

Query Understanding via Intent Description Generation

Ruqing Zhang, Jiafeng Guo, Yixing Fan, Yanyan Lan, and Xueqi Cheng
CAS Key Lab of Network Data Science and Technology, Institute of Computing Technology,
Chinese Academy of Sciences, Beijing, China
University of Chinese Academy of Sciences, Beijing, China
{zhangruqing, guojiafeng, fanyixing, lanyanyan, cxq}@ict.ac.cn

ABSTRACT

Query understanding is a fundamental problem in information retrieval (IR), which has attracted continuous attention through the past decades. Many different tasks have been proposed for understanding users' search queries, e.g., query classification or query clustering. However, it is not that precise to understand a search query at the intent class/cluster level due to the loss of many detailed information. As we may find in many benchmark datasets, e.g., TREC and SemEval, queries are often associated with a detailed description provided by human annotators which clearly describes its intent to help evaluate the relevance of the documents. If a system could automatically generate a detailed and precise intent description for a search query, like human annotators, that would indicate much better query understanding has been achieved. In this paper, therefore, we propose a novel Query-to-Intent-Description (Q2ID) task for query understanding. Unlike those existing ranking tasks which leverage the query and its description to compute the relevance of documents, Q2ID is a reverse task which aims to generate a natural language intent description based on both relevant and irrelevant documents of a given query. To address this new task, we propose a novel Contrastive Generation model, namely CtrsGen for short, to generate the intent description by contrasting the relevant documents with the irrelevant documents given a query. We demonstrate the effectiveness of our model by comparing with several state-of-the-art generation models on the Q2ID task. We discuss the potential usage of such Q2ID technique through an example application.

CCS CONCEPTS

• Information systems → Query intent;

KEYWORDS

Query understanding, Intent Description Generation, Contrastive

ACM Reference Format:

Ruqing Zhang, Jiafeng Guo, Yixing Fan, Yanyan Lan, and Xueqi Cheng. 2020. Query Understanding via Intent Description Generation. In *Proceedings of the 29th ACM International Conference on Information and Knowledge Management (CIKM '20)*, October 19–23, 2020, Virtual Event, Ireland. ACM, New York, NY, USA, 10 pages. <https://doi.org/10.1145/3340531.3411999>

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.
CIKM '20, October 19–23, 2020, Virtual Event, Ireland

© 2020 Association for Computing Machinery.
ACM ISBN 978-1-4503-6859-9/20/10...\$15.00
<https://doi.org/10.1145/3340531.3411999>

Table 1: An example from the TREC 2004 Robust dataset.

Query (Number: 400): Amazon rain forest
Description: What measures are being taken by local South American authorities to preserve the Amazon tropical rain forest?
Document 1 (DOCNO: LA050789-0087): Eight nations of South America's Amazon basin called on wealthy countries to provide money for the preservation of the world's greatest rain forest and for the economic development of the region. At the first summit meeting on the Amazon ... (Relevance: 1)
Document 2 (DOCNO: LA060890-0004): Destruction of the world's tropical forests is occurring nearly 50% faster than the best previous scientific estimates showed ... Tree burning accounts for an estimated 30% of worldwide total carbon dioxide emissions ... (Relevance: 1)
Document 3 (DOCNO: LA062589-0034): For Beverly Revness and Janice Tarr ... Scientists also caution that burning the trees amounts to a double-barreled contribution to global warming the so-called "greenhouse effect." Combustion adds pollutants and carbon dioxide to the atmosphere ... (Relevance: 0)

1 INTRODUCTION

Query understanding is a key issue in information retrieval (IR), which aims to predict the search intent given a search query. Understanding the intent behind a query offers several advantages: 1) achieving better performance on query reformulation, query refinement, query prediction and query suggestion, and 2) improving the document relevance modeling based on the search intent of the query.

There have been many related research topics dedicated to query understanding over the past decades. Early works include a huge amount of human analysis and effort to identify the intent of a search query [53]. Later, many automated query intent analysis, such as query classification and query clustering, have been proposed to understand users' information needs. Query classification [6, 7, 27] aims to classify queries into one or more pre-defined target categories depending on different types of taxonomies. However, most existing research on query classification has focused on the coarse-grained understanding of a search query at the intent category level, which may result in the loss of many detailed information. Query clustering [4, 25, 64] attempts to discover search topics/intent by mining clusters of queries. Nevertheless, it is a little difficult for a human to clearly understand what each cluster represents. Hence, it is not that precise to understand a search query at the intent class/cluster level.

As we may find in many relevance ranking benchmark datasets (e.g., TREC¹ and SemEval²), queries are often associated with a detailed description provided by human annotators which clearly

¹<https://trec.nist.gov/>

²<https://en.wikipedia.org/wiki/SemEval>

describes its intent. As shown in Table 1, given a short query “Amazon rain forest” from the TREC 2004 Robust dataset, the description precisely clarifies its search intent “what are the measures to preserve the Amazon tropical rain forest”. Based on the detailed description, the relevance of the documents to a query can be well evaluated (i.e., relevance score 1 denotes *relevant* and relevance score 0 denotes *irrelevant*). From this example, we can find that the intent description is a more accurate and informative representation of the query compared with the intent class/cluster proposed in previous works. If a system could automatically generate a detailed and precise intent description for a search query, like human annotators did, that would indicate much better query understanding has been achieved.

In this paper, we thus introduce a novel Query-to-Intent-Description (Q2ID) task for query understanding. Given a query associated with a set of relevant and irrelevant documents, the Q2ID task aims to generate a natural language intent description, which interprets the information need of the query. The Q2ID task can be viewed as a reverse task of those existing ranking tasks. Specifically, the Q2ID task generates a description based on both relevant and irrelevant documents of a given query, while existing ranking tasks explicitly utilize the query and its description to compute the relevance of documents. Note the Q2ID task is quite different from traditional query-based multi-document summarization tasks which typically do not consider the irrelevant documents. To facilitate the study and evaluation of the Q2ID task, we build a benchmark dataset³ based on several public IR collections, i.e., Dynamic Domain Track in TREC, Robust Track in TREC and Task 3 in SemEval.

To address this new task, we introduce a novel Contrastive Generation model, named CtrsGen for short, to generate the intent description by contrasting the relevant documents with the irrelevant documents of a given query. The key idea is that a good intent description should be able to distinguish relevant documents from those irrelevant ones. Specifically, our CtrsGen model employs the Seq2Seq framework [2], which has achieved tremendous success in natural text generation. In the encoding phase, rather than treat each sentence in relevant documents equally, we introduce a *query-aware encoder attention mechanism* to identify those important sentences that can reveal the essential topics of the query. In the decoding phase, we employ a *contrast-based decoder attention mechanism* to adjust the importance of sentences in relevant documents with respect to their similarity to irrelevant documents. In this way, the CtrsGen model could identify those most distinguishing topics based on contrast.

Empirical results on the Q2ID benchmark dataset demonstrate that intent description generation for query understanding is feasible and our proposed method can outperform all the baselines significantly. We provide detailed analysis on the proposed model, and conduct case studies to gain better understanding on the learned description. Moreover, we discuss the potential usage of such Q2ID technique through an example application.

2 RELATED WORK

To the best of our knowledge, the Query-to-Intent-Description task is a new task in the IR community. In this section, we briefly review

three lines of related work, i.e., query understanding, summarization and interpretability.

2.1 Query Understanding

The most closely related query understanding tasks include query classification, query clustering, and query expansion.

2.1.1 Query Classification. Query classification has long been studied for understanding the search intent behind the queries which aims to classify search queries into pre-defined categories. Query classification is quite different from traditional text classification since queries are usually short (e.g., 2-3 terms) and ambiguous [27]. Initial studies focus on the type of the queries based on the information needed by the user and various taxonomies of search intent have been proposed. [6] firstly classified queries according to their intent into 3 types, i.e., Navigational, Informational and Transactional. This trichotomy is the most widely adopted one in automatic query intent classification work probably due to its simplicity and essence. Later, many taxonomies based on [6] are established [1, 49]. Since the query is often incomplete and indirect, and the intent is highly subjective, many works are proposed to consider the topic taxonomy for queries. Typical topic taxonomies used in literature include that proposed in the KDD Cup 2005 [35], and a manually collected one from AOL [5]. Different methods have been leveraged for this task. Some works are focusing to tackle the training data sparsity by introducing an intermediate taxonomy for mapping [29, 52], while some mainly consider the difficulty in representing the short and ambiguous query [5]. Beyond the previous major classification tasks on search queries, there has been research work paying attention to other “dimensions”, e.g., information type [45], time requirement [28] and geographical location [23].

2.1.2 Query Clustering. Query clustering aims to group similar queries together to understand the user’s search intent. The very early query clustering technique comes from information retrieval studies. The key issue is that how to measure the similarity of queries. Similarity functions such as cosine similarity or Jaccard similarity were used to measure the distance between two queries. Since queries are quite short and ambiguous, these measures suffer from the sparse nature of queries. To address this problem, click-through query logs have been mined to yield similar queries. [4] first introduced the agglomerative clustering method to discover similar queries using clicked URLs and query logs. [64] analyzed both query contents and click-through bipartite graph, and then applied DBSCAN algorithm to group similar queries. [67] presented a method that groups similar queries by analyzing users’ sequential search behavior. [47] combined click and reformulation information to find intent clusters. [18] proposed the use of click patterns through a hierarchical clustering algorithm. [42] proposed to dynamically mine query intents from search query logs. [38] calculated query similarity using feature terms extracted from user clicked documents with the help of WordNet. [25] quantified query similarity based on the queries’ top ranked search results. Recently, [32] used word2vec to obtain query representations and then applies Divide Merge Clustering on top of it. [48] proposed a new document clustering prototype, where the information is periodically updated to cater to the distributed environment.

³The dataset is available at <https://github.com/daqingchong/Q2ID-benchmark-dataset>.

2.1.3 Query Expansion. Query expansion aims to reformulate the initial query by adding similar terms that help in retrieving more relevant results. It was first applied by [37] as a technique for literature indexing and searching in a mechanized library system. Early works [21, 63] expanded the initial query terms by analyzing the expansion features, e.g., lexical, semantic and syntactic relationships, from large knowledge resources, e.g., WordNet and ConceptNet. Later, some works recognized the expansion features over the whole corpus [56, 57, 68] to create co-relations between terms, while many works utilized user’s search logs for expanding original query [14, 15]. Besides, several works used both positive and negative relevance feedback for query expansion [30, 46]. Recently, natural language generation models are used to automatically expand queries. [44] designed a query expansion model on the basis of a neural encoder-decoder model for question answering. [34] proposed to introduce the conditional generative adversarial network framework to directly generate the related keywords from a given query.

2.2 Summarization

The most closely related summarization tasks include multi-document summarization and query-based multi-document summarization.

2.2.1 Multi-Document Summarization. Multi-document summarization aims to produce summaries from document clusters on the same topic. Traditional multi-document summarization models studied in the past are extractive in nature, which try to extract the most important sentences in the document and rearranging them into a new summary [8, 19, 39, 59]. Recently, with the emergence of neural network models for text generation, a vast majority of the literature on summarization is dedicated to abstractive methods which are largely in the single-document setting. Most methods for abstractive text summarization [12, 26, 50, 58] are based on the neural encoder-decoder architecture [2, 60]. Many recent studies have attempted to adapt encoder-decoder models trained on single-document summarization datasets to multi-document summarization [33, 66]. Recently, for assessing sentence redundancy, [11] introduced the improved similarity measure inspired by capsule networks and [20] incorporated MMR into a pointer-generator network for multi-document summarization.

2.2.2 Query-based Multi-document Summarization. Query-based multi-document summarization is the process of automatically generating natural summaries of text documents in the context of a given query. An early work for extractive query-based multi-document summarization is presented by [22], which ranked sentences using a weighted combination of statistical and linguistic features. [16] presented to extract sentences based on the language model, Bayesian model, and graphical model. [40] introduced the graph information to look for relevant sentences. [51] used the multi-modality manifold-ranking algorithm to extract topic-focused summary from multiple documents. Recently, some works employ the encoder-decoder framework to produce the query-based summaries. [24] trained a pointer-generator model, and [3] incorporated relevance into a neural seq2seq models for query-based abstractive summarization. [43] introduced a new diversity based attention mechanism to alleviate the problem of repeating phrases.

2.3 Interpretability

Interpretability, in machine learning has been studied under the ill-defined notion of “the ability to explain or to present in understandable terms to a human” [17]. Recently there have been some works on explaining results of a ranking model in IR. [54] utilized a posthoc model agnostic interpretability approach for generating explanations, which are used to answer interpretability questions specific to ranking. [62] explored several sampling methods in generating local explanations for a document scored with respect to a query by an IR model. [55] proposed a model-agnostic approach, which attempts to locally approximate a complex ranker by using a simple ranking model in the term space. Also, researchers have explored various approaches [9, 13, 65] towards explainable recommendation systems, which can not only provide users with the recommendation lists, but also intuitive explanations about why these items are recommended.

The Q2ID task introduced in our work is quite different from the above existing tasks. Firstly, query classification and query clustering focused on the coarse-grained understanding of queries at the intent class/cluster level, while our Q2ID task provides fine-grained understanding of queries by generating a detailed intent description. Secondly, query expansion adds additional similar terms into the initial query, while our Q2ID task generates new sentences as the intent description of a given query. Then, the Q2ID task generates the intent description based on the relevant and irrelevant documents of a given query, while there is no such consideration of irrelevant documents in multi-document summarization or query-based multi-document summarization. Finally, most previous explainable search/recommender systems aim to explain why a single document/product was considered relevant, while our Q2ID task aims to generate an intent description for a given query based on both the relevant and irrelevant documents.

3 PROBLEM FORMALIZATION

In this section, we introduce the Q2ID task, and describe the benchmark dataset in detail.

3.1 Task Description

Given a query associated with a set of relevant documents and irrelevant documents, the Q2ID task aims to generate a natural language intent description, which precisely interprets the search intent that can help distinguish the relevant documents from the irrelevant documents.

Formally, given a query $q = \{w_1^q, \dots, w_C^q\}$ with a sequence of C words, a set of M relevant documents $\mathcal{R} = \{D_1^r, \dots, D_M^r\}$ where each relevant document D_m^r is composed of T_m sentences, and a set of N irrelevant documents $\mathcal{I} = \{D_1^i, \dots, D_N^i\}$ where each irrelevant document D_n^i is composed of T_n sentences, the Q2ID task is to learn a mapping function $g(\cdot)$ to produce an intent description $y = \{w_1^y, \dots, w_Z^y\}$ which contains a sequence of Z words, i.e.,

$$g(q, \mathcal{R}, \mathcal{I}) = y, \quad (1)$$

where $M \geq 1$ and $N \geq 0$. Specifically, the collection of irrelevant documents could be empty while the relevant documents are necessary for intent description generation.

Table 2: Data statistics: #s denotes the number of sentences, #w denotes the number of words, #r denotes the number of relevant documents, and #i denotes the number of irrelevant documents.

Query	5358
Query: avg #w	4.7
Query: avg #r	10.8
Query: avg #i	65.5
Relevant documents	62,916
Irrelevant documents	196,787
Relevant documents: avg #s	20.7
Irrelevant documents: avg #s	23.2
Intent Description: avg #w	31.0

3.2 Data Construction

In order to study and evaluate the Q2ID task, we build a benchmark dataset based on the public TREC and SemEval collections.

- **TREC** is an ongoing series of workshops focusing on a list of different IR research areas. We utilize the TREC 2015, 2016 and 2017 Dynamic Domain Track⁴ and TREC 2004 Robust Track⁵.
- **SemEval** is an ongoing series of evaluations of computational semantic analysis systems. We utilize the SemEval-2015 Task 3⁶ and SemEval-2016 Task 3⁷ for English⁸.

In these IR collections, queries are associated with a human-written detailed description, and documents are annotated with multi-graded relevance labels indicating a varying degree of match with the query intent/information. As a primary study on the intent description generation, here, we convert multi-graded relevance labels into binary relevance labels, indicating a document is relevant or irrelevant to a query. We leave the Q2ID task over multi-graded relevance labels as our future work.

With regard to different tasks, we define the binary relevance labels for documents as follows:

- For Dynamic Domain Track, each passage is graded at a scale of 0 ~ 4 according to the relevance to a query (i.e., subtopic). We treat passages with rating of 1 or higher as relevant documents, and passages with rating of 0 as irrelevant documents.
- For Robust Track, the judgment of each document is on a three-way scale of relevance to a query (i.e., topic). We treat documents annotated as *highly relevant* and *relevant* as relevant documents, and documents annotated as *not relevant* as irrelevant documents.
- For Task 3 in both SemEval-2015 and SemEval-2016, each answer is classified as *good*, *bad*, or *potential* according to the relevance to a query (i.e., question). We treat answers annotated as *good* as relevant documents, and the rest as irrelevant documents.

In this way, we obtain the <query, relevant documents, irrelevant documents, intent description> quadruples, as ground-truth data for training/validation/testing. Queries with no relevant documents are removed. Table 2 shows the overall statistics of our Q2ID benchmark dataset. Note there are 743 queries without corresponding irrelevant documents in our dataset.

⁴<https://trec.nist.gov/data/dd.html>

⁵<https://trec.nist.gov/data/robust.html>

⁶<http://alt.qcri.org/semeval2015/task3/>

⁷<http://alt.qcri.org/semeval2016/task3/>

⁸We do not consider the Arabic language and SemEval-2017 [41] task which reran the four subtasks from SemEval-2016.

4 OUR APPROACH

In this section, we introduce our proposed approach for the Q2ID task in detail. We first give an overview of the model architecture, and then describe each component of our model as well as the learning procedure.

4.1 Overview

Without loss of generality, the Q2ID task needs to distill the salient information from the relevant documents and remove unrelated information from the irrelevant documents with respect to a query. For example, as shown in Table 1, given the keyword query “Amazon rain forest”, there might be different underlying users’ search intents, like “location of Amazon rain forest”, “age of Amazon rain forest” or “protection of Amazon rain forest”. From the relevant document 1 and 2, we can find several topics, e.g., “the measures to preserve Amazon rain forest” and “the effect of tree burning”. From the irrelevant document 3, it mainly talks about the topic of “the effect of tree burning”. By contrasting the relevant documents with the irrelevant documents, we find that the query “Amazon rain forest” aims to search for “the measures to preserve Amazon rain forest” instead of “the effect of tree burning”. Therefore, in this work, we formulate the Q2ID task as a novel contrastive generation problem and introduce a CtrsGen model to solve it.

Basically, our CtrsGen model contains the following four components: 1) Query Encoder, to obtain the representation of a query; 2) Relevant Documents Encoder, to obtain the representations of relevant documents by finding common salient topics through the consideration of the semantic interaction between relevant documents; 3) Irrelevant Documents Encoder, to obtain the representations of irrelevant documents by modeling each irrelevant document separately; 4) Intent Description Decoder, to generate the intent description by contrasting the relevant documents with the irrelevant documents given a query. The overall architecture is depicted in Figure 1 and we will detail our model as follows.

4.2 Query Encoder

The goal of the query encoder is to map the input query to a vector representation. Specifically, we use a bi-directional GRU [10] as the query encoder. Each word w_c^q in the query q is firstly represented by its semantic representation \mathbf{e}_c^q as the input of the encoder. Then, the query encoder represents the query q as a series of hidden vectors $\{\mathbf{h}_c^q\}_{c=1}^C$ modeling the sequence from both forward and backward directions. Finally, we use the concatenated forward and backward hidden state as the query representation \mathbf{x}^q .

4.3 Relevant Documents Encoder

Generally, the relevant documents encoder takes in the input relevant documents, and encodes them into a series of hidden representations. Relevant documents are assumed to share similar underlying topics since they are all related to the specific information need behind a query to different extent. To achieve this purpose, we propose to encode the relevant documents together by concatenating the source M relevant documents into a single relevant mega-document $D_r = \{s_1^r, \dots, s_U^r\}$, with U (i.e., $U = \sum_{m=1}^M T_m$) sentences where each sentence s_u^r contains L_u^r words.




Figure 1: The overall architecture of contrastive generation model (CtrsGen).

We adopt a hierarchical encoder framework, where a word encoder encodes the words of a sentence s_u^r , and a sentence encoder encodes the sentences of a relevant mega-document D_r . We use a bi-directional GRU as both the word and sentence encoder. Firstly, we obtain the hidden state $\mathbf{h}_{u,l}^r$ for a given word $w_{u,l}^r$ in each sentence s_u^r by concatenating the forward and backward hidden states of the word encoder. Then, we concatenate the last hidden states of the forward and backward passes as the embedding representation \mathbf{e}_u^r of the sentence s_u^r . A sentence encoder is used to sequentially receive the embeddings of sentences $\{\mathbf{e}_u^r\}_{u=1}^U$ and the hidden representation \mathbf{h}_u^r of each sentence s_u^r is given by concatenating the forward and backward hidden states of the sentence encoder.

Different from previous simple methods [24, 66] which directly concatenate or aggregate the hidden states of sentences to obtain the relevant mega-document representation, we further employ a *query-aware encoder attention mechanism* as follows, to aggregate the sentence representations according to their importance to obtain a good relevant mega-document representation.

Query-aware Encoder Attention Mechanism The key idea of the query-aware encoder attention mechanism is to identify those important sentences in the relevant documents that can reveal the essential topics of the query rather than treat each sentence equally. Different from pre-defined relevance scores using unigram overlap between query and sentences [3], we leverage the query representation \mathbf{x}^q to estimate the importance of each sentence through attention.

Specifically, the importance score (or weight) γ_u^r of each sentence s_u^r is given by $\gamma_u^r = \text{softmax}(\mathbf{x}^q \cdot \mathbf{Q} \cdot \mathbf{h}_u^r)$, where \mathbf{Q} is a parameter matrix to be learned and the softmax function ensures all the weights sum up to 1. Finally, we obtain the representation of the relevant mega-document \mathbf{x}^r by using the weighted sums of hidden states $\{\mathbf{h}_1^r, \dots, \mathbf{h}_U^r\}$, i.e., $\mathbf{x}^r = \sum_{u=1}^U \gamma_u^r \mathbf{h}_u^r$.

4.4 Irrelevant Documents Encoder

Different from the relevant documents encoder which encodes relevant documents together, the irrelevant documents encoder processes each irrelevant document separately. The key idea is that while relevant documents could be similar to each other, there might be quite different ways for documents to be irrelevant to a

query. Therefore, it is unreasonable to take the semantic interaction between irrelevant documents into consideration.

For each irrelevant document $D_n^i = \{s_{n,1}^i, \dots, s_{n,T_n}^i\}$ where each sentence $s_{n,t}^i$ contains $K_{n,t}^i$ words, we adopt a hierarchical encoder framework similar with that in relevant documents encoder. Thus, we can obtain the hidden representation of each word $w_{n,t,k}^i$ in each sentence $s_{n,t}^i$, i.e., $\{\mathbf{h}_{n,t,k}^i\}_{k=1}^{K_{n,t}^i}$, the embedding representation of each sentence $s_{n,t}^i$ in each irrelevant document D_n^i , i.e., $\{\mathbf{e}_{n,t}^i\}_{t=1}^{T_n}$, and the hidden representation of each sentence $s_{n,t}^i$, i.e., $\{\mathbf{h}_{n,t}^i\}_{t=1}^{T_n}$. Finally, we obtain all the sentence hidden representations in N irrelevant documents, i.e., $\{\mathbf{h}_{n,t}^i\}_{n=1,t=1}^{N,T_n}$.

4.5 Intent Description Decoder

The intent description decoder is responsible for producing the intent description given the representations of the query, relevant documents and irrelevant documents. To generate the intent description y , we employ 1) a *query-aware decoder attention mechanism*, which maintains a query-aware context vector to make sure more important content in the query is attended, and 2) a *contrast-based decoder attention mechanism*, which maintains a document-aware context vector for description generation to distinguish relevant documents from those irrelevant ones with respect to a query.

Specifically, the query-aware context vector \mathbf{c}_z^q and the document-aware context vector \mathbf{c}_z^d are provided as extra inputs to derive the hidden state \mathbf{h}_z^s of the z -th word w_z^y in an intent description and later the probability distribution for choosing the word w_z^y .

Concretely, \mathbf{h}_z^s is defined as,

$$\mathbf{h}_z^s = f_s(w_{z-1}^y, \mathbf{h}_{z-1}^s, \mathbf{c}_z^q, \mathbf{c}_z^d), \quad (2)$$

where f_s is a GRU unit, w_{z-1}^y is the predicted word from vocabulary at $z-1$ -th step when decoding the intent description y .

The initial hidden state of the decoder is defined as the weighted sums of query and relevant mega-document representations,

$$\mathbf{h}_0^s = \mathbf{W}_q \mathbf{x}^q + \mathbf{W}_r \mathbf{x}^r, \quad (3)$$

where \mathbf{W}_q and \mathbf{W}_r are learned parameters.

The probability for choosing the word w_z^y is defined as,

$$p(w_z^y | w_{<z}^y, q, \mathcal{R}, \mathcal{I}) = f_g(w_{z-1}^y, \mathbf{h}_z^s, \mathbf{c}_z^q, \mathbf{c}_z^d), \quad (4)$$

where f_g is a nonlinear function that computes the probability vector for all legal output words at each output time. We now describe the specific mechanism in the follows.

4.5.1 Query-aware Decoder Attention Mechanism. The key idea of the query-aware decoder attention mechanism is to make the generation of a description focusing on the query. The description should contain as much information relevant to the query as possible. We maintain a query-aware context vector \mathbf{c}_z^q for generating the z -th word w_z^y in the description. Specifically, \mathbf{c}_z^q is a weighted sum of the hidden representations of all the words in the query q ,

$$\mathbf{c}_z^q = \sum_{c=1}^C \alpha_{z,c}^q \mathbf{h}_c^q, \quad (5)$$

where $\alpha_{z,c}^q$ indicates how much the c -th word w_c^q from the query q contributes to generating the z -th word in the intent description y , and is usually computed as,

$$\alpha_{z,c}^q = \text{softmax}(\mathbf{h}_c^q \cdot \mathbf{W}_1 \cdot \mathbf{h}_{z-1}^s), \quad (6)$$

where \mathbf{W}_1 is a learned parameter.

4.5.2 Contrast-based Decoder Attention Mechanism. The intent description should cover topics in relevant documents and eliminate topics in irrelevant documents of a given query. To achieve this purpose, we introduce a contrast-based decoder attention mechanism between relevant documents and irrelevant documents with respect to a query.

Specifically, the contrast-based decoder attention mechanism contains two steps: 1) to compute the sentence-level attention weights in the relevant documents when decoding the intent description, 2) to compute the the contrast scores to adjust the sentence-level attention weights in the relevant documents.

- **Sentence-level Attention Weight** Firstly, we compute the attention weight $\alpha_{z,u}^r$ of each sentence s_u^r in relevant mega-document D_r , which indicates how much each sentence s_u^r contributes to generating the z -th word w_z^y in the intent description y ,

$$\alpha_{z,u}^r = \text{softmax}(v^T \tanh(\mathbf{W}_2 \mathbf{h}_u^r + \mathbf{W}_3 \mathbf{c}_z^q + \mathbf{W}_4 \mathbf{h}_{z-1}^s)), \quad (7)$$

where \mathbf{W}_2 , \mathbf{W}_3 and \mathbf{W}_4 are learned parameters. Specifically, the query-aware context vector \mathbf{c}_z^q is used to incorporate query relevance into the focus on the relevant documents.

- **Contrast Score** Then, we compute the contrast score for each sentence $s_u^r \in D_r$, to omit the information similar with the irrelevant documents. The contrast score $\beta_{z,u}^r$ for generating the z -th word in the description is defined as,

$$\hat{\beta}_{z,u}^r = \lambda \text{Sim}(s_u^r, y_{<z}) - (1 - \lambda) \max_{s_{n,t}^i \in \mathcal{I}} \text{Sim}(s_u^r, s_{n,t}^i), \quad (8)$$

$$\beta_{z,u}^r = \text{softmax}(\hat{\beta}_{z,u}^r), \quad (9)$$

where λ is a balancing factor. The two similarity functions are defined as,

- The similarity function $\text{Sim}(s_u^r, y_{<z})$ between each sentence s_u^r in the relevant mega-document and current generated description $y_{<z}$ is defined as,

$$\text{Sim}(s_u^r, y_{<z}) = \mathbf{h}_u^r \cdot \mathbf{W}_5 \cdot \mathbf{h}_{z-1}^s, \quad (10)$$

where \mathbf{W}_5 is a learned parameter.

- The similarity function $\text{Sim}(s_u^r, s_{n,t}^i)$ between each sentence s_u^r in the relevant mega-document and each sentence $s_{n,t}^i$ in

the irrelevant documents is defined as,

$$\text{Sim}(s_u^r, s_{n,t}^i) = \text{softmax}(\tanh(\mathbf{h}_u^r \cdot \mathbf{W}_6 \cdot \mathbf{h}_{n,t}^i), \quad (11)$$

where \mathbf{W}_6 is a learned parameter.

Finally, we maintain a document-aware context vector \mathbf{c}_z^d for generating the z -th word w_z^y in the description y ,

$$\mathbf{c}_z^d = \sum_{u=1}^U \beta_{z,u}^r \alpha_{z,u}^r \mathbf{h}_u^r. \quad (12)$$

4.6 Model Learning

We employ maximum likelihood estimation (MLE) to learn our CtrsGen model in an end-to-end way. Specifically, the training objective is a probability over the training corpus \mathcal{D} ,

$$\arg \max_{\theta} \sum_{(q, \mathcal{R}, \mathcal{I}, y) \in \mathcal{D}} \log p(y|q, \mathcal{R}, \mathcal{I}; \theta). \quad (13)$$

We apply stochastic gradient decent method Adam [31] to learn the model parameters θ .

5 EXPERIMENTS

In this section, we conduct experiments to verify the effectiveness of our proposed model.

5.1 Experimental Settings

To evaluate the performance of our model, we conducted experiments on our Q2ID benchmark dataset. In preprocessing, all the words in queries, documents and descriptions are white-space tokenized and lower-cased, and pure digit words and non-English characters are removed. We randomly divide the 5358 queries into training(5000)/validation(100)/test(258) query sets.

We keep the 120,000 most frequently occurring words in our experiments. All the other words outside the vocabularies are replaced by a special token <UNK> symbol. We implement our model in Tensorflow. Specifically, we use one layer of bi-directional GRU for each encoder and uni-directional GRU for decoder, with the GRU hidden unit size set as 256 in both the encoder and decoder. The dimension of word embeddings is 300. We use pretrained word2vec vectors trained on the same corpus to initialize the word embeddings, and the word embeddings will be further fine-tuned during training. The learning rate of Adam algorithm is set as 0.0005. The learnable parameters (e.g., the parameters \mathbf{W}_q and \mathbf{W}_1) are uniformly initialized in the range $[-0.1, 0.1]$. The mini-batch size for the update is set as 16. We clip the gradient when its norm exceeds 5. We set the $\lambda = 0.5$ to calculate the contrast score in Equ. 8. All the hyper-parameters are tuned on the validation set.

5.2 Baselines

5.2.1 Model Variants. Here, we firstly employ some variants of our model by removing some components and adopting different model architectures.

- **CtrsGen-I** removes the input irrelevant documents, and only considers the effect of the queries and relevant documents.
- **CtrsGen-I+Con** is similar with CtrsGen-I, but considers the decoder attention over the concatenation of the hidden states of the queries and relevant documents.
- **CtrsGen-Q** removes the query, and only considers the effect of the relevant and irrelevant documents.

- **CtrsGen_{Q-I}** removes the query and irrelevant documents, and only considers the effect of the relevant documents.
- **CtrsGen_{IrCon}** concatenates the irrelevant documents and encodes them as a single input.
- **CtrsGen_{ReIrCon}** concatenates the relevant and irrelevant documents into a single input as the relevant mega-document.

5.2.2 *Extractive Models.* We also apply extractive summarization models to extract a sentence from the relevant documents as the intent description.

- **LSA** [59] applies Singular Value Decomposition (SVD) to pick a representative sentence.
- **TextRank** [39] is a graph-based method inspired by the PageRank algorithm.
- **LexRank** [19] is also a graph-based method inspired by the PageRank algorithm. The difference with TextRank is to use different methods to calculate the similarity between two sentences.

5.2.3 *Abstractive Models.* Additionally, we consider neural abstractive models, including query-independent and query-dependent abstractive summarization models, to illustrate how well these systems perform on the Q2ID task.

- **Query-independent abstractive summarization models** based on the relevant documents include,
 - **ABS** [50] is the attention bag-of-words encoder based sentence summarization model.
 - **Extract+Rewrite** [58] firstly scores sentences using LexRank and then generates a title-like summary.
- **Query-dependent abstractive summarization models** based on the query and relevant documents include,
 - **PG** [24] employs a pointer-generator model for query-based summarization.
 - **RSA** [3] incorporates query relevance into the seq2seq framework for query-based summarization.
 - **SD₂** [43] introduces a diversification mechanism in query-based summarization for handling the duplication problem.

5.3 Evaluation Methodologies

We use both automatic and human evaluation to measure the quality of intent descriptions generated by our model and the baselines.

For automatic evaluation, following the previous studies [20, 33, 50], we adopt the widely used automatic metric Rouge [36] to evaluate n-grams of the generated intent descriptions with gold-standard descriptions as references. We report recall results on Rouge-1, Rouge-2 and Rouge-L.

For human evaluation, we consider two evaluation metrics: 1) Naturalness, which indicates whether the intent description is grammatically correct and fluent; and 2) Reasonableness, which measures the semantic similarity between generated intent descriptions and the golden baseline descriptions. We asked three professional native speakers to rate the 258 test quadruples in terms of the metrics mentioned above on a 1~5 scale (5 for the best).

5.4 Ablation Analysis

We conduct ablation analysis to investigate the effect of proposed mechanisms in our CtrsGen model. As shown in Table 3, we can find that: (1) By removing the irrelevant documents, the performance

Table 3: Ablation analysis of our CtrsGen model with its variants under the automatic evaluation (%). Two-tailed t-tests demonstrate the improvements of CtrsGen to the variants are statistically significant ([‡] indicates p-value < 0.01).

Model	Rouge-1	Rouge-2	Rouge-L
CtrsGen _I	23.31	4.05	19.03
CtrsGen _{I+Con}	23.26	4.03	19.02
CtrsGen _Q	23.07	3.94	18.61
CtrsGen _{Q-I}	22.55	3.59	17.25
CtrsGen _{IrCon}	24.19	4.51	19.43
CtrsGen _{ReIrCon}	22.62	3.60	17.14
CtrsGen	24.76[‡]	4.62	20.21[‡]

Table 4: Comparisons between our CtrsGen and the baselines under the automatic evaluation (%).

Model	Rouge-1	Rouge-2	Rouge-L
LexRank	13.73	1.74	10.92
LSA	18.49	2.05	14.50
TextRank	20.15	2.95	16.43
ABS	17.21	1.23	12.83
Extract+Rewrite	18.85	2.49	15.45
PG	20.81	3.04	17.25
RSA	19.65	2.39	16.17
SD ₂	21.37	3.25	18.49
CtrsGen	24.76	4.62	20.21

of *CtrsGen_I* and *CtrsGen_{I+Con}* has a significant drop as compared with *CtrsGen*. The results indicate that contrasting the relevant documents with the irrelevant documents does help generate the intent description better. (2) *CtrsGen_Q* performs worse than *CtrsGen_I*, showing that the information in the query has much bigger impact than that in irrelevant documents for extracting salient information for intent description generation. (3) *CtrsGen_{IrCon}* performs worse than *CtrsGen*. The reason might be that it tends to bring noisy information when considering the interaction between irrelevant documents. (4) The performance of *CtrsGen_{ReIrCon}* is relatively poor, indicating that considering all judged documents as relevant tends to bring noisy information that may hurt the intent description generation. (5) *CtrsGen_{Q-I}* gives the worst performance, indicating that traditional multi-document summarization model without considering the query and irrelevant documents is not suitable for the Q2ID task. (6) By including all the mechanisms, *CtrsGen* achieves the best performance in terms of all the evaluation metrics.

5.5 Baseline Comparison

The performance comparisons between our model and the baselines are shown in Table 4. We have the following observations: (1) The abstractive methods generally outperform the extractive methods, since those extractive methods are unsupervised in nature. (2) The query-independent abstractive summarization models (i.e., *ABS* and *Extract+Rewrite*) perform worse than the query-dependent abstractive summarization models (i.e., *PG*, *RSA* and *SD₂*), showing that it is necessary to generate the intent description with the guidance of the query. (3) By introducing a diversification mechanism for solving the duplication problem, *SD₂* improves Rouge scores when compared to *PG* and *RSA*. (4) As compared with the best-performing

Table 5: Results on the human evaluation. Best% is the ratio of the best score in the two metrics.

	Naturalness	Reasonableness	Best%
SD_2	3.14	1.73	17.46
<i>CtrsGen</i>	3.46	2.86	34.63
Human	4.72	4.21	70.42

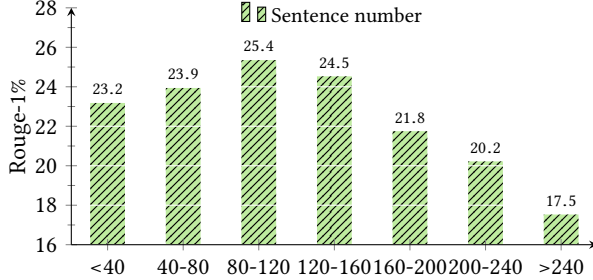


Figure 2: Performance comparison of CtrsGen for different sentence numbers of relevant documents.

baseline SD_2 , the relative improvement of *CtrsGen* over SD_2 is about 42.1% in terms of Rouge-2. (5) Our *CtrsGen* model can outperform all the baselines significantly (p -value < 0.01), demonstrating the effectiveness of the contrastive generation idea.

Table 5 shows the results of the human evaluation. We can see that our *CtrsGen* outperforms the best performing baseline SD_2 in all evaluation metrics. The results imply that our model can generate fluent and grammatically correct intent descriptions (i.e., Naturalness) which better interpret the search intent behind the queries (i.e., Reasonableness) than the baseline SD_2 .

5.6 Breakdown Analysis

Beyond above overall performance analysis, we also take some breakdown analysis for the Q2ID task.

5.6.1 Analysis on Query Types. There exist two query types in our Q2ID dataset, i.e., natural language questions from SemEval and keyword-based queries from TREC. There are 103 keyword queries and 155 natural language questions respectively in the test dataset. Here, we analyze the generated intent description for different query types from our *CtrsGen* model. As shown in Table 6, we can see that *CtrsGen* model for questions perform better than that for keywords. The major reason might be that the natural language questions are usually long and thus bring more key information needs, which could help capture the search intent better.

5.6.2 Analysis on Relevant Document Length. We also analyze the effect of relevant document length for intent description generation. Since the relevant documents are concatenated into a single relevant mega-document in our *CtrsGen* model, we depict the histogram of Rouge-1 results over different sentence numbers of relevant mega-documents. The results are shown in Figure 2. For the test dataset, the average sentence number of relevant mega-documents is 203. From the results, we can observe that by introducing less than 80 or more than 120 relevant documents, our model tends to bring insufficient information or noisy information that may hurt the intent description generation.

Table 6: Performance comparison of our CtrsGen model for different query types (%).

Query Type	Rouge-1	Rouge-2	Rouge-L
Question	24.99	4.99	21.35
Keyword	23.12	4.15	19.56

Table 7: Performance comparison of our CtrsGen model for queries with and without irrelevant documents (%).

Irrelevant Documents	Rouge-1	Rouge-2	Rouge-L
With	25.21	4.73	20.44
Without	24.03	4.14	19.12

Table 8: An example from the test Q2ID data.

Query	Community Care Centers
Relevant Documents: (D_1^P) Another issue the county faced was Ebola’s spread in community care centers. Community care centers are part of the strategy to combat the ongoing Ebola epidemic ... (D_2^P) Centre provides an alternative to Ebola Treatment Units where residents can seek ... D_{38}^P	
Irrelevant Documents: (D_1^I) Liberia is actively considering the concept of Community Care Centers (CCCs) smaller 10-20 bed units located in “hot spot” ... (D_2^I) Save the Children is constructing and operating Ebola Community Care Centers to provide “close to the community” care ...	
Intent Description:	
Ground Truth: Describe the new initiative, and also the controversy, surrounding the opening of Community Care Centers throughout West Africa and how they will aid in combating the spread of Ebola.	
SD_2: Community Care Centers is constructed to provide services and care for Ebola cases.	
CtrsGen: World Health Organization establishes Community Care Centers in Africa to combat Ebola epidemic. What measures are taken to prevent the Ebola’s spread.	

5.6.3 Analysis on Irrelevant Documents Existence. To further analyze the effect of irrelevant documents, we conduct a comparison of the generated intent descriptions from our *CtrsGen* model for queries with and without irrelevant documents. There are 223 and 35 queries with and without irrelevant documents respectively in the test dataset, As shown in Table 7, we can find that the performance of *CtrsGen* for queries with irrelevant documents is better than that for queries without irrelevant documents. These results again demonstrate the effectiveness of contrasting the relevant documents with the irrelevant documents of a given query.

5.7 Case Study

To better understand how different models perform, we show the generated intent description from our *CtrsGen* model as well as that from the best baseline model SD_2 . We take one query “Community Care Centers” from the test data as an example. Due to the limited space, we only show some key sentences. As shown in Table 8, we can see that without the consideration of irrelevant documents, SD_2 focuses on the “services for Ebola cases”, instead of “combat for Ebola spread”. On the contrary, by leveraging the irrelevant documents, our model can better distill the essential information, and then generate a much more accurate intent description which is more consistent with the ground-truth.




Figure 3: (a) and (b) is the heatmap of the sentence-level decoder attention weights in relevant documents for generating the first word in the description, given by *CtrsGen-I* and *CtrsGen* respectively. Deeper shading denotes higher value.

Furthermore, we analyze the effect of irrelevant documents in our *CtrsGen* model. As shown in Figure 3, we visualize the sentence-level decoder attention weights $\alpha_{z,u}^r$ (Eq. (7)) over the relevant documents from our model variant *CtrsGen-I*, and the adjusted weights $\beta_{z,u}^r \alpha_{z,u}^r$ (Eq. (12)) over the relevant documents from our *CtrsGen* model. From the test data, we select a new query “Mexican Air Pollution” with ground-truth intent description “Mexico City has the worst air pollution in the world. Pertinent Documents would contain the specific steps Mexican authorities have taken to combat this deplorable situation”. Due to space limitation, we only visualize sampled 6 sentences in the relevant documents for generating the first word in the description. As we can see, *CtrsGen-I* pays too much attention on the 2-th and 6-th sentences which confuse the model to generate a description mainly about the “Free Trade Agreement”. Specifically, the attention weight of the 2-th and 6-th sentence computed by *CtrsGen* has a significantly drop as compared with *CtrsGen-I* due to the high similarity with irrelevant documents. This in turn guides the decoder to pay attention to those informative sentences and generate a much better intent description.

5.8 Potential Application

Here, we discuss the potential usage of such Q2ID technique. An interesting application would be to facilitate the exploratory search, i.e., interpreting the search results in exploratory search.

In exploratory search, users are often unclear with their information needs initially and thus queries may be broad and vague. Based on the pseudo relevant feedback idea [61], we can treat the top k ranked documents as relevant documents and others as irrelevant documents, and then leverage the Q2ID technique to generate the intent description. Such description could be viewed as an explanation of how the search engine understands the query and why those documents are displayed at the top ranks. With such interpretation, users may better understand the search results and find a direction to refine their query. For instance, as shown in Figure 4, the query is “Shanghai Disneyland”, and an intent description of this ambiguous query “Describe the general information about Shanghai Disneyland, such as location, history, and guide maps. How much is a ticket to Shanghai Disneyland and where can I buy it?” is generated based on the understanding of the search intent by the search engine. If the user’s intent is to find “What are the must-see attractions at Shanghai Disneyland?”, he/she will easily find that the search engine has not captured that aspect and a refinement (e.g., Shanghai Disneyland Attractions) is necessary in the following. This can provide the user with a better overall search experience, in turn, boost the retrieval performance and the confidence of a search engine.

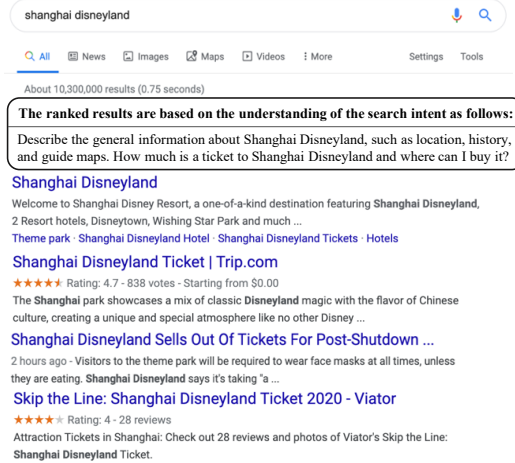


Figure 4: An example application in exploratory search using our Q2ID technique.

6 CONCLUSION

In this paper, we introduced a challenging Q2ID task for query understanding via generating a natural language intent description based on relevant and irrelevant documents of a given query. To tackle this problem, we developed a novel Contrastive Generation model to contrast the relevant documents with the irrelevant documents given a query. Empirical results over our constructed Q2ID dataset showed that our model can well understand the query through a detailed and precise intent description.

In the future work, we would like to consider the multi-graded relevance labels of documents with respect to the query intent, and realize the potential usage of such Q2ID technique in the above discussed application. Also, it is valuable to do experiments on transferring intent descriptions learned on one corpus to another due to the lack of labeled data of this kind.

7 ACKNOWLEDGMENTS

This work was supported by Beijing Academy of Artificial Intelligence (BAAI) under Grants No. BAAI2019ZD0306 and BAAI2020ZJ0303, and funded by the National Natural Science Foundation of China (NSFC) under Grants No. 61722211, 61773362, 61872338, and 61902381, the Youth Innovation Promotion Association CAS under Grants No. 20144310, and 2016102, the National Key RD Program of China under Grants No. 2016QY02D0405, the Lenovo-CAS Joint Lab Youth Scientist Project, the K.C.Wong Education Foundation, and the Foundation and Frontier Research Key Program of Chongqing Science and Technology Commission (No. cstc2017jcyjBX0059).

REFERENCES

- [1] Ricardo Baeza-Yates, Liliana Calderón-Benavides, and Cristina González-Caro. 2006. The intention behind web queries. In *SPIRE*.
- [2] Dzhmitry Bahdanau, Kyunghyun Cho, and Yoshua Bengio. 2015. Neural machine translation by jointly learning to align and translate. In *ICLR*.
- [3] Tal Baumel, Matan Eyal, and Michael Elhadad. 2018. Query focused abstractive summarization: Incorporating query relevance, multi-document coverage, and summary length constraints into seq2seq models. In *arXiv*.
- [4] Doug Beeferman and Adam Berger. 2000. Agglomerative clustering of a search engine query log. In *KDD*.
- [5] Steven M Beitzel, Eric C Jensen, Ophir Frieder, David D Lewis, Abdur Chowdhury, and Aleksander Kolcz. 2005. Improving automatic query classification via semi-supervised learning. In *ICDM*.
- [6] Andrei Broder. 2002. A taxonomy of web search. In *ACM Sigir forum*.
- [7] Huanhuan Cao, Derek Hao Hu, Dou Shen, Daxin Jiang, Jian-Tao Sun, Enhong Chen, and Qiang Yang. 2009. Context-aware query classification. In *SIGIR*.
- [8] Jaime G Carbonell and Jade Goldstein. 1998. The use of MMR, diversity-based reranking for reordering documents and producing summaries.. In *SIGIR*.
- [9] Hanxiong Chen, Chen Xu, Shi Shaoyun, and Zhang Yongfeng. 2019. Generate Natural Language Explanations for Recommendation. In *SIGIR*.
- [10] Kyunghyun Cho, Bart Van Merriënboer, Caglar Gulcehre, Dzhmitry Bahdanau, Fethi Bougares, Holger Schwenk, and Yoshua Bengio. 2014. Learning phrase representations using RNN encoder-decoder for statistical machine translation. In *EMNLP*.
- [11] Sangwoo Cho, Logan Lebanoff, Hassan Foroosh, and Fei Liu. 2019. Improving the Similarity Measure of Determinantal Point Processes for Extractive Multi-Document Summarization. In *ACL*.
- [12] Sumit Chopra, Michael Auli, and Alexander M Rush. 2016. Abstractive sentence summarization with attentive recurrent neural networks. In *NAACL*.
- [13] Felipe Costa, Sixun Ouyang, Peter Dolog, and Aonghus Lawlor. 2017. Automatic Generation of Natural Language Explanations. (2017).
- [14] Hang Cui, Ji-Rong Wen, Jian-Yun Nie, and Wei-Ying Ma. 2002. Probabilistic query expansion using query logs. In *WWW*.
- [15] Van Dang and Bruce W Croft. 2010. Query reformulation using anchor text. In *WSDM*.
- [16] Hal Daumé III and Daniel Marcu. 2006. Bayesian query-focused summarization. In *ACL*. 305–312.
- [17] Finale Doshi-Velez and Been Kim. 2017. Towards A Rigorous Science of Interpretable Machine Learning. (2017).
- [18] Huizhong Duan, Emre Kiciman, and ChengXiang Zhai. 2012. Click patterns: an empirical representation of complex query intents. In *CIKM*.
- [19] Günes Erkan and Dragomir R Radev. 2004. Lexrank: Graph-based lexical centrality as salience in text summarization. *Journal of artificial intelligence research* (2004).
- [20] Alexander R Fabbri, Irene Li, Tianwei She, Suyi Li, and Dragomir R Radev. 2019. Multi-News: a Large-Scale Multi-Document Summarization Dataset and Abstractive Hierarchical Model. In *ACL*.
- [21] Hui Fang. 2008. A re-examination of query expansion using lexical resources. In *ACL*.
- [22] Jade Goldstein, Mark Kantrowitz, Vibhu Mittal, and Jaime Carbonell. 1999. Summarizing text documents: sentence selection and evaluation metrics. In *SIGIR*.
- [23] Luis Gravano, Vasileios Hatzivassiloglou, and Richard Lichtenstein. 2003. Categorizing web queries according to geographical locality. In *CIKM*.
- [24] Johan Hasselqvist, Niklas Helmerz, and Mikael Kågebäck. 2017. Query-based abstractive summarization using neural networks. In *arXiv*.
- [25] Yuan Hong, Jaideep Vaidya, Haibing Lu, and Wen Ming Liu. 2016. Accurate and efficient query clustering via top ranked search results. In *Web Intelligence*.
- [26] Wan-Ting Hsu, Chieh-Kai Lin, Ming-Ying Lee, Kerui Min, Jing Tang, and Min Sun. 2018. A unified model for extractive and abstractive summarization using inconsistency loss. In *ACL*.
- [27] Jian Hu, Gang Wang, Fred Lochovsky, Jian-tao Sun, and Zheng Chen. 2009. Understanding user’s query intent with wikipedia. In *WWW*.
- [28] Rosie Jones and Fernando Diaz. 2007. Temporal profiles of queries. *TOIS* (2007).
- [29] Zsolt T Kardkovács, Domonkos Tikk, and Zoltán Bánvágyi. 2005. The ferretty algorithm for the KDD Cup 2005 problem. *SIGKDD* (2005).
- [30] Maryam Karimzadehgan and Cheng Xiang Zhai. 2011. Improving retrieval accuracy of difficult queries through generalizing negative document language models. In *CIKM*.
- [31] Diederik P Kingma and Jimmy Ba. 2015. Adam: A method for stochastic optimization. In *ICLR*.
- [32] SK Kolluru and Prasenjit Mukherjee. 2016. Query Clustering using Segment Specific Context Embeddings. In *arXiv*.
- [33] Logan Lebanoff, Kaiqiang Song, and Fei Liu. 2018. Adapting the neural encoder-decoder framework from single to multi-document summarization. In *EMNLP*.
- [34] Mu-Chu Lee, Bin Gao, and Ruofei Zhang. 2018. Rare query expansion through generative adversarial networks in search advertising. In *SIGKDD*.
- [35] Ying Li, Zijian Zheng, and Honghua Kathy Dai. 2005. KDD CUP-2005 report: Facing a great challenge. *SIGKDD* (2005).
- [36] Chin-Yew Lin. 2004. Rouge: A package for automatic evaluation of summaries. In *Text summarization branches out*.
- [37] Melvin Earl Maron and John Larry Kuhns. 1960. On relevance, probabilistic indexing and information retrieval. *JACM* (1960).
- [38] Lingling Meng, Runqing Huang, and Junzhong Gu. 2013. A new algorithm of web queries clustering using user feedback. *IJSIP* (2013).
- [39] Rada Mihalcea and Paul Tarau. 2004. TextRanks: Bringing order into text. In *EMNLP*.
- [40] Ahmed A Mohamed and Sanguthevar Rajasekaran. 2006. Improving query-based summarization using document graphs. In *ISSPIT*.
- [41] Preslav Nakov, Doris Hoogeveen, Lluís Màrquez, Alessandro Moschitti, Hamdy Mubarak, Timothy Baldwin, and Karin Verspoor. 2017. SemEval-2017 task 3: Community question answering. In *SemEval*.
- [42] Ya nan Qian, Tetsuya Sakai, Junting Ye, Qinghua Zheng, and Cong Li. 2013. Dynamic query intent mining from a search log stream. In *CIKM*.
- [43] Preksha Nema, Mitesh Khapra, Anirban Laha, and Balaraman Ravindran. 2017. Diversity driven attention model for query-based abstractive summarization. In *ACL*.
- [44] Atsushi Otsuka, Kyosuke Nishida, Katsuji Bessho, Hisako Asano, and Junji Tomita. 2018. Query expansion with neural question-to-answer translation for FAQ-based question answering. In *WebConf*.
- [45] Marius A Pasca and Sandra M Harabagiu. 2001. High performance question answering. In *SIGIR*.
- [46] Zhang Peng, Yuexian Hou, and Dawei Song. 2009. Approximating True Relevance Distribution from a Mixture Model based on Irrelevance Data. In *SIGIR*.
- [47] Filip Radlinski, Martin Szummer, and Nick Craswell. 2010. Inferring query intent from reformulations and clicks. In *WWW*.
- [48] Manukonda Sumathi Rani and Geddati China Babu. 2019. Efficient Query Clustering Technique and Context Well-Informed Document Clustering. In *Soft Computing and Signal Processing*.
- [49] Daniel E Rose and Danny Levinson. 2004. Understanding user goals in web search. In *WWW*.
- [50] Alexander M Rush, Sumit Chopra, and Jason Weston. 2015. A neural attention model for abstractive sentence summarization. In *ACL*.
- [51] Frank Schilder and Ravikumar Kondadadi. 2008. FastSum: fast and accurate query-based multi-document summarization. In *ACL*.
- [52] Dou Shen, Rong Pan, Jian-Tao Sun, Jeffrey Junfeng Pan, Kangheng Wu, Jie Yin, and Qiang Yang. 2005. Q 2 C@ UST: our winning solution to query classification in KDDCUP 2005. *SIGKDD* (2005).
- [53] Ben Shneiderman, Don Byrd, and W Bruce Croft. 1997. Clarifying search: A user-interface framework for text searches. *D-lib magazine* (1997).
- [54] Jaspreet Singh and Avishek Anand. 2019. EXS: Explainable Search Using Local Model Agnostic Interpretability. In *WSDM*.
- [55] Jaspreet Singh and Avishek Anand. 2020. Model Agnostic Interpretability of Raters via Intent Modelling. In *Proceedings of the 2020 Conference on Fairness, Accountability, and Transparency*.
- [56] Jagendra Singh and Aditi Sharan. 2015. Co-occurrence and semantic similarity based hybrid approach for improving automatic query expansion in information retrieval. In *ICDCIT*.
- [57] Jagendra Singh and Aditi Sharan. 2015. Context window based co-occurrence approach for improving feedback based query expansion in information retrieval. In *IJIRR*.
- [58] Kaiqiang Song, Lin Zhao, and Fei Liu. 2018. Structure-infused copy mechanisms for abstractive summarization. (2018).
- [59] Josef Steinberger and Karel Ježek. 2004. Text summarization and singular value decomposition. In *International Conference on Advances in Information Systems*. Springer, 245–254.
- [60] I Sutskever, O Vinyals, and QV Le. 2014. Sequence to sequence learning with neural networks. In *NIPS*.
- [61] Tao Tao and Cheng Xiang Zhai. 2006. Regularized estimation of mixture models for robust pseudo-relevance feedback. In *SIGIR*.
- [62] Manisha Verma and Debasis Ganguly. 2019. LIRME: Locally Interpretable Ranking Model Explanation. In *SIGIR*.
- [63] Ellen M Voorhees. 1994. Query expansion using lexical-semantic relations. In *SIGIR*.
- [64] Ji-Rong Wen, Jian-Yun Nie, and Hong-Jiang Zhang. 2002. Query clustering using user logs. *TOIS* (2002).
- [65] Chen Xu, Qin Zheng, Yongfeng Zhang, and Xu Tao. 2016. Learning to Rank Features for Recommendation over Multiple Categories. (2016).
- [66] Jianmin Zhang, Jiwei Tan, and Xiaojun Wan. 2018. Adapting neural single-document summarization model for abstractive multi-document summarization: A pilot study. In *INLG*.
- [67] Zhiyong Zhang and Olfa Nasraoui. 2006. Mining search engine query logs for query recommendation. In *WWW*.
- [68] Zhiwei Zhang, Qifan Wang, Luo Si, and Jianfeng Gao. 2016. Learning for efficient supervised query expansion via two-stage feature selection. In *SIGIR*.