

Explaining Documents' Relevance to Search Queries

RAZIEH RAHIMI, Center for Intelligent Information Retrieval, University of Massachusetts Amherst

YOUNGWOON KIM, Center for Intelligent Information Retrieval, University of Massachusetts Amherst

HAMED ZAMANI, Center for Intelligent Information Retrieval, University of Massachusetts Amherst

JAMES ALLAN, Center for Intelligent Information Retrieval, University of Massachusetts Amherst

We present **GenEx**, a generative model to explain search results to users beyond just showing matches between query and document words. Adding GenEx explanations to search results greatly impacts user satisfaction and search performance. Search engines mostly provide document titles, URLs, and snippets for each result. Existing model-agnostic explanation methods similarly focus on word matching or content-based features. However, a recent user study shows that word matching features are quite obvious to users and thus of slight value. GenEx explains a search result by providing a terse description for the query aspect covered by that result. We cast the task as a sequence transduction problem and propose a novel model based on the Transformer architecture. To represent documents with respect to the given queries and yet not generate the queries themselves as explanations, two *query-attention layers* and *masked-query decoding* are added to the Transformer architecture. The model is trained without using any human-generated explanations. Training data are instead automatically constructed to ensure a tolerable noise level and a generalizable learned model. Experimental evaluation shows that our explanation models significantly outperform the baseline models. Evaluation through user studies also demonstrates that our explanation model generates short yet useful explanations.

CCS Concepts: • **Computer systems organization** → **Embedded systems**; *Redundancy*; Robotics; • **Networks** → Network reliability.

Additional Key Words and Phrases: Relevance explanation, Black-box explainer, Content-based explanation

ACM Reference Format:

Razieh Rahimi, Youngwoo Kim, Hamed Zamani, and James Allan. 2018. Explaining Documents' Relevance to Search Queries. *J. ACM* 37, 4, Article 111 (August 2018), 24 pages. <https://doi.org/10.1145/1122445.1122456>

1 INTRODUCTION

We focus on a new class of explanations for search results that aims to help users gain deeper understanding of search results and that is suitable for different search scenarios, from ad-hoc to conversational information seeking, and from desktop computers to voice-based systems.

Search engine result pages currently provide document snippets in addition to document titles for each result, where the snippets are typically 2- or 3-line extracts highlighting the query words in the documents' contents. Although it is known that the quality of that summary can have a significant effect on user interactions [36], search snippets oftentimes are not coherent [64] and fail

Authors' addresses: Razieh Rahimi, Center for Intelligent Information Retrieval, University of Massachusetts Amherst, 140 Governors Dr, Amherst, Massachusetts, 01002, rahimi@cs.umass.edu; Youngwoo Kim, Center for Intelligent Information Retrieval, University of Massachusetts Amherst, youngwookim@cs.umass.edu; Hamed Zamani, Center for Intelligent Information Retrieval, University of Massachusetts Amherst, zamani@cs.umass.edu; James Allan, Center for Intelligent Information Retrieval, University of Massachusetts Amherst, allan@cs.umass.edu.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

© 2018 Association for Computing Machinery.

0004-5411/2018/8-ART111 \$15.00

<https://doi.org/10.1145/1122445.1122456>

to explain the documents' relevance to the submitted query [62]. Thus the users may be confused as why a document is presented in search results.

To address all these issues, we propose GENEx, an approach that generates terse explanations – on the level of noun phrases – for search results describing what aspect of the query is covered by a retrieved document. For example, suppose that in response to the query “OBAMACARE” a document is listed that discusses how income is subject to an additional tax. A desired explanation for the document is “IMPACTS ON MEDICARE TAX”, which provides information beyond that of the snippet: “..TAX TO OFFSET THE COSTS OF THE OBAMACARE. THIS TAX FIRST TOOK EFFECT IN 2013...” (as automatically generated by the Indri search engine).

We start this work by describing a set of studies to assess the usefulness of explanations like those produced by GENEx – that is, to explore whether the proposed explanations can help users make more accurate and/or faster relevance decisions? In Section 3, we describe and then compare two presentations of search results: 1) showing documents' snippets only, and 2) showing snippets and explanations. We demonstrate that when participants have the explanations, they reach consensus on relevant document in 23% more cases than when they have snippets alone. In addition, participants could detect the relevant document in 7 fewer seconds on average (22% faster) when explanations are provided.

The problem of constructing such explanations does not appear to have been studied previously in works on snippet generation or model-agnostic search result explanation. Prior work on model-agnostic explanation of information retrieval models is either *local*, focusing on explaining individual rankings [59], or *global*, explaining the model behavior as a whole by training a simpler “interpretable” ranker, such as decision trees or a linear ranker [58]. Both types of work focus on explaining content-based features in ranking: which words or word-matching features contributed more in the provided rankings. Although useful in some cases, the recent study by Thomas et al. [62] suggests that users of search systems benefit more from explanations describing documents' relevance *beyond* word matching.

Explaining documents' relevance faces the major challenging step of extracting and conceptually representing the query-related part(s) of the document content (possibly a small part of the document [65]) with respect to generally vague short-keyword queries. In addition, as with the simpler task of general document explanation (such as headline generation [53]), obtaining substantial amounts of manually-labeled training data is often costly and time consuming.

In Section 4, we cast the task of search result explanation as a sequence transduction problem, where an attention-based encoder-decoder architecture first provides a topic-focused contextual representation of a document and then generates desired explanations. We specifically extend the Transformer architecture [67] by introducing a query attention layer in the encoder to represent query-focused parts of documents. We then mask the query in the decoder to generate coherent textual explanations of information in documents that satisfy the user information need.

We propose solutions to automatically obtain training data from the Web to bypass the expensive human labeling process. Our model is thus trained with no manually labeled training data. To build weakly-labeled training data with a noise level that a supervised encoder-decoder model can tolerate and that learns a model that generalizes to open-domain input texts, we combine samples from two sources: (1) Wikipedia articles as more controlled edited content than the entire Web, and (2) anchor texts in a collection of general web pages.

In Section 5, we describe extensive experiments on multiple datasets to evaluate GENEx. Our results show that GENEx significantly outperforms the baselines: it improves BLEU-1 by 67%-73% as well as ROUGE-1 and ROUGE-L by 47% - 51%, all over the original transformer architecture on general test samples from the Web.

Since BLEU and ROUGE do not always correlate with human judgments, we continue in Section 6 with a final study asking people to evaluate the quality of generated explanations, incorporating both relevance to the query and match to the documents' content. The results show that GenEx explanations are preferred over the strongest baseline by a majority of workers in 73% of samples, and tied in another 7%.

All datasets and user annotations collected in this study will be publicly available.

2 RELATED WORK

We review the related previous work on document summarization, snippet generation, and explainable search and recommendation.

2.1 Document Summarization

Document summarization is related to the defined task of search result explanation. As we aim to generate coherent and grammatically readable explanations, abstractive models for text summarization [8, 31, 40, 53] better suit as finishing components of explanation generation models, compared to extractive summarization models [50]. Neural abstractive summarization started by generating headlines from the first sentence of news articles [53] and was then applied to different settings, such as longer text inputs [8, 31, 40]. Although search result explanation is beyond summarizing document contents, even exploiting abstractive summarization techniques to different types of documents in the open-domain Web is not straightforward, as they need domain-specific fine-tuning at the very least to produce reasonable outputs [9].

Dealing with no readily available labeled data, unsupervised training of sequence transduction tasks has been studied in some recent work. Since language understanding is required to generate fluent sequences (texts), using pre-trained language representation models, such as Word2vec and GloVe embeddings [38, 43], ELMo [44], BERT [14], and UNILM [15] can reduce the amount of training data required. The pre-trained language representation model Bert is used for ensuring fluency of generated text in the decoding step or representing input documents in the encoding step of text summarization models [32, 78]. Beyond using pre-trained components in conjunction with or as part of a sequence-to-sequence model, there are some pre-trained models for the sequence generation task [28, 61]. Unsupervised training of machine translation models as another example of sequence transduction has also been investigated in some studies [3, 4, 27, 47].

2.2 Snippet Generation

Snippet generation has been known a special type of document summarization, in which sentences, or sentence fragments, are selected to be presented in a search engine result page (SERP) [66]. It was also called query-biased summarization by Tombros and Sanderson [64]. Snippet generation is an active area of research and has been studied in the context of Web search [70], XML retrieval [23], semantic search [74], and more recently dataset search [75]. Early Web search engines presented query-independent snippets consisting of the first tokens of the result document. Google was the first Web search engine to provide query-biased summaries [66, 70]. Bast and Celikik [5] proposed an efficient solution for extractive snippets by taking advantage of inverted index, a popular data structure used in most information retrieval systems. Recently, Chen et al. [7] proposed abstractive snippet generation as a potential solution to circumvent copyright issues. The authors demonstrated that despite the popularity of extractive snippets in the current search engines, abstractive summarization is equally powerful in terms of user acceptance and expressiveness.

Although snippet generation, and in general query-focused document summarization, is closely related to GenEx explanation, they are fundamentally different. GenEx explanations are terse, consisting of a few words, while snippet generation models try to select or generate a few sentences

or even a paragraph. In addition to the length, the goal of these tasks are different. Query-focused summarization tries to find a set of sentences or passages containing frequent and close occurrences of query tokens, while GenEx aims to describe what such sentences convey about the query, explicitly avoiding the query words themselves. While previous work on snippet generation tries to select or generate a few sentences or even a paragraph, explaining the documents' relevance to a query in a few words is a challenging task.

2.3 Explainable Search and Recommendation

How to define and evaluate interpretability and explainability of machine learning models are discussed in several studies [16, 20, 30, 39, 49]. But, application of machine learning techniques to different tasks can impose task-specific requirements on definition and evaluation of suitable explanations. Explaining search results has been briefly studied in recent years. Describing relationships between entities in queries is considered an explanation of search results and is generated based on incomplete descriptions of relationships in knowledge graph [71, 72]. As users' information needs are very diverse, search result explanations requires description of much more relevance factors than relationships between entities, which we aim to extract and describe. Some models explain search results by providing a set of keywords (with their estimated weights) for each document by training another ranker to simulate the scores of a black-box ranker [51, 57–59, 68]. These models mainly explain search results following the posthoc explanation method for classifiers - LIME [48]. Specifically, Singh and Anand [58] interpret a base ranker by training a second tree-based learning-to-rank model with an interpretable subset of content-based ranking features, such as frequency and TF-IDF. Training data for the second ranker is generated from the outputs of the base ranker. Sen et al. [56] also use basic retrieval heuristics (frequency of a term in a document, frequency of a term in a collection, and length of a document) as explanation features. While these interpretable features seem to be useful for system engineers, how to use them to provide explanations for users is unexplored.

Singh and Anand [59] uses LIME to explain the output of a ranker, which is based on perturbing the instance to be explained. The authors cast the ranking task as a classification task and obtain binary labels of relevant or non-relevant for perturbed documents based on three ways: top-k binary, score-based, and rank-based. These perturbed instances with their binary labels are then fed to the LIME explainability model and visualized as in the original, using bar-charts to show word contributions to the model's decision.

Fernando et al. [17] explore a model-introspective explainability method for neural ranking models. They use the DeepSHAP [33] model to generate explanations and defined five different reference to generate explanations: 1) document only containing **OOV** words, 2) document built by sampling words with low **IDF** values, 3) document consisting of words with low **query-likelihood** scores, 4) document sampled from the **collection** that is not available in the top-1000 ranked list, and 5) document that is sampled from the **bottom** of the top-1000 documents retrieved. They found that DeepSHAP's explanations highly depends on a reference input that needs to be further investigated. The authors also compared DeepSHAP's explanations with those generated by EXS [59] based on LIME and found that they are significantly different. They note that this difference by the two explanation models is concerning, especially in the absence of gold explanations. Verma and Ganguly [69] propose a model-agnostic approach based on a weighted squared loss to explain rankers as well as three sampling approaches to perturb the document in the instance to be explained: uniform, biased, and masked sampling. Explanation features in their work consist of words.

Recently, Singh et al. [60] propose a local model-agnostic method for explaining learning-to-rank models. They define interpretability features based on IR heuristics and propose two metrics,

validity and *completeness*, to generate explanations. They propose a greedy approach to find a subset of the features such that there is a high correlation between the rankings produced by the selected features and the original black-box model, i.e., high validity. They try to jointly maximize completeness, which is defined as the negative of the correlation between non-explanation features and the original ranking. Correlations between ranked lists are measured using the Kendall's Tau.

Helping users interpret search results can be different than, and thus not possible through, presenting (partial) information about how search engines work, such as providing (interpretable) ranking features to users. In a recent study, Thomas et al. [62] investigated how users perceive rankings provided by a search engine, with the goal of finding out what forms of explanations may help users. The authors identified six core concepts that used in ranking at Web scale such as *relevance* and *diversity*. Participants are asked about why each result is chosen. Although *diversity* had been identified as important ranking factor before collecting user responses, less than 1% of users found that search results are presented because of diversification to cover different intents and facets of queries. In more details, *diversity* has the lowest mentions in the collected responses. They mainly performed a set of user studies and online surveys to better understand the mental model of users while using the web search engines. In this work, we propose a model for explaining to what aspect (or facet) of the query, the document is relevant. We believe adding our explanations to SERP could address the issue found by Thomas et al. [62] about query intents and facets.

Explainable recommendation has recently attracted considerable attention [80]. For instance, Ai et al. [2] recently proposed a model based on dynamic relation embedding to produce explanation for recommendation in the context of e-commerce. Content-based models for explanation of recommender systems are closer to explanation of search results than those of collaborative filtering, yet structures in items and user-item interactions are not available in Web search.

2.4 Query Aspects

Mining query aspects to diversify search results [54] is also related to our work. Most approaches in this category are based on query reformulations found in a query log [6, 46] or existing taxonomy or knowledge bases such as Open Directory Project [1]. Other than these sources, Wang et al. [73] cluster the phrases from top retrieved documents to extract query aspects. Ruotsalo et al. [52] propose to model query intent by incorporating feedback from users. Although expected explanations to be generated can reveal query aspects, the defined task is different than existing models for mining query aspects as we are interested in explaining the relevant part of a document's content to a given query without using other documents or other sources of information.

Iwata et al. [24] propose *AspecTiles* to present the degree of relevance of a document to each query aspect where the query aspects are given. Our focus in this work is not about what is a good way to present the generated explanations to users, but about how to generate high-quality explanations. Having explanations, *AspecTiles* can be one way to present them to users.

3 USABILITY STUDY

We start with a user study¹ to investigate whether GENEx-style explanations can actually improve the effectiveness and efficiency of search, as two main goals of explanations proposed by Tintarev and Masthoff [63]. More specifically, we investigate whether providing the explanation to users helps them predict the relevance status of documents faster and/or more accurately. We developed two schemes for presentation of search results, illustrated in Figure 1: in one scheme (on the left), documents are represented by their snippets alone; in the other, documents are presented using the same snippets but also an explanation in a column on the right side of the page.

¹Location omitted for review Institutional Review Board number omitted-for-review.

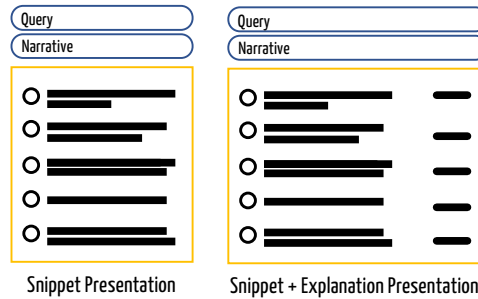


Fig. 1. Different schemes of search result presentation.

We simulate the output of a search engine and creation of explanations for this study. We randomly selected 40 articles from Wikipedia that have at least five sections with headers other than “stop headers” such as *references* or *see also*. Each section of that article is treated as an “aspect” of the query that a user might be interested in. We manually developed TREC-style narrative descriptions of the aspects to reduce ambiguity. For four selected articles, the narratives were sufficiently difficult to construct that we discarded the articles. The snippets were obtained using the Indri toolkit².

For each query (article title), we created a document per section where the section was deemed relevant to its heading (aspect). By construction, each query had at least five aspects: in the end we had 36 queries leading to 240 unique query-aspect pairs each of which had its section content as a single relevant document. (We will use Wikipedia similarly later in Section 5.)

Each search result in this study was associated with a query-aspect pair. We select the aspect’s relevant document and four other documents from the same article so exactly one of the five is relevant. The ordering is random for each pair but is the same for both conditions (with and without the explanation).

We carried out the study on Amazon Mechanical Turk using master workers in the United States who had a high task approval rate. Subjects were provided with search results using one of these presentation schemes and asked to select the document that was relevant to the user’s intent (as described by the aspect narrative). A worker was presented with four search results in sequence but could opt to do additional sets. A set comprised four distinct queries and either all included or all excluded explanations. We captured the selected documents as well as workers’ response time. If the mouse was idle for 2 minutes, we assume the worker is not active and reset the timer. Each presentation of the result list of a query-aspect pair is judged by three different workers.

By construction, every query-aspect pair has a relevant document. Users who selected the wrong document more than twice in a set were deemed to have failed – possibly because they were randomly clicking and not attending to the task – and the set was rejected (and put back in the pool for annotation). However, if workers felt that there was not a relevant document, we required them to describe why they felt that way. If their reasoning was solid, we accepted the set that would otherwise have been rejected. In the end, 240 query-aspect pairs were presented in two styles (with or without explanation) and annotated by three distinct workers, for a total of 1,440 accepted judgments. A total of 51 unique workers participated.

²<https://www.lemurproject.org/indri.php>

Table 1. Usability study results.

Style	\mathcal{K}_f	Correct relevant	Majority relevant	Avg. res. time (s)
Snippet only	0.67	66%	168 (70%)	35.7
Snippet + Expl.	0.92	91%	224 (93%)	23.1

Recall that our question is whether the explanation helped users identify relevant documents faster and/or more accurately. Table 1 summarizes the results of this study. To consider accuracy, we first look at the agreement among the three annotators, measured by Fleiss' Kappa (\mathcal{K}_f) [55]. We find that agreement is substantially greater with the explanation: \mathcal{K}_f increases by 37% from 0.67 (substantial agreement) to 0.92 (almost perfect). The fraction of all judgments that are correct increases by 38% with the explanation and the proportion of the 240 instances where the majority judgment is correct show a 33% climb from 168 to 224. We conclude that the explanation presented greatly increased the consistency and accuracy of identifying relevant documents.

We also compare the average response time for query-aspect pairs. The average time for selecting the correct relevant document decreases from 32.4 to 25.4 seconds when explanations are also provided, which is a 22% decrease. To be sure we were ignoring times when the worker was perhaps unfaithful to the task (so rapidly clicking), we only include cases where the majority of workers selected the relevant document in *both* presentation schemes. There was one case that the majority voted on non-relevant documents when explanations are provided, but voted on the relevant document given just the snippets. The average response time is thus evaluated over 167 query-aspect pairs. Each query-aspect pair thus had 2 or 3 correct responses. We report the macro-average of response times for correct responses in Table 1, showing a 35% decrease when explanations are also provided for search results.

We highlight that this study used manually generated search results and relevance judgments. Nonetheless, the results strongly suggest that adding explanations to snippets can greatly improve both the accuracy and speed of judging documents for relevance. Buoyed by those results, we next propose GENEx, an approach for creating these explanations.

4 GENERATING TERSE EXPLANATIONS

Neural approaches to explanation generation conceptualize the task as a sequence transduction problem. A major approach to this problem is based on the encoder-decoder architecture, where an encoder processes the input tokens and a decoder generates explanation tokens, autoregressively. In the defined task of explaining documents' relevance to a search query, the encoder is expected to learn a contextual representation of the document, capturing the query related parts of the document. The decoder, on the other hand, should generate a text that explains how the document is relevant to the query. In this section, we introduce **GenEx**, our approach for **generating** documents' relevance **explanations**.

4.1 Problem Formulation

Given a query \mathbf{q} and a document \mathbf{d} , the goal is to learn an abstractive explanation model $\mathbf{e} = \mathcal{F}(\mathbf{q}, \mathbf{d})$ to generate a text sequence \mathbf{e} that explains how the given document \mathbf{d} is relevant to the query \mathbf{q} . The explanation generated by the model may include tokens that do not actually occur in the document. A training instance for this task is thus denoted as the triplet $(\mathbf{q}, \mathbf{d}, \mathbf{e})$.

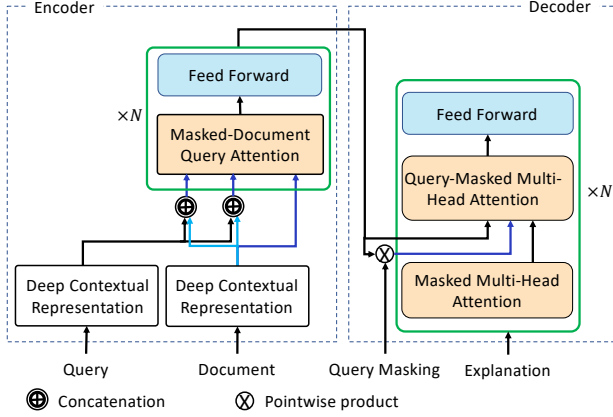


Fig. 2. GenEx Architecture (residual connection and layer normalization for each sub-layer in green boxes have been omitted to enhance the clarity of the figure).

4.2 Input Representation

We first tokenize queries, documents, and target explanations using subword tokenization following Wu et al. [76] and encode tokens with a pre-trained vocabulary embeddings [14].

We then add **positional encodings** to the **token embedding** using the sine function following the Transformer model [67] to capture the relative positions of tokens in input and target sequences given to the encoder and decoder, respectively.

Finally, as the input to the explanation task consists of query and document parts, we add a special token at the end of each part of the input. We also construct a **segment embedding** in two ways as the input to our explanation models: (1) segment embeddings indicating whether a token belongs to the query or to the document, and (2) embeddings indicating whether a token occurs in the query or not. The choice of segment embedding is based on the model architecture and will be discussed later.

4.3 Contextual Encoding of Query and Document

A straightforward solution to handle the two-input characteristic of the explanation task is to concatenate query and document tokens separated using a special token, and then feed the obtained vector to the encoder of a sequence-to-sequence model, such as the Transformer encoder [67]. The input sequence to the model is thus as follows:

$$(q^1, \dots, q^m, \sigma, d^1, \dots, d^n), \quad (1)$$

where σ is a special separation token, similar to the input of several Bert based models for different tasks such as passage/document ranking [41] or question answering. We observed, however, that concatenation of short keyword queries with long documents does not lead to sufficient attention weights from document tokens to the query tokens. This is while learning a query-focused representation of documents is crucial for generating explanations.

To learn proper attention from document tokens to query tokens, we first build two input sequences for a given query and document. These inputs are then separately fed to the Transformer-based encoder model with shared weights. The self-attention mechanism in the Transformer leads to contextual representations of query and document tokens. The query and document

representations obtained by this encoding component are denoted as z_q and z_d , respectively. This architecture also allows pre-computation of document representations similar to the recent Bert-based rankers [25, 34].

4.4 Query Attention Layer

To obtain query-focused representations of documents, we propose another encoder model on top of the learned contextual representations of document tokens. The architecture is shown in Figure 2. This second encoder model is based on the Transformer encoder where self-attention layers are replaced with masked-document query attention layers. This layer consists of three parameter matrices W_Q , W_K , and W_V . The query attention layer computes the following attention matrices based on inputs different from the self-attention mechanism:

$$Q = z_d \times W_Q, \quad (2)$$

$$K = (z_q \parallel z_d) \times W_K, \quad (3)$$

$$V = (z_q \parallel z_d) \times W_V, \quad (4)$$

where \parallel donates concatenation of representations build from previous encoder layers. The representation of each document token is then updated using the scaled dot-product attention function of the Transformer as below:

$$\text{Attention}(Q, K, V) = \text{softmax}\left(\frac{QK^T}{\sqrt{k}}\right)V, \quad (5)$$

where k is the dimension of the keys K . However, we mask the scaled dot product in a way to prevent a document token from attending to other document tokens. Therefore, the output representation of a document token is computed as a weighted sum of the attention values corresponding to itself and query tokens. Lastly, representations are obtained by concatenation of outputs from multiple attention functions, each projecting its inputs to a different subspace. The input and output dimensionality of query attention layers are the same and a stack of N query attention layers is used in the second encoder component. The representations of document tokens from the second encoder component constitute the input of the decoder.

4.5 Query-Masked Decoding

Documents on top of a search engine result page often contain query tokens with high frequencies [35]. In addition, our encoder consists of masked-document query attention layers, which highlight the query related part of the document. Therefore, it is likely that the decoder generates query tokens as the explanation of a document's relevance. However, this is not desired; the model should generate explanations that provide more information other than the query itself, otherwise the generated explanation is useless. Therefore, we generate explanations by extending the Transformers Decoder architecture [67] using a query masking mechanism to reduce the probability of generating query tokens by the decoder. To achieve this, we use a *masked* multi-headed attention for the encoder-decoder attention layer in which the query tokens in both query and document are masked. This makes every decoder position to attend over only non-query tokens in the input sequence.

Note that query tokens are not replaced with a special mask token, because these tokens are central in building the representation of documents. Instead, the representations of query tokens are masked during the decoding process to allow the model to describe the query-related information that the document provides.

Table 2. Wiki dataset statistics, where each sample consists of a query-document-explanation triple, and average values of length are calculated on the specified element of samples.

	Train	Dev	Test
number of samples	6,386,916	336,425	10,000
average length of queries	4.2	4.2	4.2
average length of documents	231.8	231.8	233.2
average length of explanations	2.4	2.5	2.5

5 EXPERIMENTS

5.1 Datasets

We constructed **Wiki training data** using *Wikipedia articles*. In this dataset, each section of a Wikipedia article is treated as a document. The section documents extracted from a Wikipedia article are all relevant to the query built from the article’s title, with their section headers as their explanation labels. For the work discussed here, we use a March 2019 dump of English Wikipedia. Only Wikipedia articles are used for building the dataset, and Wikipedia pages such as disambiguation and redirect pages as well as administration pages such as talk, user, and maintenance pages are removed. Articles with URL links as their titles are also removed. We extract all text from each Wikipedia article, filtering out unwanted data such as HTML markup, using WikiExtractor³ and then discard any article that has fewer than 500 characters of text. A section is not included unless it contains at least 20 tokens. Sections with highly-frequent headers such as “references” and “see also” are removed. The obtained samples are randomly divided into train, validation, and test samples referred to as the Wiki test set. Statistics of Wiki dataset are reported in Table 2. The reported length values are based on tokens obtained by using sub-word tokenization. Explanations having a token not occurred in the query or document are counted as abstractive explanations. Almost 98% of explanations have a token not occurred in the input sequence.

We also built **Anchor training data** to improve the generalizability of learned explanation models. Ideally, we could have used Web search logs for mining different facets of a head query, but they are not publicly available due to concerns about user privacy. However, it has been shown that *anchor text* can effectively simulate real user logs for query reformulation techniques [13]. Based on that work, we approximate query facets using widely available anchor text data, where different anchor texts linking to a particular page approximate query facets relevant to that page. For this purpose, we follow the *subtopic clarification by keyword* observation through user log analysis [22]. We used the external version of the *anchor text for ClueWeb09* dataset⁴ to build training data for the explanation task. All anchor texts to one page that start with the same prefix words are grouped together. The common prefix is then considered as a query, and each suffix is considered as a facet of the query. The linked page is considered to be a relevant document to the query. We performed a number of post-processing steps on the extracted query facets. For examples, facets containing words such as “homepage” or “website” are disregarded. Stopwords such as “of” or “and” are removed from the beginning of query facets.

Clue-Res test data is generated from the ClueWeb09 category B dataset⁵ which is a standard TREC collection and has been used in the TREC Web track for several years [10, 11]. Topics from TREC Web track 2009 to 2012 are used to build test samples. We use the title of a TREC topic as the query, and subtopics of relevant documents in judgments as the set of explanations. Some topics

³<https://github.com/attardi/wikiextractor>

⁴<http://lemurproject.org/clueweb09/anchortext-querylog/>

⁵<http://lemurproject.org/clueweb09/>

have multiple subtopics. Navigational subtopics as well as the first subtopic of queries are discarded since first subtopics mainly provide general description of queries and are not focused around a query facet. Subtopic descriptions are manually rephrased by removing phrases such as “I'd like to find information about”. This test set contains 543 samples.

We built **Passex test data** using the passage ranking dataset of TREC 2019's Deep Learning Track⁶. The TREC passage ranking dataset is built based on the MS MARCO dataset where passages retrieved with respect to the test questions are judged for relevance in much more details. Questions having more than 30 relevant passages in the dataset are chosen for building our evaluation samples. Questions are then manually rephrased as short keyword queries and explanations. Some questions whose conversion to query and explanation were not straightforward are ignored in this step. Also relevant passages with less than 100 tokens are removed. In the end, there are 27 unique queries and 188 evaluation samples in the built Passex dataset.

Note that target outputs in the built test samples are not precise explanations. These test sets are constructed to have approximate out-of-domain samples to guide the development of a generalized explanation model. In the current study, we focus on explaining why a document is relevant to a query by providing the query aspect that the document covers as supportive evidence. Explanation of non-relevant documents to a query is left for future work.

Pre-processing steps. Document texts, queries, and ground truth explanations are lowercased, and encoded using sub-word tokenization following Wu et al. [2016] with Bert vocabulary⁷ of size 30,522. We consider only documents that have more than 20 tokens, and truncate long documents. Documents of the Wikipedia dataset are truncated if required by keeping the first L tokens. We do not truncate the output sequence, but we ignore training samples whose explanations (sub-headings) are longer than 15 tokens. Documents chosen from the ClueWeb dataset for building anchor dataset are extractively summarized as below.

Extractive summarization of long documents. Documents are segmented into sentences using the NLTK toolkit⁸. Sentences with exact occurrences of query terms where document and query terms are stemmed using the Porter stemmer, are chosen to be in the query-biased extractive summary of the document. To capture all query-related information of the document, we also compare the Word2vec pre-trained embeddings of all pairs of query-document terms. A sentence whose most similar term to a query term has a score higher than a defined threshold of 0.8 is retained for the summarized version of the document. If the length of selected sentences is less than L , then additional sentences with the highest contextual similarity to the query are added in the order of their occurrence until no other sentences can be added to the summary while keeping the length lower than or equal to the length cap L .

5.2 Evaluation Metrics

Following the text summarization and machine translation communities, we use BLEU [42] and different variations of the ROUGE metric [29] such as ROUGE-1 (unigrams), ROUGE-2 (bigrams), and ROUGE-L (longest-common substring)⁹ for evaluation of generated explanations: the more words and n-grams that are in common between the predicted output and the target subtopics, the more likely the explanation is to be good. We also use BERTScore [79] to semantically compare generated explanations with reference explanations. BERTScore uses pre-trained contextual embeddings from BERT or its variants to compute cosine similarity between tokens in candidate and reference text segments. Zhang et al. [79] showed that BERTScore better correlates with human judgements

⁶<https://microsoft.github.io/TREC-2019-Deep-Learning/>

⁷https://storage.googleapis.com/bert_models/2018_10_18/uncased_L-12_H-768_A-12.zip

⁸<https://www.nltk.org/>

⁹<https://github.com/google-research/google-research/tree/master/rouge>

and thus is a stronger metric for comparing text generation models. Two-tailed paired t -test is used to test whether the differences between performance of models are statistically significant. Content-based explanation of a single document with respect to a query does not produce any rankings, therefore metrics for evaluation of rankings such as Mean Average Precision are not suitable to measure the quality of explanations.

5.3 Baseline Models

To the best of our knowledge, this is the first study to explain search results from a black-box ranker by generating abstractive and concise explanations beyond term matching. The unique characteristics of the defined problem make models for related task not suitable where the differences are discussed in Section 2. The first baseline is to use TextRank algorithm [37] to extract document keywords. The TextRank algorithm represents a document as a graph of terms linked by co-occurrence relation. Two terms are connected if they co-occur within a window of fixed size, set to 10 terms in our experiments. Then, the PageRank algorithm is applied to the graph to rank terms. The top ranked terms of a document are identified as its keywords. For the defined explanation problem, we are interested in describing a document with respect to a query. To accommodate this setting, we also tried the TextRank algorithm by using topic-sensitive PageRank [21] to rank graph vertices. We refer to this modified version as *topic-sensitive TextRank*. We also compare with KeyBERT [19] which uses BERT to extract keyphrases from a document. This baseline demonstrates the necessity of attending to queries for explanation of document relevance, therefore models for keyword or keyphrase extraction from documents cannot fulfill the task of relevance explanation with noun phrases. Due to the lack of training data for query-focused keyword extraction, we chose the unsupervised and widely used TextRank model as well as recent BERT-base model KeyBERT.

The next group of baselines is based on using local model-agnostic explanation methods to describe the relevance of a document with respect to a query. For this group, we use LIME [48] and Sensitivity [77] that explain the prediction of a black-box classifier for a given input sample. For the purpose of our task, explanation features are defined as document tokens [48, 59]. Specifically, the LIME method generates a number of perturbed samples of the input and learns a linear model, based on the classifier's predictions for perturbed samples. We trained a linear SVM model as the explanation model in our experiments. Features with large coefficients in the learned linear model are considered as explanation. The Sensitivity method estimates the score of a token by measuring the change in the predicted relevance probabilities of perturbed samples that do not include that token [77]. The higher the contribution of a token in the relevance probability is, the more important the token is in understanding the relevance of the document with respect to the query.

We used two different types of rankers as the black-box model whose predictions are to be explained by LIME and sensitivity. These rankers are used to get the relevance probabilities of perturbed documents with respect to the query. The first ranker for this purpose is the classic TF.IDF ranker. For this ranker, documents are represented as bag of words and perturbed samples are obtained by randomly changing one dimension of document vectors to zero. We also used the BERT-based ranker [41] since it has been shown to achieve the current state-of-the-art performance [12, 18]. In this ranker, document tokens that are semantically similar to query tokens can impact the relevance probability of a document, in contrast to the TF.IDF ranker which is only based on exact term matching between a query and document.

The input of the BERT-based ranker is built by concatenating the query and document with a separate token. The pre-trained BERT-based model is fine-tuned for the ranking task using the training data of the MS MARCO passage-ranking dataset [41]. As the BERT-based ranker uses sub-word tokens, the term-level score is obtained by summing the score of each token's subword.

Table 3. Performance of different explanation models. Symbols \blacktriangledown and \triangledown show statistical significant differences with GenEx at levels 0.01 and 0.05, respectively.

	Wiki		Clue-Res		Passex	
	BLEU-1	R-1	BLEU-1	R-1	BLEU-1	R-1
TextRank	0.0862 \blacktriangledown	0.1427 \blacktriangledown	0.0331	0.0435	0.0319 \blacktriangledown	0.0880 \blacktriangledown
TS-TextRank	0.0736 \blacktriangledown	0.1169 \blacktriangledown	0.0313 \blacktriangledown	0.0369	0.0248 \blacktriangledown	0.0715 \blacktriangledown
KeyBERT	0.0567 \blacktriangledown	0.0666 \blacktriangledown	0.0437	0.0404	0.0160 \blacktriangledown	0.0222 \blacktriangledown
LIME + TF.IDF	0.0454 \blacktriangledown	0.0366 \blacktriangledown	0.0281 \blacktriangledown	0.0288 \blacktriangledown	0.0059 \blacktriangledown	0.0044 \blacktriangledown
LIME + BERT	0.0104 \blacktriangledown	0.0085 \blacktriangledown	0.0178 \blacktriangledown	0.0257 \blacktriangledown	0.0089 \blacktriangledown	0.0418 \blacktriangledown
Sensitivity + BERT	0.0062 \blacktriangledown	0.0046 \blacktriangledown	0.0330 \triangledown	0.0398 \blacktriangledown	0.0019 \blacktriangledown	0.0018 \blacktriangledown
GenEx	0.2313	0.3582	0.0520	0.0617	0.1264	0.1179

Table 4. Performance of generative explanation models based on R-1 F-measure, R-2 F-measure, and R-L F-measure metrics. Symbols \blacktriangle and \triangle show statistical significant differences with Orig-Trans at levels 0.01 and 0.05, respectively.

	Wiki				Passex		
	R-1	R-2	R-L		R-1	R-2	R-L
Orig-Trans	0.3180	0.0940	0.3173	Orig-Trans	0.0780	0.0319	0.0752
Seg-q-toks	0.4103\blacktriangle	0.1251\blacktriangle	0.4086\blacktriangle	Seg-q-toks	0.0975	0.0372	0.0993
Sep-q-doc	0.4056 \blacktriangle	0.1247 \blacktriangle	0.4044 \blacktriangle	Sep-q-doc	0.1330\triangle	0.0426	0.1339
GenEx	0.3582 \blacktriangle	0.0868	0.3575	GenEx	0.1179 \triangle	0.0479	0.1179 \triangle

	Clue-Res				Passex		
	R-1	R-2	R-L		R-1	R-2	R-L
Orig-Trans	0.0421	0.0005	0.0418	Orig-Trans	0.0780	0.0319	0.0752
Seg-q-toks	0.0515	0.0031	0.0510	Seg-q-toks	0.0975	0.0372	0.0993
Sep-q-doc	0.0572	0.0043	0.0562	Sep-q-doc	0.1330\triangle	0.0426	0.1339
GenEx	0.0617\triangle	0.0055	0.0614\triangle	GenEx	0.1179 \triangle	0.0479	0.1179 \triangle

For BERT-based ranker, a given input sample which is a query-document pair in the setting of our task, is perturbed by masking a token from a random position in the document. The final score of each document token is calculated by summing the scores obtained for all occurrences of the token in the document. Document tokens are sorted by the obtained scores and the top-ranked tokens are considered as the explanation of the document. The top-ranked tokens are considered to be tokens whose scores are not less than 10% of the token with the highest score. We consider the top three tokens if more than 3 are selected by the thresholding function. The cut-off value is chosen based on the average length of gold explanations in the test sets. Note that the average length of gold explanations is not used by the GenEx model.

The last baseline is training the original Transformer model where input sequences are constructed by concatenating query and document tokens in the training dataset as Eq. (1). The model input is constructed by adding segment embeddings differentiating query and document segments in addition to the separator token σ between them. These segment embeddings have the value of 0 for query tokens and 1 for all document tokens. We refer to this model as Orig-Trans.

Ablation study. There are two main architectural choices in the proposed explanation model; query attention layer, and query-masking during the decoding. To show the effectiveness of each choice, we compare our final model with the following variants, gradually constructed on top of

Table 5. Performance of generative explanation models based on BLEU metrics. Symbols \blacktriangle and \triangle show statistical significant differences with Orig-Trans at levels 0.01 and 0.05, respectively.

	Wiki		Clue-Res		Passex	
	BLEU-1	BLEU-2	BLEU-1	BLEU-2	BLEU-1	BLEU-2
Orig-Trans	0.2269	0.1357	0.0301	0.0056	0.0755	0.0468
Seg-q-toks	0.2918 \blacktriangle	0.1791 \blacktriangle	0.0385	0.0108	0.1038	0.0595
Sep-q-doc	0.2979\blacktriangle	0.1845\blacktriangle	0.0473 \triangle	0.0156\blacktriangle	0.1212 \blacktriangle	0.0651
GenEx	0.2313 \blacktriangle	0.1236	0.0520\blacktriangle	0.0099 \blacktriangle	0.1264\blacktriangle	0.0718

Table 6. Performance of generative explanation models based on F1-score of BERTScore. Statistical significant differences between the GenEx and Orig-Trans at the levels of 0.01 and 0.1 are shown with \blacktriangledown and \triangledown , respectively.

	Wiki	Clue-Res	Passex
Orig-Trans	0.4071 \blacktriangledown	0.2985 \blacktriangledown	0.3596 \triangledown
Seg-q-toks	0.4796	0.3172	0.3989
Sep-q-doc	0.4705	0.2898	0.3786
GenEx	0.4472	0.3303	0.4430

Table 7. Examples of explanations generated by GenEx, TextRank, and LIME-TF.IDF.

Query	diversity
Doc. title	Equality & Diversity: Athena Report and Action Plan
TextRank	“oxford”, “athena”, “women”
LIME+TF.IDF	diversity
GenEx	women s career development
Query	OCD
Doc. title	OCD Obsessive-Compulsive Disorder - Mahalo
TextRank	“ocd”, “disorder”, “help”
LIME+TF.IDF	ocd
GenEx	obsessive

the Orig-Trans model. The first variant uses separate encoders for queries and documents with an additional encoder for documents which consists of query attention layers. However, query tokens are not masked in the decoder. We refer to this model as Sep-q-doc. We also test another variant by segment embeddings that differentiate query tokens from other tokens in the input sequence, having 0 for query tokens in the input sequence and a value of 1 for other tokens. Therefore, query tokens in a document are also masked during the decoding process, and only the final hidden vectors corresponding to non-query tokens in a document are fed into the decoder. This is similar to the decoder input of the GenEx model. This variant is referred to as Seg-q-toks model.

5.4 Experimental Details

All baseline based on the Transformer architecture and GenEx are all built and trained using the same settings of hyperparameters and on the same training data to ensure a fair comparison between them. Hyperparameters are mainly set following the base Transformer model. Each

Table 8. Example documents from the Clue-Res dataset. Only the beginning of the first document is copied here.

-
- Doc. 1 university of oxford athena project 2000 1 action plan encouraging applications from women scientists summary in 1999 2000 the university of oxford applied successfully to the national athena project 1 for funding to assist with a programme of positive action aimed at encouraging applications from women scientists for academic appointments at the university . positive action with the objective of encouraging applications from an under represented sex is defined and authorised by the sex discrimination act 1975 . the university s application was based on an analysis of data from its recruitment monitoring scheme . this consistently demonstrates that women are appointed to academic posts , including those in science , engineering and technology set , at least in proportion to their applications , but that the rate of applications from women is low compared with numbers suggested by the available data to exist in recruitment pools such as contract research staff at oxford and elsewhere and lecturers at other institutions . the acceptance of athena funding commits the university to develop and carry out an action plan based on the experience of its project . this action plan pdf file , 12kb is annexed in tabular form and further information on the oxford athena project is provided below . although the oxford athena project , and therefore this report , deal with the scientific disciplines , there is evidence that women are similarly under represented at oxford in some areas of the humanities and social sciences and , where appropriate , the action plan is intended to cover all disciplines .
-
- Doc. 2 obsessive compulsive disorder ocd is an acronym for the mental disorder known specifically as obsessive compulsive disorder . ocd is a type of anxiety disorder . ocd can persist throughout a persons life . the symptoms of ocd can be mild to severe . if severe , they can interfere with a persons ability to function at work , school and home . symptoms ocd involves uncontrollable urges , rituals , or thoughts that cannot be put out of a persons mind . such behavior , may be all consuming , and eventually take over the person s life . a person with ocd feels the need to repeat the same thing over and over to keep bad things from happening . the symptoms of ocd consist of obsession , the constant idea that something bad is going to happen , and compulsion , the constant action to try to prevent the bad things . for example , someone that is worried about germs will wash their hands over and over . cause the precise cause of ocd is still not known . while some researchers believe it is a chemical imbalance , others see it as a physical condition .
-

encoder/decoder module consists of a stack of 6 identical layers. We used 8 attention heads in each layer, each with depth of 96. The input dimension is 768. Deep contextualized representations of queries and documents can be obtained by fine-tuning BERT. However, our pilot experiments showed that we can obtain reasonable contextual representation using 6-layer Transformer encoder, mainly because these representations will be updated by the next encoder component. Therefore, considering computational constraint and having large amount of weakly labeled training data, we decided to use Transformer encoder with pre-trained input embeddings from BERT, instead of using the entire pre-trained BERT model. Training samples are batched by their total length. Each batch consists of samples with approximately 2048 tokens in total. The input sequence to encoder modules are truncated if required by keeping the first 256 tokens. The models are trained for 10

epochs. We used the Adam optimizer [26] with $\beta_1 = 0.9$, $\beta_2 = 0.98$, $\epsilon = 10^{-9}$, and a fixed learning rate of 0.00001. Dropout rate of 0.1 is used during the training of models. In addition, target outputs are smoothed with a value of 0.1. Explanations for test samples are generated by greedy decoding.

5.5 Results and Analysis

In this section, we provide the evaluation results of the proposed model. Performance of GenEx and baseline models on Wikipedia test fold (Wiki), Clue-Res, and Passex datasets in terms of ROUGE and BLEU metrics is shown in Table 3. As baseline models do not consider the order of words that are selected as explanation, we do not compare them with GenEx based on the performance metrics that depend on higher order n-grams. The difficulty of the task at hand, the limitations of the noisy automatically-generated training data, and imprecise labels of test data are the main reasons for low performance values. GenEx outperforms all baselines over all three test sets where the improvements are mostly substantial and statistically significant. The BLEU improvements of GenEx over the best performing baseline are 168%, 18%, and 296% over Wiki, Clue-Res, and Passex datasets respectively. In terms of the BLEU metric, the best performing baseline over Wiki and Passex is TextRank, and KeyBERT shows the highest performance over the Clue-Res test set. In terms of the ROUGE metric, the best performing baselines over all datasets is TextRank. The ROUGE improvements of GenEx over TextRank are 151%, 41%, and 34%.

The first group of baselines is keyword extraction by TextRank and its topic-sensitive variant. TextRank provides a query-independent explanation of a document, while topic-sensitive TextRank is developed to explain documents with respect to a given query. As a small part of a document may be relevant to the query [65] and in such cases we are not interested in the document keywords related to its general topic, the topic-sensitive variant of the TextRank algorithm should intuitively select document words that are more suitable for explanation of document relevance. The reported results in Table 3 show that these models are among the best performing baselines. However, the TextRank algorithms shows higher performance than its topic-sensitive variant based on BLEU and ROUGE metrics. Further investigation of TextRank and its topic-sensitive variant, we observed that topic-sensitive TextRank provides superior keywords when all query tokens occur in the document to be explained. When the document contains only one or a subset of query tokens, topic-sensitive PageRank leads to lower weights for document tokens related to non-present query tokens in the document compared to those by the original PageRank algorithm, and this hurts the quality of explanations for such cases by topic-sensitive TextRank.

LIME and Sensitivity methods underperform the GenEx model where the performance difference is considerable and statistically significant. These models extract document tokens that highly impact the relevance probability of a document with respect to a query. Consequently, tokens that are exact match or semantically similar to query tokens, get the highest importance scores. While this type of explanation show document relevance in terms of keyword matching, they mostly fail to provide words related to query aspects in their top keywords. These models thus provide different but complementary explanation than the GenEx model. LIME explanation of the TF.IDF ranker achieves better performance than that of the BERT-based ranker over the Wiki and Clue-Res test sets. The TF.IDF and BERT-based rankers are used to get the relevance probability of perturbed documents. The lower performance of LIME for the BERT-based ranker could be related to the complex structure of BERT, where the widely used explanation based on linear approximation is not accurate enough to predicate its function.

GenEx significantly outperforms KeyBERT that shows explanations of interest are beyond finding important noun phrases of a text segment.

Tables 4 and 5 report the performance of the ablations of the GenEx model. As the results show, all model variants almost always outperform the original transformer model. The obtained

Table 9. Comparing GenEx, TextRank, and LIME-TF.IDF explanations based on human evaluation.

\mathcal{K}_f	Majority prefer		$\frac{\# \text{individual prefer (GenEx)}}{\# \text{individual prefer (TextRank)}}$
	GenEx	TextRank	
0.50	73%	20%	3.3

\mathcal{K}_f	Majority prefer		$\frac{\# \text{individual prefer (GenEx)}}{\# \text{individual prefer (LIME)}}$
	GenEx	LIME	
0.37	50%	32%	2.3

improvements are also mostly statistically significant, with exceptional cases of the Bleu-2/R-2 performance. This observation demonstrates the effectiveness of adding *query-attention* layers and *masked-query* decoding.

Seg-q-toks and Sep-q-doc models show similar performance on test sets, however our analysis of their generated explanations reveals that the two model perform well on almost disjoint sets of query-document pairs. Comparing the generated explanations by these two models, the main difference seems to arise from query ambiguity. The Seg-q-toks model generates better explanations for ambiguous queries than the Sep-q-doc model. One example of ambiguous queries in the Clue-Res dataset is query 73 of TREC Web track 2010, “the sun”, with documents about the star in solar system, the U.K. newspaper, and the Baltimore Sun newspaper. In case of ambiguous queries, we believe that it would be helpful to encode queries using document tokens as their contexts. That is likely the main reason that the Seg-q-toks model outperforms the Sep-q-doc model for ambiguous queries.

Semantic evaluation of explanations. To semantically compare the generated explanations with respect to gold explanations, we use the BERTScore metric. This evaluation is necessary for abstractive text generation as BLEU and ROUGE metrics only consider the exact matching between generated and gold text segments, while a concept can be expressed in different ways. The default setting of BERTScore, “roberta-large_L17_no-idf_version = 0.3.2 (hug_trans = 2.8.0)-rescaled”, is used in our evaluations. The results are shown in Table 6. GenEx outperforms the original Transformer by 9.9%, 10.7%, and 23.2% over Wiki, Clue-Res, and Passex datasets, respectively, and all improvements are statistically significant. These results demonstrate the higher quality of generated explanations by GenEx compared to those by the original Transformer. As both models are trained on the same data, these improvements also demonstrate the suitability of GenEx architecture for the task in hand.

The semantic evaluation results over Clue-Res and Passex datasets are almost consistent with those of Rouge and BLEU metrics; GenEx greatly outperforms the original transformer and its variants, and is not the best performing model variant on the Wiki test set. Although GenEx is not the best performing variant on the Wiki test set, it still outperforms the original Transformer and the improvements are statistically significant. The difference between model variants could be related to the special structure of Wikipedia articles that is not the case for all pages in the Web. As the goal of our study is to explain a relevant document to a query, a model that is not dependent on the structure of document’s content is more desirable. Thus, GenEx is designed and trained to perform well on general documents in Clue-Res and Passex without being trained on them.

Explanations of keyword-based baselines are not evaluated using this metric as the BERT representations of input sequences in BERTScore are sensitive to input grammaticality and the order of words.

Table 10. Quality of GenEx explanations.

	Grammaticality	Relevance to query	Relevant to document
Avg. score	4.06	3.23	3.07
\mathcal{K}_f (binary)	0.49	0.42	0.51

5.6 Example Generated Explanations

Table 7 shows two examples of explanations generated by the GenEx model. Document contents are shown in Table 8. In addition to the generated explanations, the table shows the document titles and the top 3 keywords of documents obtained by the TextRank and LIME+TF.IDF models. Document titles show the general topic of the documents, which are not necessarily good explanations for queries that seek to find these documents. The GenEx model generates a reasonable explanation for the first example in the table. As this example shows query-focused explanation of the retrieved document is different than the document title which can be considered as general explanation of the document. The generated explanation for example 1 is more useful than the document title to reveal the relevancy of the document to the query. Note that document titles are not used as input of explanation generation models. Although it may be helpful to have document titles for contextual representation of documents, we do not want the model to rely on a feature that may not be available for some webpages and reduces the applicability of the model in practice.

Second example in Table 7 shows a failure of the GenEx model, where the generated explanation is part of the expansion of the given abbreviated query. A reasonable explanation for the second document with respect to the given query can be "symptoms". Desired explanations to be generated by a model should not repeat query terms as long as explanation readability is not sacrificed. This is the reason that masked query decoding is proposed in our GenEx model. However, expansion tokens of abbreviated queries in documents have high similarity values with query tokens, and they are not masked during the decoding since query masking is done based on exact matching of tokens. This input characteristic makes it highly probable that the decoder generates expansion tokens of abbreviations as explanations. Abbreviated queries thus constitute one category of failure cases of our GenEx model.

6 HUMAN EVALUATION

In Section 3 we found that people could use GENEX-style synthetic explanations to more accurately and more efficiently identify relevant documents. In Section 5 we evaluated the effectiveness of GENEX explanations using automatic evaluations. We now close the loop by exploring whether the automatically generated explanations serve to help people as the initial study suggested. The GenEx model as well as existing baseline models explain a single document with respect to a query, and thus the user studies in this section are designed according to the nature of models, and are different than the one conducted in Section 3. We leave the explanation of the entire search results for future work.

We randomly selected 25 query-document pairs from each of the Clue-Res, Passex, and Wiki test sets (Section 5). For Clue-Res, we had the additional requirements that the documents must have a maximum length of 600 tokens (for ease of human review). In this study¹⁰, performed on Amazon Mechanical Turk, we presented each worker with a query-document pair and asked questions addressing preference between explanations of different approaches, their linguistic quality, and their relevance to the query and document pair. The worker is shown two illustrative examples for

¹⁰Institutional Review Board number 1381.

orientation. Each query-document pair is evaluated by three different master workers located in the United States to reduce subjectivity.

6.1 Comparison with Baseline Models

For each query-document pair, workers are asked to answer whether: (1) GenEx explanation is better, (2) Explanation by a baseline (LIME-TF.IDF or TextRank) is better, (3) both are equally good, or (4) both are equally bad.

TextRank is mostly the best performing baseline in terms of BLEU and ROUGE metrics according to the results in Table 3. LIME is a successful and widely-used model for explanation of black-box models, which is also the state-of-the-art model in the explanation of black-box rankers [59]. LIME has shown a better performance in explaining the TF.IDF ranker in our experiments with results reported in Table 3. LIME does not use any information about the structure of the TF.IDF ranker, it only uses its outputs for the perturbed instances of the document to be explained. This baseline is chosen as GenEx explains why a document is relevant to a query without having any knowledge of the underlying ranker model.

Table 9 summarizes the obtained results which demonstrate that the majority of workers preferred explanations generated by GenEx over those by LIME-TF.IDF and TextRank in 18%-53% more samples, respectively. In addition, the ratios of individual workers who preferred GenEx explanations over LIME-TF.IDF and TextRank are 2.3 and 3.3, respectively. Results of human evaluation show the higher quality of explanations generated by GenEx compared to baseline models.

6.2 Explanations Quality

We also conducted a user study to capture human evaluation of quality of the GenEx explanations. Three different dimensions are considered for this evaluation: linguistic quality, relevance of explanation to the query topic, and relevance of explanation to the document content. The last two questions provide a proxy for the utility of generated explanations' ability to connect the query and document. We use the 5-point Likert Scale to evaluate the subjective quality of generated explanations, as Very poor, Poor, Acceptable, Good, and Very good, with assigned scores from 1 to 5, respectively. Because differences between the levels can be subtle, we calculated Fleiss' Kappa agreement by collapsing negative scores (1 and 2) to *no* and the others to *yes*. Table 10 summarizes the obtained results.

Linguistic quality. We evaluated the grammaticality and coherence of the explanations generated by the GenEx model. As shown in Table 10, we observed that the average of grammaticality scores is 4.06. Only 9% of all answers to the grammaticality questions were Poor or Very poor. Part of this strong result can be due to the short length of explanations, as all generated explanations are terse (up to four terms). Yet as explanations are generated in an abstractive way, this result strongly indicates the potential of our GenEx model in generation relevance explanations.

Relevance to query and document pairs. We asked workers if the generated explanations were relevant to the query topic and document content. Note that, following annotation guidelines, explanations that do not provide additional information compared to the query, such as explanations that are a subsequence of queries, should be rated as very poor for both questions. An explanation that is relevant to the topic of the query without repeating the query, provides additional useful information with respect to the query that can help users in understanding the information space. When an explanation is relevant to the document, it means that the explanation is describing the content of the document, and the model is not generating a general or a high-frequent phrase in the training data, which is a common issue of text generation models.

The obtained results in Table 10 show that on average, workers rated GenEx explanations have acceptable degree of relevance to both query topic without repeating the query and document

content. These results demonstrate that generated explanations can reasonably describe document content with respect to the query topic and provide more information than the given query.

We chose not to run this part of the study for LIME explanations, since by their nature, they are unlikely to be grammatical. On the other hand, extractive explanations by LIME are likely to be relevant to the document, and do not require the same evaluations as abstractive explanations by GenEx. Note that abstractive generation of texts at the level of noun-phrase has a lower risk of topic drift compared to abstractive snippet generation which has been motivated recently [7], and is confirmed by the results of human evaluation on relevance to document content.

7 CONCLUSIONS AND FUTURE WORK

We studied how a retrieved document can be explained with respect to the given query. We proposed a Transformer-based architecture with *query attention* layers and *masked-query* decoding, called GenEx. The proposed solution is not trained using any manually labeled training data. Comprehensive evaluation of GenEx demonstrated its superior performance.

We believe that this work opens up new directions towards explainable document retrieval. In the recent search scenarios with limited bandwidth interfaces, such as conversational search systems using speech-only or small-screen devices [45], presenting result lists with long snippets is not plausible, emphasizing the need for a new form of explanation conforming to their characteristics. We intend to incorporate explanation into conversational search systems with limited bandwidth interfaces. Another interesting direction to pursue is to design an end-to-end explanation model that can handle documents of any length. Given the memory constraints of current hardware, GenEx works on query-biased extraction of long documents which may not be optimal. We also would like to extend GenEx to make it more robust with respect to diverse types of queries in Web search. Furthermore, the proposed solution generates an explanation for a query-document pair. Future work can explore document-level explanation based on the top results. Incorporating the generated explanations on a web search interface in order to improve search experience for users is another interesting future direction.

ACKNOWLEDGMENTS

This work was supported in part by the Center for Intelligent Information Retrieval, in part by Amazon.com, in part by NSF grant number 1813662 and in part by NSF grant #IIS-2039449. Any opinions, findings and conclusions or recommendations expressed in this material are those of the authors and do not necessarily reflect those of the sponsor.

REFERENCES

- [1] Rakesh Agrawal, Sreenivas Gollapudi, Alan Halverson, and Samuel Ieong. 2009. Diversifying Search Results. In *Proceedings of the Second ACM International Conference on Web Search and Data Mining (WSDM '09)*. 5–14.
- [2] Qingyao Ai, Yongfeng Zhang, Keping Bi, and W. Bruce Croft. 2019. Explainable Product Search with a Dynamic Relation Embedding Model. *ACM Trans. Inf. Syst.* 38, 1, Article 4 (Oct. 2019), 29 pages. <https://doi.org/10.1145/3361738>
- [3] Mikel Artetxe, Gorka Labaka, and Eneko Agirre. 2019. An Effective Approach to Unsupervised Machine Translation. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*. Association for Computational Linguistics, 194–203.
- [4] Mikel Artetxe, Gorka Labaka, Eneko Agirre, and Kyunghyun Cho. 2018. Unsupervised Neural Machine Translation. In *International Conference on Learning Representations*.
- [5] Hannah Bast and Marjan Celikic. 2014. Efficient Index-Based Snippet Generation. *ACM Trans. Inf. Syst.* 32, 2, Article 6 (April 2014), 24 pages. <https://doi.org/10.1145/2590972>
- [6] Gabriele Capannini, Franco Maria Nardini, Raffaele Perego, and Fabrizio Silvestri. 2011. Efficient Diversification of Web Search Results. *Proc. VLDB Endow.* 4, 7 (April 2011), 451–459.
- [7] Wei-Fan Chen, Shahbaz Syed, Benno Stein, Matthias Hagen, and Martin Potthast. 2020. Abstractive Snippet Generation. In *Proceedings of The Web Conference 2020 (Taipei, Taiwan) (WWW '20)*. Association for Computing Machinery, New

- York, NY, USA, 1309–1319. <https://doi.org/10.1145/3366423.3380206>
- [8] Sumit Chopra, Michael Auli, and Alexander M. Rush. 2016. Abstractive Sentence Summarization with Attentive Recurrent Neural Networks. In *Proceedings of the 2016 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*. Association for Computational Linguistics, San Diego, California, 93–98.
- [9] Eric Chu and Peter Liu. 2019. MeanSum: A Neural Model for Unsupervised Multi-Document Abstractive Summarization. In *Proceedings of the 36th International Conference on Machine Learning (Proceedings of Machine Learning Research, Vol. 97)*. 1223–1232.
- [10] Kevyn Collins-Thompson, Paul Bennett, Fernando Diaz, Charles L. A. Clarke, and Ellen M. Voorhees. 2014. TREC 2013 Web Track Overview. In *Proceedings of the 22nd Text REtrieval Conference (TREC 2013)*.
- [11] Kevyn Collins-Thompson, Craig Macdonald, Paul Bennett, Fernando Diaz, and Ellen M. Voorhees. 2015. Overview of the TREC 2014 Web Track. In *In Proceedings of the 23rd Text REtrieval Conference (TREC '14)*.
- [12] Nick Craswell, Bhaskar Mitra, Emine Yilmaz, Daniel Campos, and Ellen M Voorhees. 2020. Overview of the trec 2019 deep learning track. *arXiv preprint arXiv:2003.07820* (2020).
- [13] Van Dang and Bruce W. Croft. 2010. Query Reformulation Using Anchor Text. In *WSDM '10*. 41–50.
- [14] Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding. In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*. Association for Computational Linguistics, 4171–4186.
- [15] Li Dong, Nan Yang, Wenhui Wang, Furu Wei, Xiaodong Liu, Yu Wang, Jianfeng Gao, Ming Zhou, and Hsiao-Wuen Hon. 2019. Unified Language Model Pre-training for Natural Language Understanding and Generation. *arXiv preprint arXiv:1905.03197* (2019).
- [16] Finale Doshi-Velez and Been Kim. 2017. Towards a rigorous science of interpretable machine learning. *arXiv preprint arXiv:1702.08608* (2017).
- [17] Zeon Trevor Fernando, Jaspreet Singh, and Avishek Anand. 2019. A Study on the Interpretability of Neural Retrieval Models Using DeepSHAP. In *Proceedings of the 42nd International ACM SIGIR Conference on Research and Development in Information Retrieval (Paris, France) (SIGIR'19)*. 1005–1008.
- [18] Luyu Gao, Zhuyun Dai, and Jamie Callan. 2020. Modularized Transformer-based Ranking Framework. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*. Association for Computational Linguistics, Online, 4180–4190.
- [19] Maarten Grootendorst. 2020. KeyBERT: Minimal keyword extraction with BERT. <https://doi.org/10.5281/zenodo.4461265>
- [20] Riccardo Guidotti, Anna Monreale, Salvatore Ruggieri, Franco Turini, Fosca Giannotti, and Dino Pedreschi. 2018. A Survey of Methods for Explaining Black Box Models. *ACM Comput. Surv.* 51, 5, Article 93 (Aug. 2018), 42 pages.
- [21] Taher H. Haveliwala. 2002. Topic-Sensitive PageRank (*WWW '02*). 517–526.
- [22] Yunhua Hu, Yanan Qian, Hang Li, Daxin Jiang, Jian Pei, and Qinghua Zheng. 2012. Mining Query Subtopics from Search Log Data. In *SIGIR '12*. 305–314.
- [23] Yu Huang, Ziyang Liu, and Yi Chen. 2008. Query Biased Snippet Generation in XML Search. In *Proceedings of the 2008 ACM SIGMOD International Conference on Management of Data (Vancouver, Canada) (SIGMOD '08)*. Association for Computing Machinery, New York, NY, USA, 315–326. <https://doi.org/10.1145/1376616.1376651>
- [24] Mayu Iwata, Tetsuya Sakai, Takehiro Yamamoto, Yu Chen, Yi Liu, Ji-Rong Wen, and Shojiro Nishio. 2012. AspectTiles: Tile-Based Visualization of Diversified Web Search Results (*SIGIR '12*). 85–94.
- [25] Omar Khattab and Matei Zaharia. 2020. ColBERT: Efficient and Effective Passage Search via Contextualized Late Interaction over BERT. In *Proceedings of the 43rd International ACM SIGIR Conference on Research and Development in Information Retrieval (SIGIR '20)*. 39–48.
- [26] Diederik Kingma and Jimmy Ba. 2014. Adam: A Method for Stochastic Optimization. *International Conference on Learning Representations* (12 2014).
- [27] Guillaume Lample and Alexis Conneau. 2019. Cross-lingual language model pretraining. *arXiv preprint arXiv:1901.07291* (2019).
- [28] Mike Lewis, Yinhan Liu, Naman Goyal, Marjan Ghazvininejad, Abdelrahman Mohamed, Omer Levy, Ves Stoyanov, and Luke Zettlemoyer. 2019. BART: Denoising Sequence-to-Sequence Pre-training for Natural Language Generation, Translation, and Comprehension. *arXiv preprint arXiv:1910.13461* (2019).
- [29] Chin-Yew Lin. 2004. ROUGE: A Package for Automatic Evaluation of Summaries. In *Text Summarization Branches Out*. Association for Computational Linguistics, Barcelona, Spain, 74–81.
- [30] Zachary C. Lipton. 2018. The Mythos of Model Interpretability. *Queue* 16, 3, Article 30 (June 2018), 27 pages.
- [31] Fei Liu, Jeffrey Flanigan, Sam Thomson, Norman Sadeh, and Noah A. Smith. 2015. Toward Abstractive Summarization Using Semantic Representations. In *Proceedings of the 2015 Conference of the North American Chapter of the Association*

- for *Computational Linguistics: Human Language Technologies*. Association for Computational Linguistics, 1077–1086.
- [32] Yang Liu and Mirella Lapata. 2019. Text summarization with pretrained encoders. *arXiv preprint arXiv:1908.08345* (2019).
- [33] Scott M. Lundberg and Su-In Lee. 2017. A Unified Approach to Interpreting Model Predictions. In *Proceedings of the 31st International Conference on Neural Information Processing Systems (NIPS'17)*. 4768–4777.
- [34] Sean MacAvaney, Franco Maria Nardini, Raffaele Perego, Nicola Tonello, Nazli Goharian, and Ophir Frieder. 2020. Efficient Document Re-Ranking for Transformers by Precomputing Term Representations. In *Proceedings of the 43rd International ACM SIGIR Conference on Research and Development in Information Retrieval (SIGIR '20)*. 49–58.
- [35] Christopher D. Manning, Prabhakar Raghavan, and Hinrich Schütze. 2008. *Introduction to Information Retrieval*. Cambridge University Press.
- [36] Siyu Mi and Jiepu Jiang. 2019. Understanding the Interpretability of Search Result Summaries. In *Proceedings of the 42nd International ACM SIGIR Conference on Research and Development in Information Retrieval (SIGIR'19)*. 989–992.
- [37] Rada Mihalcea and Paul Tarau. 2004. TextRank: Bringing Order into Text. In *Proceedings of the 2004 Conference on Empirical Methods in Natural Language Processing*. Association for Computational Linguistics, Barcelona, Spain, 404–411. <https://www.aclweb.org/anthology/W04-3252>
- [38] Tomas Mikolov, Ilya Sutskever, Kai Chen, Greg Corrado, and Jeffrey Dean. 2013. Distributed Representations of Words and Phrases and Their Compositionality. In *Proceedings of the 26th International Conference on Neural Information Processing Systems - Volume 2 (NIPS'13)*. 3111–3119.
- [39] Tim Miller. 2019. Explanation in artificial intelligence: Insights from the social sciences. *Artificial Intelligence* 267 (2019), 1 – 38.
- [40] Ramesh Nallapati, Bowen Zhou, Cicero dos Santos, Çağlar Gülçehre, and Bing Xiang. 2016. Abstractive Text Summarization using Sequence-to-sequence RNNs and Beyond. In *Proceedings of The 20th SIGNLL Conference on Computational Natural Language Learning*. Association for Computational Linguistics, Berlin, Germany, 280–290. <https://doi.org/10.18653/v1/K16-1028>
- [41] Rodrigo Nogueira and Kyunghyun Cho. 2019. Passage Re-ranking with BERT. *arXiv preprint arXiv:1901.04085* (2019).
- [42] Kishore Papineni, Salim Roukos, Todd Ward, and Wei-Jing Zhu. 2002. BLEU: A Method for Automatic Evaluation of Machine Translation. In *Proceedings of the 40th Annual Meeting on Association for Computational Linguistics (ACL '02)*. 311–318.
- [43] Jeffrey Pennington, Richard Socher, and Christopher Manning. 2014. Glove: Global Vectors for Word Representation. In *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing (EMNLP)*. Association for Computational Linguistics, 1532–1543.
- [44] Matthew Peters, Mark Neumann, Mohit Iyyer, Matt Gardner, Christopher Clark, Kenton Lee, and Luke Zettlemoyer. 2018. Deep Contextualized Word Representations. In *NAACL*. 2227–2237.
- [45] Filip Radlinski and Nick Craswell. 2017. A Theoretical Framework for Conversational Search. In *CHIIR '17*. 117–126.
- [46] Filip Radlinski and Susan Dumais. 2006. Improving Personalized Web Search Using Result Diversification. In *Proceedings of the 29th Annual International ACM SIGIR Conference on Research and Development in Information Retrieval (SIGIR '06)*. Association for Computing Machinery, New York, NY, USA, 691–692.
- [47] Shuo Ren, Yu Wu, Shujie Liu, Ming Zhou, and Shuai Ma. 2019. Explicit Cross-lingual Pre-training for Unsupervised Machine Translation. In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*. Association for Computational Linguistics, 770–779.
- [48] Marco Tulio Ribeiro, Sameer Singh, and Carlos Guestrin. 2016. "Why Should I Trust You?": Explaining the Predictions of Any Classifier. In *KDD '16 (San Francisco, California, USA)*. 1135–1144.
- [49] Marco Tulio Ribeiro, Sameer Singh, and Carlos Guestrin. 2018. Anchors: High-Precision Model-Agnostic Explanations.
- [50] Gaetano Rossiello, Pierpaolo Basile, and Giovanni Semeraro. 2017. Centroid-based Text Summarization through Compositionality of Word Embeddings. In *Proceedings of the MultiLing 2017 Workshop on Summarization and Summary Evaluation Across Source Types and Genres*. Association for Computational Linguistics, 12–21.
- [51] Dwaipayan Roy, Sourav Saha, Mandar Mitra, Bihan Sen, and Debasis Ganguly. 2019. I-REX: A Lucene Plugin for EXplainable IR. In *Proceedings of the 28th ACM International Conference on Information and Knowledge Management (Beijing, China) (CIKM '19)*. ACM, 2949–2952.
- [52] Tuukka Ruotsalo, Jaakko Peltonen, Manuel J. A. Eugster, Dorota Glowacka, Patrik Floréen, Petri Myllymäki, Giulio Jacucci, and Samuel Kaski. 2018. Interactive Intent Modeling for Exploratory Search. *ACM Trans. Inf. Syst.* 36, 4, Article 44 (Oct. 2018), 46 pages.
- [53] Alexander M. Rush, Sumit Chopra, and Jason Weston. 2015. A Neural Attention Model for Abstractive Sentence Summarization. In *Proceedings of the 2015 Conference on Empirical Methods in Natural Language Processing*. Association for Computational Linguistics, Lisbon, Portugal, 379–389.

- [54] Rodrygo L. T. Santos, Craig Macdonald, and Iadh Ounis. 2015. Search Result Diversification. *Found. Trends Inf. Retr.* 9, 1 (March 2015), 1–90.
- [55] Hubert J. A. Schouten. 1986. Nominal scale agreement among observers. *Psychometrika* 51, 3 (1986), 453–466.
- [56] Procheta Sen, Debasis Ganguly, Manisha Verma, and Gareth J.F. Jones. 2020. *The Curious Case of IR Explainability: Explaining Document Scores within and across Ranking Models*. 2069–2072.
- [57] Jaspreet Singh and Avishek Anand. 2018. Interpreting search result rankings through intent modeling. *arXiv preprint arXiv:1809.05190* (2018).
- [58] Jaspreet Singh and Avishek Anand. 2018. Posthoc Interpretability of Learning to Rank Models using Secondary Training Data. In *Workshop on Explainable Recommendation and Search (EARS 2018) at SIGIR 2018*.
- [59] Jaspreet Singh and Avishek Anand. 2019. EXS: Explainable Search Using Local Model Agnostic Interpretability. In *WSDM '19 (Melbourne VIC, Australia)*. 770–773.
- [60] Jaspreet Singh, Megha Khosla, and Avishek Anand. 2021. Valid Explanations for Learning to Rank Models. In *Proceedings of the 2021 ACM SIGIR on International Conference on Theory of Information Retrieval (ICTIR '21)*.
- [61] Kaitao Song, Xu Tan, Tao Qin, Jianfeng Lu, and Tie-Yan Liu. 2019. MASS: Masked Sequence to Sequence Pre-training for Language Generation. In *ICML (Proceedings of Machine Learning Research, Vol. 97)*. PMLR, Long Beach, California, USA, 5926–5936.
- [62] Paul Thomas, Bodo Billerbeck, Nick Craswell, and Ryen W. White. 2019. Investigating Searchers' Mental Models to Inform Search Explanations. *ACM Trans. Inf. Syst.* 38, 1, Article 10 (Dec. 2019), 25 pages. <https://doi.org/10.1145/3371390>
- [63] Nava Tintarev and Judith Masthoff. 2015. *Explaining Recommendations: Design and Evaluation*. Springer US, Boston, MA, 353–382.
- [64] Anastasios Tombros and Mark Sanderson. 1998. Advantages of Query Biased Summaries in Information Retrieval. In *SIGIR '98*. 2–10.
- [65] TREC. 2000. Text REtrieval Conference (TREC) Data - English Relevance Judgements. https://trec.nist.gov/data/rejudge_eng.html.
- [66] Andrew Turpin, Yohannes Tsegay, David Hawking, and Hugh E. Williams. 2007. Fast Generation of Result Snippets in Web Search. In *SIGIR '07 (Amsterdam, The Netherlands)*. 127–134.
- [67] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N. Gomez, Lukasz Kaiser, and Illia Polosukhin. 2017. Attention is All You Need. In *Proceedings of the 31st International Conference on Neural Information Processing Systems* (Long Beach, California, USA). 6000–6010.
- [68] Manisha Verma and Debasis Ganguly. 2019. LIRME: Locally Interpretable Ranking Model Explanation. In *SIGIR'19*. 1281–1284.
- [69] Manisha Verma and Debasis Ganguly. 2019. LIRME: Locally Interpretable Ranking Model Explanation. In *Proceedings of the 42nd International ACM SIGIR Conference on Research and Development in Information Retrieval (Paris, France) (SIGIR'19)*. 1281–1284.
- [70] Alexandre A Verstak and Anurag Acharya. 2012. Generation of document snippets based on queries and search results. US Patent 8,145,617.
- [71] Nikos Voskarides, Edgar Meij, and Maarten de Rijke. 2017. Generating Descriptions of Entity Relationships. In *Advances in Information Retrieval*. Springer International Publishing, 317–330.
- [72] Nikos Voskarides, Edgar Meij, Manos Tsagkias, Maarten de Rijke, and Wouter Weerkamp. 2015. Learning to Explain Entity Relationships in Knowledge Graphs. In *ACL*. 564–574.
- [73] Qinglei Wang, Yanan Qian, Ruihua Song, Zhicheng Dou, Fan Zhang, Tetsuya Sakai, and Qinghua Zheng. 2013. Mining Subtopics from Text Fragments for a Web Query. *Inf. Retr.* 16, 4 (Aug. 2013), 484–503.
- [74] Xiaxia Wang, Jinchi Chen, Shuxin Li, Gong Cheng, Jeff Z. Pan, Evgeny Kharlamov, and Yuzhong Qu. 2019. A Framework for Evaluating Snippet Generation for Dataset Search. In *The Semantic Web – ISWC 2019*. Springer International Publishing, 680–697.
- [75] Xiaxia Wang, Jinchi Chen, Shuxin Li, Gong Cheng, Jeff Z. Pan, Evgeny Kharlamov, and Yuzhong Qu. 2019. A Framework for Evaluating Snippet Generation for Dataset Search. In *The Semantic Web – ISWC 2019*. 680–697.
- [76] Yonghui Wu, Mike Schuster, Zhifeng Chen, Quoc V Le, Mohammad Norouzi, Wolfgang Macherey, Maxim Krikun, Yuan Cao, Qin Gao, Klaus Macherey, et al. 2016. Google's neural machine translation system: Bridging the gap between human and machine translation. *arXiv preprint arXiv:1609.08144* (2016).
- [77] Matthew D. Zeiler and Rob Fergus. 2014. Visualizing and Understanding Convolutional Networks. In *Computer Vision – ECCV 2014*, David Fleet, Tomas Pajdla, Bernt Schiele, and Tinne Tuytelaars (Eds.). 818–833.
- [78] Haoyu Zhang, Jingjing Cai, Jianjun Xu, and Ji Wang. 2019. Pretraining-Based Natural Language Generation for Text Summarization. In *CoNLL*. 789–797.
- [79] Tianyi Zhang*, Varsha Kishore*, Felix Wu*, Kilian Q. Weinberger, and Yoav Artzi. 2020. BERTScore: Evaluating Text Generation with BERT. In *International Conference on Learning Representations*. <https://openreview.net/forum?id=SkeHuCVFDr>

- [80] Yongfeng Zhang and Xu Chen. 2018. Explainable recommendation: A survey and new perspectives. *arXiv preprint arXiv:1804.11192* (2018).