{yloss}(f_{M,q}(c_i^t), y) + \frac{\lambda_1}{k} \sum_{i=1}^k \text{dist}(c_i^t, d) - \lambda_2 \text{div}(c_1^t, \dots, c_k^t)$   
3 Update  $c_i^t$ s using gradient descent  
4 end for  
5 return  $d'_{vec}$ ,  $d'_{vec}$  is a  $|V|$  dimensional vector randomly chosen from the subset of  $c_i^{maxIter}$ s for which  $c_{i,j}^{maxIter} \geq tf_j^d \forall j = 1, \dots, |V|$

## 4 Experiment Setup

**Dataset** We use three ranking datasets for our experiments: MS MARCO passage dataset for passage ranking (Bajaj et al., 2016) and MS MARCO document ranking dataset for longer documents (Craswell et al., 2023) and TREC Robust (Voorhees, 2005) dataset. The MS MARCO passage and document ranking datasets contain



queries from Bing<sup>1</sup> and the queries of TREC Robust are manually chosen. For each dataset, we randomly selected 100 queries from the test set and chose 5 documents not ranked in the top 10 results for each query, resulting in a test set of 500 query-document pairs. The details of the dataset are given in Table 4 in Appendix C.

We use five different retrieval models BM25, DRMM Guo et al. (2016), DSSM (Huang et al., 2013), ColBERT Khattab and Zaharia (2020), MonoT5 (Nogueira et al., 2020) and Splade (Formal et al., 2021) in our experiments. The details of each retrieval model is given in Appendix A.

**Baselines** To the best of our knowledge, this is the first work which attempts to provide counterfactual explanations in IR. Consequently, there exists no baseline for our proposed approach. However we have used a query word and top-K word based intuitive baseline to compare with our proposed approach. In query word baseline (*QW*), we use query words not originally present in a document to enhance its ranking. For Top-K' (*Top - K'*) baseline we use the top  $k'$  words extracted from top 5 documents corresponding to a query as relevance set. Words appearing in the relevance set but not appearing in a document are added to the document to improve its ranking. For different retrieval models we have corresponding versions of *QW* and *Top - K'* baselines.

**Evaluation Metrics** There exists no standard evaluation framework for exIR approaches. The three different evaluation metrics in our experiment setup are described as follows.

**Fidelity (FD):** Existing xAI approaches in IR use Fidelity (Anand et al., 2022) as one of the metrics to evaluate the effectiveness of the proposed explainability approach. Intuitively speaking, Fidelity measures the correctness of the features obtained from a xAI approach. In the context of the CFIR setup described in this work, we define this fidelity score as the number of times the words predicted by the counterfactual algorithm could actually improve the rank of a document. Let  $n$  be total number of query document pairs in our test case and  $x$  be number of query document pairs for which the the rank of the document improved after adding the counterfactuals obtained from the optimization setup described in Equation 2. Then the

Fidelity score is mathematically defined with respect to a test dataset  $D$  and retrieval model  $M$  is defined as follows.

$$FD(D, M) = \frac{x}{n} * 100 \quad (3)$$

**Avg. New Words:** Here we compute the average number of new words added by the counterfactual approach for a set of query document pairs.

**Avg. Query Overlap:** Here we report on an average how many of the words suggested by the counterfactual algorithm come from the query words.

**Parameters and Implementation Details** The details of implementation about retrieval models are shown in Appendix B. We employed two popular classical machine learning methods, Logistic Regression (LR) and Random Forest (RF) for the classifier described in Section 3.1. For Logistic Regression, the learning rate was set to 0.001. For Random Forest, the number of estimators was set to 100. As described in Section 3.1, all the words present in a document are not used as input to the classifier. We use the top 10 ( $n' = 10$ ) most important words from a document. As described in Section 3.1, we explored three different ways to implement  $Imp(d)$  function a) TF-IDF weight based word extraction, b) BERT based keyword extraction (Grootendorst, 2020) and c) Similarity between the BERT representation of query and the document tokens. We found that BERT representation-based similarity computation worked the best for our approach. More details on the implementation of  $Imp(d)$  function are shown in Appendix I. The value of  $K'$  for *Top - K'* baseline was set to 5. More details on the parameter configuration are shown in Appendix G.

## 5 Results

Table 1 shows the performance of the counterfactual approach across different retrieval models (i.e. BM25, DRMM, DSSM, ColBERT, MonoT5 and Splade). We conducted experiments on MS MARCO passage and document ranking dataset and TREC Robust dataset to observe the effectiveness of our proposed explanation approach for different types of documents. Mainly four different observations can be made from Table 1. **Firstly**, It can be clearly observed that the CFIR model for each retrieval model has performed better compared to its corresponding query word or top-K'

<sup>1</sup><https://bing.com>

| Model Description    |            | MS MARCO Passage |                |                    | MS MARCO Document |                |                    | Trec Robust |                |                    |
|----------------------|------------|------------------|----------------|--------------------|-------------------|----------------|--------------------|-------------|----------------|--------------------|
| Retrieval Model      | Classifier | FD(%)            | Avg. New Words | Avg. Query Overlap | FD(%)             | Avg. New Words | Avg. Query Overlap | FD(%)       | Avg. New Words | Avg. Query Overlap |
| $QW_{BM25}$          | NA         | 50%              | 5.61           | 100%               | 48%               | 6.14           | 100%               | 56%         | 6.12           | 100%               |
| $Top - K'_{BM25}$    | NA         | 42%              | 11.28          | 100%               | 40%               | 9.61           | 100%               | 41%         | 12.34          | 100%               |
| $CFIR_{BM25}$        | RF         | 65%              | 10.64          | 66%                | 52%               | <b>16.81</b>   | 56%                | <b>64%</b>  | 11.12          | 57%                |
| $CFIR_{BM25}$        | LR         | <b>69%</b>       | <b>17.14</b>   | 58%                | <b>57%</b>        | 14.15          | 56%                | 58%         | <b>13.25</b>   | 56%                |
| $QW_{DRMM}$          | NA         | 48%              | 5.12           | 100%               | 47%               | 6.14           | 100%               | 49%         | 7.12           | 100%               |
| $Top - K'_{DRMM}$    | NA         | 42%              | <b>15.11</b>   | 100%               | 31%               | 14.12          | 100%               | 33%         | <b>16.12</b>   | 100%               |
| $CFIR_{DRMM}$        | RF         | <b>72%</b>       | 11.31          | 48%                | 56%               | 8.12           | 46%                | 62%         | 12.56          | 47%                |
| $CFIR_{DRMM}$        | LR         | 68%              | 12.37          | 62%                | <b>62%</b>        | <b>14.53</b>   | 45%                | <b>65%</b>  | 13.47          | 43%                |
| $QW_{DSSM}$          | NA         | 49%              | 5.32           | 100%               | 45%               | 6.64           | 100%               | 52%         | 7.12           | 100%               |
| $Top - K'_{DSSM}$    | NA         | 35%              | 12.51          | 100%               | 32%               | 12.62          | 100%               | 34%         | 13.14          | 100%               |
| $CFIR_{DSSM}$        | RF         | 57%              | 11.52          | 58%                | 46%               | 18.14          | 57%                | <b>59%</b>  | 12.46          | 100%               |
| $CFIR_{DSSM}$        | LR         | <b>62%</b>       | <b>15.78</b>   | 54%                | <b>53%</b>        | <b>18.52</b>   | 63%                | 58%         | <b>17.24</b>   | 64%                |
| $QW_{ColBERT}$       | NA         | 56%              | 4.78           | 100%               | 34%               | 5.64           | 100%               | 38%         | 6.14           | 100%               |
| $Top - K'_{ColBERT}$ | NA         | 48%              | <b>15.63</b>   | 100%               | 36%               | <b>13.42</b>   | 100%               | 38%         | 11.32          | 100%               |
| $CFIR_{ColBERT}$     | RF         | 72%              | 12.41          | 56%                | <b>72%</b>        | 11.05          | 49%                | 71%         | 10.35          | 52%                |
| $CFIR_{ColBERT}$     | LR         | <b>75%</b>       | 14.12          | 61%                | 71%               | 10.23          | 62%                | <b>74%</b>  | <b>16.45</b>   | 65%                |
| $QW_{MonoT5}$        | NA         | 52%              | 10.15          | 100%               | 54%               | 12.23          | 100%               | 63%         | 10.15          | 100%               |
| $Top - K'_{MonoT5}$  | NA         | 75%              | <b>14.11</b>   | 100%               | 68%               | 10.13          | 100%               | 75%         | 11.12          | 100%               |
| $CFIR_{MonoT5}$      | RF         | 80%              | 12.13          | 64%                | 72%               | 11.23          | 61%                | 73%         | 10.95          | 66%                |
| $CFIR_{MonoT5}$      | LR         | <b>82%</b>       | 13.15          | 65%                | <b>74%</b>        | <b>12.23</b>   | 63%                | <b>75%</b>  | <b>11.45</b>   | 68%                |
| $QW_{Splade}$        | NA         | 49%              | 10.15          | 100%               | 51%               | 11.51          | 100%               | 62%         | 11.11          | 100%               |
| $Top - K'_{Splade}$  | NA         | 71%              | <b>13.05</b>   | 100%               | 65%               | 9.23           | 100%               | 74%         | 12.22          | 100%               |
| $CFIR_{Splade}$      | RF         | 78%              | 11.23          | 62%                | 69%               | 12.11          | 60%                | 71%         | 9.81           | 65%                |
| $CFIR_{Splade}$      | LR         | <b>80%</b>       | 12.15          | 63%                | <b>71%</b>        | <b>14.11</b>   | 64%                | <b>73%</b>  | <b>10.55</b>   | 67%                |

Table 1: CFIR model Performance for BM25, DRMM, DSSM, ColBERT, MonoT5 and Splade in MSMARCO Passage and Document Collection and TREC Robust. The Best Performing Counterfactual Explanation Method for every retrieval model is boldfaced; the overall best performance across all rows is underlined. All the results reported in Table 1 are statistically significant with  $p < 0.05$ .

| Retrieval Model | Query Text                             | docId   | Explanation Terms  |
|-----------------|--|---------|--|
| DRMM            | What law repealed prohibition ?        | 3686955 | working, strict, Maine, 1929, law, resentment, New York City, Irish, immigrant, prohibition, repeal, fall, Portland, temperance, riot, visit |
| DSSM            | What is the role of lipid in the cell? | 6159679 | phospholipid, fluidity, storage, triglyceride, fatty receptor  |
| ColBERT         | what type of wave is electromagnetic?  | 5217641 | directly ,oscillations, medium, wave, properties, speed  |
| MonoT5          | what is a caret?                       | 6338711 | display, diamond, weight   |
| Splade          | which vitamins help heal bruises?      | 3465680 | minerals, body, eat, cut   |

Table 2: CFIR explanation terms for DRMM, DSSM, ColBERT, MonoT5 and Splade in MS MARCO passage.

words baseline in terms of Fidelity score(FD). The above-mentioned observation is consistent for both passages and long documents (i.e. in MS-MARCO passage, Document and TREC Robust). **Secondly**, it can be observed from Table 1 that mostly CFIR approach provided the highest number of new terms (terms not already present in the documents) as part of the explanation to improve ranking. Consequently, we can say the overall set of explanation terms are more diverse for CFIR approach compared to others. It can also be also observed from Table 1 that the Fidelity scores are generally better in the MS MARCO passages compared to MSMARCO document and TREC Robust dataset. One likely explanation for this phenomenon is that documents in MSMARCO document and TREC Robust are longer in length compared to passages. Consequently, it is easier for shorter documents to change the ranking compared to longer documents. **Thirdly**, another interesting observation from Table 1 is that the maximum query word overlap by our proposed ap-

proach is 68%. This implies that the counterfactual algorithm is suggesting new words that are not even present in a query. **Fourthly**, the performance of representation learning based retrieval models (i.e. ColBERT, MonoT5) are significantly better than the other models for Fidelity metric. One potential reason can be that, the counterfactual generator suggests words which are similar to the content of the document. Because of using better embedding representation (BERT (Devlin et al., 2019) and T5 compared to Word2Vec (Mikolov et al., 2013) in DRMM) these retrieval models give more priority to similar words than other retrieval models.

Prior work in information retrieval has explored adversarial attacks, where document content or embeddings are perturbed to manipulate rankings with malicious intent (Liu et al., 2023; Wu et al., 2022a). In contrast, the goal of counterfactual explanations is to provide interpretability for IR models by revealing how document rankings can be improved. A key distinction lies in the nature of

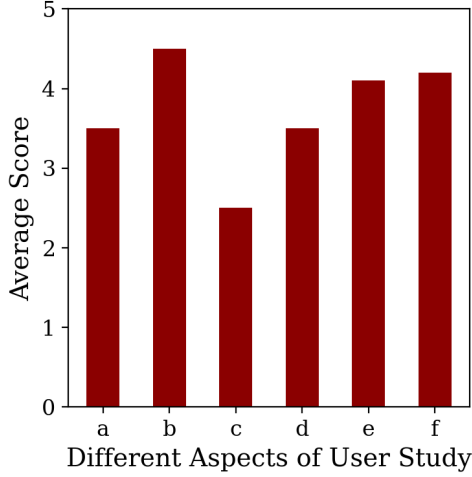


Figure 2: Average Rank shift by CFIR for BM25, DRMM, DSSM, ColBERT, MonoT5 and Splade




Figure 3: Average Semantic Similarity between original documents and the corresponding counterfactual explanation Terms for BM25, DRMM, DSSM, ColBERT, MonoT5 and Splade


intervention: adversarial methods typically aim to introduce minimal perturbations often by substituting content, including important terms to preserve the original semantics while deceiving the model. In our case, CFIR explicitly seeks to identify new terms that, when added to a document, improve its rank, thereby highlighting what informative aspects were absent. Replacing important terms is not useful in counterfactual setup, as it fails to address what the document was lacking from the model’s perspective. This formulation is particularly relevant for understanding model behavior, including uncovering potentially problematic model preferences (e.g., prior studies have observed gender bias in ranking systems). By identifying helpful additions, such as gendered terms, CFIR can reveal latent model sensitivities. Importantly, unlike adversarial attacks, the size of the added term set is also not constrained in CFIR (Avg. New Words column in Table 1 shows maximum 16.81 new words per explanation), as the focus is on explanatory sufficiency rather than minimality. However, for comparison, we have evaluated the performance of CFIR against the PRADA (Wu et al., 2022a) model which replaces certain words in a document to improve its ranking. Table 8 in Appendix H shows that CFIR performs better than PRADA for both ColBERT and MonoT5 in terms of Fidelity score.

Table 2 shows a sample of example terms extracted by our proposed approach. The words shown in Table 2 have improved the ranking of a docID with respect to the queries shown.


**Further Analysis** Figure 2 shows the average

change in rank after introducing the explanation terms suggested by the CFIR setup. Figure 2 essentially demonstrates the actionability introduced by the counterfactual explanation terms. The two things to observe from Figure 2 are firstly, the average rank shift is greater for documents than for passages. Table 1 shows that ColBERT achieved a significantly higher fidelity score (16<sup>th</sup> row) and a larger average rank shift compared to the other models, as also seen in Figure 2. Figure 3 shows the average cosine similarity computed between documents and the corresponding explanation terms. For both documents and the explanation terms we use pretrained BERT representations to compute the similarity. It can be observed from Figure 3 that the cosine similarity for the representation learning based retrieval models (i.e. ColBERT, MonoT5) are higher than the other retrieval models in general.

**Parameter Sensitivity Analysis** In Table 1, we observed that for most of the retrieval models the performance of the counterfactual explainer follows similar trend both in MSMARCO passage and document dataset (i.e. the best performing model in terms of fidelity score is same in most of the cases). As a result, we conducted parameter sensitivity experiments only on MSMARCO passage dataset. Figure 4 (a) shows the variance in Fidelity score with respect to the K value in Top-K. In Figure 4 (b) we show the variance of FD score with respect to the number of most significant words (i.e.  $n$ ) used to construct the document vector. It is clearly visible from Figure 4



(a)



(b)

Figure 4: Counterfactual Classifier Performance Variance with Top-K and Counterfactual Performance Variance with variation of number of Counterfactuals

(b) that with an increase in the number of counterfactuals, there is a decrease in the performance of the counterfactual classifier. It can be observed that for  $n = 10$  the best performance is achieved. Intuitively, as the number of words increases, the feature vector grows exponentially, making it challenging to train the classifier effectively.

**Qualitative Evaluation of Explanations** We conducted a user study involving three researchers with doctoral degrees in IR to estimate the quality of explanations. Each annotator was provided with 30 documents from the MS MARCO passage collection, along with the corresponding queries, ranked lists, and explanation terms generated by CFIR applied to the best-performing model, MonoT5 (shown in Table 1). Further details about the experiment setup is given in Appendix J.

Users were asked to assess the quality of explanations across six dimensions: (a) *Intuitiveness* how intuitive the explanation terms appeared given the query, document, and ranking context, with knowledge of the retrieval model; (b) *Non-intuitiveness* the extent to which explanations felt unexpected or misaligned with the query-document pair; (c) *Query Relatedness* whether the explanation terms were semantically related to the query; (d) *Document Relatedness* whether the explanation terms aligned with the overall topic of the document; (e) *Informativeness* whether the terms were meaningful and content-rich rather than generic or uninformative (e.g. mostly stop words); and (f) *Diversity* whether the explanation terms covered varied semantic aspects. For each aspect the users were asked to put a score between 0 to 5. Figure 5 shows that in general the explanation terms are intuitive and more similar to the document topic compared query topic (as ex-




Figure 5: Qualitative Assessment of Generated Explanations over a) Intuitiveness b) Non-Intuitiveness c) Query Relatedness d) Document Relatedness e) Informativeness f) Diversity)

pected due to use of document similarity criteria in the loss function in Equation 1). The explanation terms are also quite diverse. The non-intuitiveness score is quite low which shows that most of explanation terms follow an IR practitioner’s intuition.

## 6 Conclusion

In this paper, we propose a counterfactual setup for a query-document pair and a retrieval model. Our experiments show that the proposed approach on an average 70% cases for both in short and long documents could successfully improve the ranking. In the future, we would like to explore different explanation units for the counterfactual setup.

## 7 Limitations

One of the limitations of this work is that we assume that top 10 or 20 words (based on tf-idf



weights) within a document play the most important part in improving the rank of a document. However, theoretically speaking we should consider all the words present in a document to determine the most influential words for a retrieval model. We have used top tf-idf words (Similar to statistical retrieval models) to reduce the computational complexity of our experiments and we have seen that increasing the number of top words doesn't affect the performance of the model that much.

## 8 Ethical Considerations

In this work, we have used publicly available search query log and document collection to demonstrate counterfactual explanation. No sensitive data was used in this experiment. As a result of this there is no particular ethical concern associated with this work. If there is any kind of bias present in the search log data that effect can be observed within our approach. However mitigating that bias was beyond the scope of this work

## References

- Avishek Anand, Lijun Lyu, Maximilian Idahl, Yumeng Wang, Jonas Wallat, and Zijian Zhang. 2022. Explainable information retrieval: A survey.
- Payal Bajaj, Daniel Campos, Nick Craswell, Li Deng, Jianfeng Gao, Xiaodong Liu, Rangan Majumder, Andrew McNamara, Bhaskar Mitra, Tri Nguyen, Mir Rosenberg, Xia Song, Alina Stoica, Saurabh Tiwary, and Tong Wang. 2016. Ms marco: A human generated machine reading comprehension dataset. In *InCoCo@NIPS*.
- Alexander Bondarenko, Maik Fröbe, Jan Heinrich Reimer, Benno Stein, Michael Völske, and Matthias Hagen. 2022. Axiomatic retrieval experimentation with ir axioms. In *Proc. of SIGIR 2022*, pages 3131–3140.
- Miguel Á Carreira-Perpiñán and Suryabhan Singh Hada. 2021. Counterfactual explanations for oblique decision trees: Exact, efficient algorithms. In *Proceedings of the AAAI conference on artificial intelligence*, volume 35, pages 6903–6911.
- Nick Craswell, Bhaskar Mitra, Emine Yilmaz, and Daniel Campos. 2021. [Overview of the TREC 2020 deep learning track](#). *CoRR*, abs/2102.07662.
- Nick Craswell, Bhaskar Mitra, Emine Yilmaz, Daniel Campos, Jimmy Lin, Ellen M. Voorhees, and Ian Soboroff. 2023. Overview of the trec 2022 deep learning track. In *Text REtrieval Conference (TREC)*. NIST, TREC.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. BERT: Pre-training of deep bidirectional transformers for language understanding. In *NAACL-HLT*.
- Gokhan Egri and Coskun Bayrak. 2014. [The role of search engine optimization on keeping the user on the site](#). *Procedia Computer Science*, 36:335–342. Complex Adaptive Systems Philadelphia, PA November 3-5, 2014.
- Anett Erdmann, Ramón Arilla, and José M. Ponzoa. 2022. [Search engine optimization: The long-term strategy of keyword choice](#). *Journal of Business Research*, 144:650–662.
- Thibault Formal, Benjamin Piwowarski, and Stéphane Clinchant. 2021. Splade: Sparse lexical and expansion model for first stage ranking. In *Proceedings of the 44th International ACM SIGIR Conference on Research and Development in Information Retrieval*, page 2288–2292.
- Maarten Grootendorst. 2020. [Keybert: Minimal keyword extraction with bert](#).
- Jiafeng Guo, Yixing Fan, Qingyao Ai, and W. Bruce Croft. 2016. A deep relevance matching model for ad-hoc retrieval. In *Proceedings of the 25th ACM International on Conference on Information and Knowledge Management, CIKM '16*, page 55–64, New York, NY, USA. Association for Computing Machinery.
- Jiafeng Guo, Yixing Fan, Xiang Ji, and Xueqi Cheng. 2019. Matchzoo: A learning, practicing, and developing system for neural text matching. In *Proceedings of the 42nd International ACM SIGIR Conference on Research and Development in Information Retrieval, SIGIR'19*, pages 1297–1300.
- Faisal Hamman, Erfaun Noorani, Saumitra Mishra, Daniele Magazzeni, and Sanghamitra Dutta. 2023. Robust counterfactual explanations for neural networks with probabilistic guarantees. In *Proceedings of the 40th International Conference on Machine Learning, ICML'23*. JMLR.org.
- Po-Sen Huang, Xiaodong He, Jianfeng Gao, Li Deng, Alex Acero, and Larry Heck. 2013. Learning deep structured semantic models for web search using clickthrough data. In *Proceedings of the 22nd ACM International Conference on Information & Knowledge Management, CIKM '13*, page 2333–2338.
- Kentaro Kanamori, Takuya Takagi, Ken Kobayashi, Yuichi Ike, Kento Uemura, and Hiroki Arimura. 2021. Ordered counterfactual explanation by mixed-integer linear optimization. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 35, pages 11564–11574.
- Amir-Hossein Karimi, Gilles Barthe, Borja Balle, and Isabel Valera. 2020. Model-agnostic counterfactual explanations for consequential decisions. In *International Conference on Artificial Intelligence and Statistics*, pages 895–905. PMLR.

|     |   |     |
|-----|---|-----|
| 697 | Omar Khattab and Matei Zaharia. 2020. Colbert: Efficient and effective passage search via contextualized late interaction over bert. In <i>Proceedings of the 43rd International ACM SIGIR Conference on Research and Development in Information Retrieval</i> , page 39–48.  | 753 |
| 698 |   | 754 |
| 699 |   | 755 |
| 700 |   |     |
| 701 |   |     |
| 702 |   |     |
| 703 | Jimmy Lin, Xueguang Ma, Sheng-Chieh Lin, Jheng-Hong Yang, Ronak Pradeep, and Rodrigo Nogueira. 2021. Pyserini: A python toolkit for reproducible information retrieval research with sparse and dense representations. In <i>Proceedings of the 44th International ACM SIGIR Conference on Research and Development in Information Retrieval</i> , page 2356–2362.                                      | 756 |
| 704 |   | 757 |
| 705 |   | 758 |
| 706 |   | 759 |
| 707 |   | 760 |
| 708 |   |     |
| 709 |   |     |
| 710 |   |     |
| 711 | Yu-An Liu, Ruqing Zhang, Jiafeng Guo, Maarten de Rijke, Wei Chen, Yixing Fan, and Xueqi Cheng. 2023. Black-box adversarial attacks against dense retrieval models: A multi-view contrastive learning method. In <i>Proceedings of the 32nd ACM International Conference on Information and Knowledge Management</i> , CIKM '23, page 1647–1656, New York, NY, USA. Association for Computing Machinery. | 761 |
| 712 |   | 762 |
| 713 |   | 763 |
| 714 |   | 764 |
| 715 |   | 765 |
| 716 |   |     |
| 717 |   |     |
| 718 |   |     |
| 719 |   |     |
| 720 | Lijun Lyu and Avishek Anand. 2023. Listwise explanations for ranking models using multiple explainers. In <i>Advances in Information Retrieval: 45th European Conference on Information Retrieval, ECIR 2023, Dublin, Ireland, April 2–6, 2023, Proceedings, Part I</i> , page 653–668, Berlin, Heidelberg. Springer-Verlag.  | 766 |
| 721 |   | 767 |
| 722 |   | 768 |
| 723 |   | 769 |
| 724 |   |     |
| 725 |   |     |
| 726 |   |     |
| 727 | Tomas Mikolov, Ilya Sutskever, Kai Chen, Greg S Corrado, and Jeff Dean. 2013. Distributed representations of words and phrases and their compositional-ity. In <i>NeurIPS</i> .   | 770 |
| 728 |   | 771 |
| 729 |   | 772 |
| 730 |   | 773 |
| 731 |   | 774 |
| 732 | Ramaravind K Mothilal, Amit Sharma, and Chenhao Tan. 2020. Explaining machine learning classifiers through diverse counterfactual explanations. In <i>Proceedings of the 2020 Conference on Fairness, Accountability, and Transparency</i> , pages 607–617.   | 775 |
| 733 |   | 776 |
| 734 |   | 777 |
| 735 |   | 778 |
| 736 |   | 779 |
| 737 | Rodrigo Nogueira, Zhiying Jiang, Ronak Pradeep, and Jimmy Lin. 2020. Document ranking with a pre-trained sequence-to-sequence model. In <i>Findings of the Association for Computational Linguistics: EMNLP 2020</i> , pages 708–718, Online. Association for Computational Linguistics.  | 780 |
| 738 |   | 781 |
| 739 |   | 782 |
| 740 |   | 783 |
| 741 |   | 784 |
| 742 |   |     |
| 743 | Axel Parmentier and Thibaut Vidal. 2021. Optimal counterfactual explanations in tree ensembles. In <i>International conference on machine learning</i> , pages 8422–8431. PMLR.   | 785 |
| 744 |   | 786 |
| 745 |   |     |
| 746 | Martin Pawelczyk, Chirag Agarwal, Shalmali Joshi, Sohini Upadhyay, and Himabindu Lakkaraju. 2022. Exploring counterfactual explanations through the lens of adversarial examples: A theoretical and empirical analysis. In <i>International Conference on Artificial Intelligence and Statistics</i> , pages 4574–4594. PMLR.   | 787 |
| 747 |   | 788 |
| 748 |   | 789 |
| 749 |   | 790 |
| 750 |   | 791 |
| 751 |   |     |
| 752 |   |     |
|     | Judea Pearl. 2018. Theoretical impediments to machine learning with seven sparks from the causal revolution. <i>arXiv preprint arXiv:1801.04016</i> .   | 792 |
|     |   | 793 |
|     |   | 794 |
|     | Gustavo Penha, Eyal Krikon, and Vanessa Murdock. 2022. Pairwise review-based explanations for voice product search. In <i>ACM SIGIR Conference on Human Information Interaction and Retrieval</i> , pages 300–304.  | 795 |
|     |   | 796 |
|     |   | 797 |
|     |   | 798 |
|     | Navid Rekabsaz and Markus Schedl. 2020. Do neural ranking models intensify gender bias? In <i>Proceedings of the 43rd International ACM SIGIR Conference on Research and Development in Information Retrieval</i> , page 2065–2068.   | 799 |
|     |   | 800 |
|     |   | 801 |
|     |   | 802 |
|     |   | 803 |
|     |   | 804 |
|     | Marco Tulio Ribeiro, Sameer Singh, and Carlos Guestrin. 2016. "why should i trust you?": Explaining the predictions of any classifier. In <i>Proc. of SIGKDD 2016</i> , page 1135–1144.   |     |
|     |   |     |
|     | Jaspreet Singh and Avishek Anand. 2019. Exs: Explainable search using local model agnostic interpretability. In <i>Proceedings of the Twelfth ACM International Conference on Web Search and Data Mining</i> , WSDM '19, page 770–773.  |     |
|     |   |     |
|     | Jaspreet Singh and Avishek Anand. 2020. Model agnostic interpretability of rankers via intent modelling. In <i>Proceedings of the 2020 Conference on Fairness, Accountability, and Transparency</i> , pages 618–628.  |     |
|     |   |     |
|     | Arnaud Van Looveren and Janis Klaise. 2021. Interpretable counterfactual explanations guided by prototypes. In <i>Joint European Conference on Machine Learning and Knowledge Discovery in Databases</i> , pages 650–665. Springer.   |     |
|     |   |     |
|     | Ellen Voorhees. 2005. <a href="#">Overview of the trec 2004 robust retrieval track</a> .  |     |
|     |   |     |
|     | Chen Wu, Ruqing Zhang, Jiafeng Guo, Maarten de Rijke, Yixing Fan, and Xueqi Cheng. 2022a. <a href="#">Prada: Practical black-box adversarial attacks against neural ranking models</a> . <i>ArXiv preprint</i> , abs/2204.01321.  |     |
|     |   |     |
|     | Chen Wu, Ruqing Zhang, Jiafeng Guo, Yixing Fan, and Xueqi Cheng. 2022b. Are neural ranking models robust? <i>ACM Trans. Inf. Syst.</i> , 41(2).   |     |
|     |   |     |
|     | Zhichao Xu, Hemank Lamba, Qingyao Ai, Joel Tetreault, and Alex Jaimes. 2024. <a href="#">Counterfactual editing for search result explanation</a> . <i>Preprint</i> , arXiv:2301.10389.   |     |
|     |   |     |
|     | Puxuan Yu, Razieh Rahimi, and James Allan. 2022. Towards explainable search results: a listwise explanation generator. In <i>Proceedings of the 45th International ACM SIGIR Conference on Research and Development in Information Retrieval</i> , pages 669–680.   |     |

## A Retrieval Models

The five different retrieval models used in our experiment are described as follows.

**BM25:** BM25<sup>2</sup> is a statistical retrieval model where the similarity between a query and a document is computed based on the term frequency of the query words present in the document, document frequency of the query words and also the document length.

**DRMM:** Deep Relevance Matching Model (DRMM) Guo et al. (2016) is a neural retrieval model where the semantic similarity between each pair of tokens corresponding to a query and a document is computed to estimate the final relevance score of a document.

**DSSM:** Deep Semantic Similarity Model (DSSM) Huang et al. (2013) is another neural retrieval model which uses word hashing techniques to compute the semantic similarity between a query and a document.

**ColBERT:** Contextualized Late Interaction over BERT (ColBERT) (Khattab and Zaharia, 2020), is an advanced neural retrieval model which exploits late interaction techniques based on BERT (Devlin et al., 2019) based representations of both query and document for retrieval.

**MonoT5:** MonoT5 (Nogueira et al., 2020) is a sequence-to-sequence model fine-tuned to predict the relevance of a query-document pair.

**Splade:** Splade (Formal et al., 2021)(Sparse Lexical and Expansion Model for Information Retrieval) combines the sparse interpretability of traditional IR models (like BM25) with the semantic power of deep learning. Unlike dense retrieval models that rely on vector similarity in embedding space, SPLADE encodes queries and documents into sparse high-dimensional vectors—essentially performing learned term expansions in a way that mimics the inverted index structure used in classic IR systems.

## B Retrieval Performance of IR Models

We use Lin et al. (2021) toolkit for implementing BM25 and MonoT5 and Splade. For DRMM and DSSM, we use the implementation released by the study in Guo et al. (2019). For passage ranking we varied the parameters in a grid search and we took the configuration producing best MRR@10 value on TREC DL (Craswell et al., 2021) test

set. For both DRMM and DSSM experiments on MSMARCO data, the parameters were set as suggested in (Wu et al., 2022b). The MRR@10 values are reported in Table 7 in Appendix B. For DRMM and DSSM, we use randomly chosen 100K query pairs from the MSMARCO training dataset to train the model.

The machine used to run counterfactual experiments on retrieval model has 1 A100 GPU and 40 GB memory.

| Model   | MRR@10          |                  |
|---------|-----------------|------------------|
|         | MSMARCO Passage | MSMARCO Document |
| BM25    | 0.1874          | 0.2184           |
| DRMM    | 0.1623          | 0.1168           |
| DSSM    | 0.1320          | 0.1168           |
| ColBERT | 0.3481          | 0.3469           |
| MonoT5  | 0.3904          | 0.3827           |
| Splade  | 0.3813          | 0.3721           |

Table 3: Retrieval Model Performance on MSMARCO passage and document

## C Dataset Statistics

The dataset statistics for all the experiments are given in Table 4

|          |            | MS MARCO Passage | MS MARCO Document | TREC Robust |
|----------|------------|------------------|-------------------|-------------|
| Query    | Avg Length | 5.9              | 6.9               | 7.18        |
| Document | Avg Length | 64.9             | 1134.2            | 150.12      |
| Query    | #Instances | 100              | 100               | 100         |
| Document | #Instances | 500              | 500               | 500         |

Table 4: Dataset Details for Counterfactual Setup

## D Example of Input and Output to Classifier

Given an input query, we employ a Lucene-Searcher with MSMARCO Index to retrieve the Top-K documents. The feature vector construction process follows these steps:

For each document, we:

1. Extract the top n words based on their Imp(d) values
2. Construct a vocabulary  $V$  as the union of all top 10 words across documents
3. Note that  $|V|$  typically falls in the range of 150-180 words

The feature vector for each document has dimension  $|V|$ , where each component represents the value from the Imp(d) of the corresponding word from the vocabulary. Formally:

<sup>2</sup>[https://en.wikipedia.org/wiki/Okapi\\_BM25](https://en.wikipedia.org/wiki/Okapi_BM25)

$$d_{vec} \in R^{|V|}$$

Labels are assigned according to the following criterion:

$$\text{label} = \begin{cases} 1 & \text{for top } K \text{ documents} \\ 0 & \text{for remaining documents} \end{cases}$$

Example feature vectors and their corresponding counterfactuals generated using (Mothilal et al., 2020) are shown in Table 6. Since  $|V|$  is 150 in our experiments, hence in Table 6 we have only shown the term frequencies of the words present in each document. For other words the terms frequency values will be zero in  $d_{vec}$ .

| Existing Explanation Methods                  | Word Overlap |
|---|--------------|
| PointWise Explanation (Singh and Anand, 2019) | 21.46%       |
| ListWise Explanation (Lyu and Anand, 2023)    | 9.57%        |

Table 5: Comparison of CFIR with Existing ExIR Approaches

## E Existing EXIR approaches vs. CFIR

The existing literature aims to explain the significance of a document, a set of documents, or a pair of documents through various explanation methods. Nonetheless, our proposed approach diverges fundamentally from prior work in that we seek to demonstrate how the absence or frequency of certain tokens impacts document relevance. In this section, we examine whether there is any intersection between the two sets of tokens described earlier.

**Pointwise Explanation Approach** As outlined in Section 2.2, existing pointwise explanation methods elucidate why a specific document aligns with a given query within a retrieval model. Similarly, our proposed approach operates on individual documents and queries, albeit with a distinct objective. Here, we analyze the overlap between the explanations generated by the pointwise explanation method and those derived from our model, as presented in Table 7. This comparison was conducted on 50 pairs of documents.

**Listwise Explanation Approach** In Section 2, it is explained that listwise explanations typically aim to demonstrate the relevance of a list of documents to a given query. In listwise setup, one set of explanation terms are extracted for a list of documents, a query, and a retrieval model. Conversely, in our approach, we generate distinct explanations for each query-word pair. Therefore, to compare

listwise explanations with our method, we aggregate all individual explanations obtained for each document-query pair in the list to create a unified explanation set for the entire list corresponding to a query. The resulting overlap is presented in Table 7.

## F Counterfactual Optimization Framework

The different parts of Equation 2 are described here. The  $y_{loss}$  in Equation 2 is a hinge loss function as defined in Equation 4. In Equation 4  $z$  is  $-1$  when  $y = 0$  otherwise,  $z = 1$ .  $\text{logit}(f_{\{M,q\}}(c_i))$  is the logit values obtained from the classifier ( $f_{\{M,q\}}$ ) when the counterfactual  $c_i$  is given as input.

$$y_{loss} = \max(0, 1 - z * \text{logit}(f_{\{M,q\}}(c_i))) \quad (4)$$

The distance function ( $\text{dist}(c_i, d)$ ) in Equation 2 is computed using the formula given in Equation 5. In Equation 5,  $V$  represents the vocabulary set used to represent the document vectors ( $d_{vec}$ ). In Equation 5, the value of  $I$  is equal to 1 if the corresponding term is present in both the counterfactual input  $c$  and the original input  $d$ , otherwise it is set to 0.

$$\text{dist}(c, d) = \sum_{p=1}^V I(c_p \neq d_p) \quad (5)$$

The diversity in above equation is defined by the formula described in Equation 6. In equation 6,  $K_{i,j}$  is equal to  $\frac{1}{1+\text{dist}(c_i, c_j)}$ .  $\text{dist}(c_i, c_j)$  calculates the distance between two counterfactuals  $c_i$  and  $c_j$ .

$$\text{div}(c_1, \dots, c_k) = \sum_{i,j} \det(K_{i,j}) \quad (6)$$

## G Parameters for Counterfactual Setup

The value of  $\lambda_1$  and  $\lambda_2$  is set to 1 and 0.5 respectively in Equation 2. The value of  $k$  in Equation 2 is set to  $k = 3$ . In all our experiments in Table 1, we have observed that for  $K = 3$  and onward we have always found a counterfactual explanation for each query-document pair where only words were added for the desired counterfactual outcome.

## H Adversarial Attacks vs. Counterfactual Explanation

Here we show the performance of our proposed counterfactual explanation approach with an exist-



| docID   | Feature Vector   |
|---------|--|
| 3686955 | [prohibition:2.0, amendment:2.0, under:1.0, dwindled:1.0, eighteenth:1.0, repeal:1.0, repealed:3.0, states:1.0, 1933: 1.0, ratification: 1.0]      |
| 6159679 | [membrane:5.0, lipids:3.0, remainder:2.0, proteins:3.0, biochemical:2.0, 80:2.0, role:2.0, percent:2.0]  |
| 5217641 | [waves:6.0, transverse:5.0, electromagnetic:3.0, oscillations:2.0, vibrations:2.0, travel:2.0, radiation:2.0, angles:2.0, transfer:2.0, types:3.0] |

Table 6: Sample Feature Vector Corresponding to three different documents

| Existing Explanation Methods                  | Word Overlap |
|---|--------------|
| PointWise Explanation (Singh and Anand, 2019) | 21.46%       |
| ListWise Explanation (Lyu and Anand, 2023)    | 9.57%        |

Table 7: Comparison of CFIR with Existing ExIR Approaches

| Retrieval Model | FD in PRADA | FD in CFIR |
|-----------------|-------------|------------|
| ColBERT         | 74%         | 75%        |
| MonoT5          | 80%         | 82%        |

Table 8: Performance of CFIR vs. Adversarial Attack Model PRADA (Wu et al., 2022a)

ing adversarial model named PRADA (Wu et al., 2022a). We use the MSMARCO passage dataset as the target corpus. We use same test set (as described in Table 4) as used in the first column of Table 1 in this experiment. Table 8 shows the results in terms of Fidelity score.

## I Implementation of Imp(d)

We explored three ways to compute the top  $n$  words from each document. Each one of them is described as follows.

**TF-IDF Approach:** In this approach we choose top  $n$  words from a document based on their TF-IDF weight.

**KEYBERT Approach:** In this approach we use the model proposed in (Grootendorst, 2020) to extract keywords from a string.

**BERT-Based Similarity(BERTSim):** In this approach we compute the similarity between the BERT based representation of the query text and each token of the document and then we sort all the tokens based on the similarity.

Table 9 shows the performance of the above-mentioned three approaches in MSMARCO passage dataset and ColBERT retrieval model.  $n = 10$  for the experiments shown in Table 9. From Table 9, we can conclude that the BERT-based similarity approach works the best for the  $Imp(d)$  function. hence for all the results reported in Table 1, we use the BERTSim approach in the  $Imp(d)$  function.

| $Imp(d)$ Approach | FD  |
|-------------------|-----|
| TFIDF             | 74% |
| KeyBERT           | 70% |
| BERTSim           | 75% |

Table 9: Performance of Different Approaches in  $Imp(d)$ .

## J User Study

In the user study we didn’t record any personal information of any of the users. We only recorded their judgment about the output of the proposed methodology for the study. We have also used data which is publicly available for IR research. Hence no ethics approval was required for the study. All the researchers were made aware of the of the use of their assessment in this research.