



MacAvaney, S., Yates, A., Feldman, S., Downey, D., Cohan, A. and Goharian, N. (2021) Simplified Data Wrangling with `ir_datasets`. In: SIGIR 2021: 44th International ACM SIGIR Conference on Research and Development in Information Retrieval, 11-15 July 2021, pp. 2429-2436. (doi:[10.1145/3404835.3463254](https://doi.org/10.1145/3404835.3463254))

There may be differences between this version and the published version. You are advised to consult the publisher's version if you wish to cite from it.

© 2021 Association for Computing Machinery. This is the author's version of the work. It is posted here for your personal use. Not for redistribution. The definitive Version of Record was published in the Proceedings of the SIGIR 2021: 44th International ACM SIGIR Conference on Research and Development in Information Retrieval, 11-15 July 2021, pp. 2429-2436. (doi:[10.1145/3404835.3463254](https://doi.org/10.1145/3404835.3463254))

<http://eprints.gla.ac.uk/240149/>

Deposited on: 10 May 2021

Enlighten – Research publications by members of the University of
Glasgow

<http://eprints.gla.ac.uk>

Simplified Data Wrangling with `ir_datasets`

Sean MacAvaney*

University of Glasgow
IR Lab, Georgetown University
sean.macavaney@glasgow.ac.uk

Andrew Yates

Max Planck Institute for Informatics
ayates@mpi-inf.mpg.de

Sergey Feldman, Doug Downey, Arman Cohan

Allen Institute for AI
{sergey,doug,d,armanc}@allenai.org

Nazli Goharian

IR Lab, Georgetown University
nazli@ir.cs.georgetown.edu

ABSTRACT

Managing the data for Information Retrieval (IR) experiments can be challenging. Dataset documentation is scattered across the Internet and once one obtains a copy of the data, there are numerous different data formats to work with. Even basic formats can have subtle dataset-specific nuances that need to be considered for proper use. To help mitigate these challenges, we introduce a new robust and lightweight tool (`ir_datasets`) for acquiring, managing, and performing typical operations over datasets used in IR. We primarily focus on textual datasets used for ad-hoc search. This tool provides both a Python and command line interface to numerous IR datasets and benchmarks. To our knowledge, this is the most extensive tool of its kind. Integrations with popular IR indexing and experimentation toolkits demonstrate the tool’s utility. We also provide documentation of these datasets through the `ir_datasets` catalog: <https://ir-datasets.com/>. The catalog acts as a hub for information on datasets used in IR, providing core information about what data each benchmark provides as well as links to more detailed information. We welcome community contributions and intend to continue to maintain and grow this tool.

CCS CONCEPTS

• Information systems → Information retrieval.

KEYWORDS

information retrieval, datasets, benchmarks

1 INTRODUCTION

The datasets and benchmarks we use are a cornerstone of Information Retrieval (IR) research. Unfortunately, many of these datasets remain frustrating to find and manage. Once obtained, the variety of data formats can be a challenge to work with. Even data formats that seem simple can hide subtle problems. For example, the TSV files used by the MS-MARCO [66] has a double-encoding problem that affects special characters in roughly 20% of documents.

Recently, several tools have begun to incorporate automatic dataset acquisition. These include Capreolus [93], PyTerrier [58] and OpenNIR [55]. These reduce the user burden of finding the dataset source files and figuring out how to parse them correctly. However, the dataset coverage of each individually is patchy, as shown in Table 1. Further, using the dataset interfaces outside of these tools can be difficult, as they are often tightly coupled

with the tool’s primary functionality. Finally, each of these tools keep their own copy of data, leading to wasted storage. Thus, it is advantageous to have a lightweight tool that focuses on data acquisition, management, and typical operations like lookups.

Many tools rely on manual instructions for downloading, extracting, and processing datasets.¹ We believe providing a tool to automatically perform as much of this work as possible is clearly preferable to this approach since it ensures proper processing of data. A *common* automatic tool has additional advantages, such as reducing redundant copies of datasets and easily allowing tools to be run on alternative or custom datasets with little effort.

Anserini [91] and its Python interface Pyserini [53] use a hybrid approach by distributing copies of queries and relevance judgments in the package itself and primarily relying on manual instructions for document processing. Sometimes Anserini provides document content via downloadable indices.

Other dataset distribution tools are not well-suited for IR tasks. For instance, packages like HuggingFace Datasets [90] and TensorFlow Datasets [3] take a record-centric approach that is not well-suited for relational data like documents, queries, and query-document relevance assessments. Furthermore, IR work involves additional important use cases when working with datasets, such as efficiently looking up a document by ID, for which the designs of prior libraries is not conducive. Dataset schemata, such as DCAT and schema.org, provide a common format machine-readable dataset documentation, which could be supported in the future.

In this work, we present `ir_datasets`, a tool to aid IR researchers in the discovery, acquisition, and management of a variety of IR datasets. The tool provides a simple and lightweight Python and command line interface (see Figure 1) allowing users to iterate the documents, queries, relevance assessments, and other relations provided by a dataset. This is useful for indexing, retrieval, and evaluation of ad-hoc retrieval systems. A document lookup API provides fast access to source documents, which is useful for recent text-based ranking models, such as those that use BERT [27]. PyTerrier [58], Capreolus [93], and OpenNIR [55] recently added support for `ir_datasets`, greatly expanding the number of datasets they support, and other tools like Anserini [91] can utilize our tool using the command line interface. Finally, the `ir_datasets` catalog² acts as a documentation hub, making it easy to find datasets and learn about their characteristics. We intend to continue to backfill prior

¹Such as <https://github.com/castorini/anserini/blob/master/docs/experiments-msmarco-passage.md>, <https://github.com/thunlp/OpenMatch/blob/master/docs/experiments-msmarco.md>, <https://github.com/microsoft/ANCE#data-download>, etc.

²<https://ir-datasets.com/>

*This work was conducted in part during an internship at the Allen Institute for AI.

Table 1: Dataset support in Capreolus [93] (Cap.), PyTerrier [58] (PT), OpenNIR [55] (ONIR), Anserini [91] (Ans.), and ir_datasets (IRDS). ✓ indicates built-in support that automatically provides documents, queries, and query relevance judgments (i.e., as an automatic download). ◊ indicates support for a dataset with some manual effort (e.g., specifying the document parser and settings to use). Datasets marked with * have licenses that require manual effort (e.g., requesting from NIST), and therefore can at most have ◊.

Dataset/Benchmark	Cap.	PT	ONIR	Ans.	IRDS
News					
NYT* [57, 75]		◊	◊	◊	◊
TREC Arabic* [33–35]		◊	◊	◊	◊
TREC Common Core* [7]		◊		◊	◊
TREC Mandarin* [73, 78, 89]		◊	◊	◊	◊
TREC News* [79, 80]		◊		◊	◊
TREC Robust* [83, 85]	◊	◊	◊	◊	◊
TREC Spanish* [36, 37, 74]		◊	◊	◊	◊
Question Answering					
ANTIQUA [38]	✓	◊	✓	◊	✓
MS-MARCO Doc. [66]		✓	◊	✓	✓
MS-MARCO Pass. [66]	✓	✓	✓	✓	✓
MS-MARCO QnA [66]					✓
Natural Questions [48, 50]				◊	✓
TREC CAR [28, 29]			✓	◊	✓
TREC DL [25, 26]	✓	✓	✓	✓	✓
TREC DL-Hard [59]		◊	◊	◊	✓
TriviaQA [47, 48]				◊	✓
Scientific, Bio-medical, Health					
Cranfield [1]					✓
CLEF eHealth* [64, 94]		◊		◊	✓
NFCorpus [9]	✓				✓
TREC CDS [71, 72, 77]					✓
TREC COVID [84, 88]	✓	✓	✓	✓	✓
TREC Genomics [40–43]					✓
TREC Health Misinfo.* [4]		◊		◊	✓
TREC PM [68–70]		◊			✓
TripClick* [67]		◊		◊	✓
Vaswani [2]		✓			✓
Web					
NTCIR WWW* [54, 62]		◊		◊	◊
ORCAS [21]		◊	◊	◊	✓
TREC Million Query* [5, 6, 11]		◊		◊	◊
TREC Terabyte* [10, 12, 13]		◊		◊	◊
TREC Web* [14–17, 19, 20, 22–24]		◊		◊	◊
Other/Miscellaneous					
BEIR [8, 9, 18, 30, 39, 44, 50, 60, 66, 81, 82, 84, 86–88, 92]					✓
CodeSearchNet [45]	✓				✓
TREC Microblog [51, 52, 76]				◊	✓
WikIR [31, 32]			✓		✓

```

1 import ir_datasets
2 dataset = ir_datasets.load('msmarco-passage/train')
3 for doc in dataset.docs_iter(): # documents
4     print(doc)
5     > # GenericDoc(doc_id='0', text='The presence of commun...
6     > # GenericDoc(doc_id='1', text='The Manhattan Project ...
7     > # ...
8
9 for query in dataset.queries_iter(): # queries
10    print(query)
11    > # GenericQuery(query_id='121352', text='define extreme')
12    > # GenericQuery(query_id='634306', text='what does chatt...
13    > # ...
14
15 for qrel in dataset.qrels_iter(): # relevance judgments
16    print(qrels)
17    > # TrecQrel(query_id='1185869', doc_id='0', relevance=1)
18    > # TrecQrel(query_id='1185868', doc_id='16', relevance=1)
19    > # ...
20
21 # Look up documents by ID
22 docs_store = dataset.docs_store()
23 docs_store.get("16")
24 > # GenericDoc(doc_id='16', text='The approach is based...

```

```

1 $ ir_datasets export msmarco-passage docs
2 > # 0 The presence of communication amid scientific ...
3 > # 1 The Manhattan Project and its atomic bomb hel...
4 > # ...
5
6 $ ir_datasets export msmarco-passage/train queries
7 > # 121352 define extreme
8 > # 634306 what does chattel mean on credit history
9 > # ...
10
11 $ ir_datasets export msmarco-passage/train qrels
12 > # 1185869 0 0 1
13 > # 1185868 0 16 1
14 > # ...
15
16 # Look up documents by ID
17 $ ir_datasets lookup msmarco-passage/train 16
18 > # 16 The approach is based on a theory of justice ...

```

Figure 1: Parallel examples of common use cases in ir_datasets using Python and the command line interface.

datasets and add support for new datasets as they are released. The package is open source,³ and we welcome contributions.

2 IR_DATASETS

ir_datasets is a lightweight tool focused on providing easy access to a variety of IR datasets and benchmarks. It provides both a Python and command line interface (see Figure 1), allowing it to be easily used by a variety of toolkits, or simply for ad-hoc data exploration.

To achieve these goals, ir_datasets adheres to several design principles. First, to stay lightweight, the tool is focused on core dataset operations, such as downloading content, iterating through queries or documents, and performing document lookups by ID. This policy explicitly leaves functionality like full-text indexing or neural network processing to other tools. Further, to be practical in a variety of environments, ir_datasets attempts to keep a low memory footprint by using inexpensive data structures and iterators. Finally, in order to leave maximum flexibility to the tool’s users, we attempt to perform “just enough” processing of the data to account for various formats, while not removing information

³https://github.com/allenai/ir_datasets/

Table 2: Entity types in `ir_datasets`.

Entity Type	Python API Example	Description
<code>docs</code>	<code>ds.docs_iter()</code>	A document (or passage for passage retrieval). Contains a <code>doc_id</code> and one or more text fields.
<code>queries</code>	<code>ds.queries_iter()</code>	A query (topic). Contains a <code>query_id</code> and one or more text fields.
<code>qrels</code>	<code>ds.qrels_iter()</code>	A query relevance assessment. Maps a <code>query_id</code> and <code>doc_id</code> to a relevance score or other human assessments.
<code>scoreddocs</code>	<code>ds.scoreddocs_iter()</code>	(uncommon) A scored document (akin to a line from a run file). Maps a <code>query_id</code> and <code>doc_id</code> to a ranking score from a system. Available for datasets that provide an initial ranking (for testing reranking systems).
<code>docpairs</code>	<code>ds.docpairs_iter()</code>	(uncommon) A pair of documents (useful for training). Maps a <code>query_id</code> to two or more <code>doc_ids</code> . Available for datasets that provide suggested training pairs.

that is potentially useful. We hope that this commitment to being lightweight and flexible makes `ir_datasets` an attractive tool to jump-start or enhance other tools for doing IR research.

2.1 Dataset Identifiers

Since no standard identifiers (IDs) exist for datasets in IR, we propose hierarchical dataset IDs. These IDs allow datasets to be looked up in the Python API, command line interface, and online documentation. IDs are usually in the format of `corpus/benchmark`. For instance, the TREC COVID [84] benchmark uses the CORD-19 [88] document corpus and is given an ID of `cord19/trec-covid`. In this case, `cord19` provides documents, while `cord19/trec-covid` provides queries and relevance judgments for those documents.

2.2 Simple & Memorable Python API

A dataset object can be obtained simply by calling:

```
import ir_datasets
ds = ir_datasets.load("dataset-id")
```

Each dataset objects provides access to a number of entity types (see Table 2). Dataset objects are stateless; they simply define the capabilities and the procedures for obtaining and processing the data.

Most ad-hoc retrieval datasets consist of 3 main entity types: documents (`docs`), queries/topics (`queries`), and query relevance assessments (`qrels`). In the spirit of being simple, lightweight, and low-memory, entities are provided as `namedtuple` instances from iterators. For each entity type provided by a particular dataset, there is a corresponding `ds.{entity}_iter()` function that returns an iterator (e.g., `ds.docs_iter()`). Since the particular attributes returned for an entity differ between datasets (e.g., some provide only an ID and text for a document, while others also include a title field), type definitions can be accessed via `ds.{entity}_cls()`. The type definitions include type annotations for each field, and try to adhere to conventions when possible (e.g., the ID of documents is the first field and named `doc_id`).

The iterator approach is versatile. In some cases, it is only necessary to operate over a single entity at a time, minimizing the memory overhead. In other cases, particularly in neural networks, operations happen in batches, which can also be accomplished trivially through an iterator. And finally, in cases where all data needs

to be loaded, all entities can be easily loaded, e.g., by passing the iterator into the Python list constructor, or the dataframe constructor in Pandas [65].

Some datasets provide other entity types, such as sample document rankings or training sequences. For the former, we have a `scoreddocs` entity type, which by default is a tuple containing a query ID, a document ID, and a score. For the latter, we have a `docpairs` entity, which consists of a query and a pair of contrasting document IDs (e.g., one relevant and one non-relevant).

2.3 Command Line Interface

`ir_datasets` also provides a Command Line Interface (CLI) for performing basic operations over supported datasets. This is helpful for integration with tools not written in Python, or simply for ad-hoc data exploration. The primary operations of the CLI are `export` (corresponding to Python’s `dataset.*_iter()` functions) and `lookup` (corresponding to Python’s `docstore.get_many_iter()`). Examples of these operations are shown in right-hand side of Figure 1. The command line interface supports multiple output formats, including TSV and JSON lines. The output fields can also be specified, if only certain data is desired.

2.4 Data Acquisition

When possible, `ir_datasets` downloads content automatically from the original public sources as needed. In cases where a data usage agreement exists, the user is notified before the file is downloaded. The download process is robust; it verifies the integrity of the downloaded content via a hash and is resilient to interrupted downloads by re-issuing the request if the connection is broken (using Range HTTP requests, if supported by the server). Further, the access to and integrity of downloadable content is automatically checked periodically using a continuous integration job so that if access to some resources are lost (e.g., a file is moved) the problem can be quickly investigated and fixed. There are nearly 350 downloadable files supporting the current datasets in `ir_datasets`, each validated weekly.

Some data are not publicly available. For instance, due to its size, the ClueWeb 2009 and 2012 collections (used for tasks like the TREC WebTrack and NTCIR WWW tasks) are obtained via hard drives. Other datasets, like the Arabic Newswire collection (used for the TREC Arabic tasks) contain copyrighted material and are only available with a usage agreement and subscription to the Linguistic Data Consortium. In these cases, the user is presented

with instructions on how to acquire the dataset and where to put it. Once acquired by the user, `ir_datasets` will take care of any remaining processing. There are currently 12 document collections that require a manual process to acquire.

2.5 Supported datasets

`ir_datasets` supports a wide variety of datasets (see Table 1). These include some of the most popular evaluation benchmarks (e.g., TREC Robust [83]), large-scale shallow datasets (e.g., MS-MARCO [66]), biomedical datasets (e.g., TREC CDS [71, 72, 77]), multi- and cross-lingual datasets (e.g., TREC Arabic [33, 34]), a content-based weak supervision dataset (NYT [57]), a large-scale click dataset (ORCAS [21]), and a ranking benchmark suite (BEIR [81]). To our knowledge, this represents the largest collection and variety of IR datasets supported by any tool. To facilitate experiments with custom datasets, the Python API provides an easy mechanism to build a dataset object from files that use simple data formats:

```
ds = ir_datasets.create_dataset(docs_tsv="path/docs.tsv",
                               queries_tsv="path/queries.tsv",
                               qrels_trec="path/qrels")
```

2.6 Document lookups

It is a common task to look up documents by their ID. For instance, when training or running a neural IR model, it is often necessary to fetch the text of the current document to perform processing. Another example would be a researcher who is looking into cases in which their model fails may want to see the text of the offending documents.

One option is to load all documents into an in-memory hashmap. This may be appropriate in some cases, such a long-running process where the large upfront cost is negligible and memory is plentiful (enough for the entire collection). Building an in-memory hashmap for a collection is trivial with the Python interface:

```
doc_map = {doc.doc_id: doc for doc in dataset.docs_iter()}
```

To support other cases, `ir_datasets` provides a `docs_store` API that simplifies the process of looking up documents from disk. This API supports fetching individual or multiple documents by their ID:

```
docs_store = dataset.docs_store()
docs_store.get_many(['D1', 'D2'])
# {'D1': GenericDoc('D1', ...), 'D2': GenericDoc('D2', ...)}
it = docs_store.get_many_iter(['D1', 'D2'])
# An iterator of D1 and D2 (order not guaranteed)
```

The implementation of `docs_store()` varies based on the dataset. For many small datasets (those with up to a few million documents), we build a specialized lookup structure for the entire collection on disk as needed. A specialized structure was built for this package to provide a good trade-off between lookup speed and storage costs. All documents are compressed using lz4 and stored in sequence. A separate sorted document ID and corresponding index offset structure is also built on disk. Although simple, we found that this structure enables lookups that exceed the performance of leading indexes and databases (see Table 3). In this experiment, we used the metadata lookup functionality of Anserini [91] and Terrier [63] and key-value storage with SQLite and MongoDB. The average duration was computed per query for TREC DL 2019 passage task [26] (with the official set of reranking documents), and for TREC COVID

Table 3: Document lookup benchmarks on small datasets.

System	Time/query			Size
	HDD	SSD	Warm	
msmarco-passage/trec-dl-2019 (avg. 949 docs/query)				
<code>ir_datasets</code>	2.34 s	66 ms	7 ms	2.8 GB
MongoDB	3.62 s	130 ms	14 ms	2.7 GB
SQLite	3.72 s	94 ms	27 ms	4.1 GB
Pyserini	2.34 s	85 ms	51 ms	2.4 GB
PyTerrier	3.40 s	138 ms	68 ms	2.3 GB
cord19/fulltext/trec-covid (avg. 1,386 docs/query)				
<code>ir_datasets</code>	1.19 s	0.11 s	36 ms	1.3 GB
MongoDB	3.65 s	0.19 s	65 ms	1.8 GB
SQLite	5.99 s	0.19 s	50 ms	2.8 GB
Pyserini	2.05 s	0.32 s	51 ms	1.5 GB
PyTerrier	3.72 s	1.70 s	1,620 ms	4.2 GB

Table 4: Document lookup benchmarks on large datasets. Storage costs are listed as space beyond the source files.

Strategy	Time/query		Size
	HDD	Warm	
clueweb12/trec-web-2014 (avg. 289 docs/query)			
<code>ir_datasets</code>	44.4 s	14 ms	4.5 GB
(w/o checkpoints)	369.4 s	14 ms	0.3 GB
Pyserini	19.7 s	1,210 ms	6,041.5 GB
tweets2013-ia/trec-mb-2013 (avg. 1,188 docs/query)			
<code>ir_datasets</code>	23.3 s	24 ms	120 GB
Pyserini	17.6 s	115 ms	323 GB

complete [84] (using the judged documents). We also find that the storage cost is reasonable, with a total storage size comparable to MongoDB for the MS-MARCO passage collection and smaller than all others for the CORD19 collection.

For large collections, it is impractical and undesirable to make a copy of all documents. For instance, the ClueWeb09 and ClueWeb12 collections (for TREC Web Track) are several TB in size, even when heavily compressed. Luckily, for these datasets, their directory structure mimics the structure of the document IDs, which allows the desired source file containing a given document ID to be easily identified. To speed up lookups within a given file, we use `zlib-state`⁴ to take periodic checkpoints of the zlib decoding state of the source files. This eliminates the need to read all the source file contents up to the desired document and greatly speeds up lookups of documents that appear late in the source files. The pre-built checkpoints are automatically downloaded and used when appropriate. Furthermore, we cache fetched documents on disk for even faster subsequent lookups. Different approaches are taken for other large collections, such as Tweets2013-ia [76] (for the TREC Microblog task [51, 52]). See Table 4 for a comparison between document

⁴<https://github.com/seanmacavaney/zlib-state>

lookup times using `ir_datasets` and Pyserini (from stored document source). Even though `ir_datasets` is slower than Pyserini on the first lookup, the cache greatly speeds up subsequent fetches (see “Warm”). Since experiments in neural IR frequently only work with a small subset of documents, this is very beneficial for these pipelines. We also observe that the checkpoint files for ClueWeb12 speed up lookups considerably, without adding much overhead in terms of storage; since Anserini keeps a copy of all documents, it accumulates around 6TB of storage overhead, compared to 4.5GB using `ir_datasets`. Note that the other approaches explored in Table 1 would accumulate similar storage overheads, as they also copy the data. Tweets2013-ia accumulates considerable storage costs, as the source hierarchy is not conducive to document lookups. In this case, `ir_datasets` builds an ID-based lookup file hierarchy.

2.7 Fancy slicing

In many cases, it is beneficial to select a segment of a document collection. For instance, some techniques involve pre-computing neural document representations to speed up reranking [56] or for performing first-stage retrieval [49]. In this case, dividing the operation over multiple GPUs or machines can yield substantial speed gains, as the process is embarrassingly parallel. To divide up the work, it is helpful to be able to select ranges of the document collection for processing.

The Python standard library `islice` function is not ideal for this task because I/O and processing of documents would be performed for skipped indices. Instead, all objects returned from `doc_iter` can themselves be sliced directly. The implementation of the slicing depends on the particular dataset, but in all implementations avoid unnecessary I/O and processing by seeking to the appropriate location in the source file. This *fancy slicing* implementation mostly follows typical Python slicing semantics, allow for different workers to be assigned specific ranges of documents:

```
dataset.docs_iter()[:10] # the first 10 docs
dataset.docs_iter()[-10:] # the last 10 docs
dataset.docs_iter()[100:110] # 10 docs starting at index 100
dataset.docs_iter()[3::5] # every 5 docs, starting at index 3
dataset.docs_iter()[:1/3] # the first third of the collection
```

2.8 Documentation

Documentation about datasets are available from the `ir_datasets` catalog.⁴ An overview list shows all available datasets and their capabilities (Figure 2). The documentation page for each individual dataset includes a brief description, relevant links (e.g., to shared task website and paper), supported relations, citations, and code samples. An example is shown in Figure 3 for the TREC COVID dataset [84].

2.9 Automated Testing

`ir_datasets` includes several suites of automated tests to ensure the package works as expected, functionality does not regress as changes are made, and to ensure that downloaded content remains available and unchanged. The automated testing suite includes include unit tests, integration/regression tests, and tests to ensure downloadable content remains available and unchanged.

Dataset	docs	queries	qrels
antique	✓		
antique/test	✓	✓	✓
antique/test/non-offensive	✓	✓	✓
antique/train	✓	✓	✓
antique/train/split200-train	✓	✓	✓
antique/train/split200-valid	✓	✓	✓
aquaint	⚠		
aquaint/trec-robust-2005	⚠	✓	✓

docs available from LDC

Figure 2: Example from the `ir_datasets` catalog. Users can easily check which datasets are available for automatic downloads (green checkbox) and those that require obtaining data from a third party (yellow triangle).

"cord19/trec-covid"

The TREC COVID collection. Queries related to COVID-19, including deep relevance judgments.

- [Shared task site](#)
- [Shared task paper](#)

queries | docs | qrels | Citation

Language: **en**

Query type:

```
TreeQuery: (namedtuple)
[0] query_id: str
[1] title: str
[2] description: str
[3] narrative: str
```

Example

```
import ir_datasets
dataset = ir_datasets.load('cord19/trec-covid')
for query in dataset.queries_iter():
    query # namedtuple=query_id, title,
    description, narrative>
```

Figure 3: Example documentation for `cord19/trec-covid`.

3 INTEGRATION WITH OTHER TOOLS

The CLI makes `ir_datasets` easy to use with various tools (e.g., the PISA engine [61] can index using the document export). However, deeper integration can provide further functionality, as we demonstrate in this section with four tools. Note that `ir_datasets` does not depend on any of these tools; instead they use `ir_datasets`.

Capreolus [93] is a toolkit for training and evaluating neural learning-to-rank models through Python and command line interfaces. In terms of data, it includes components for “collections” (sets of documents) and “benchmarks” (sets of queries and qrels). Though it has some built-in datasets, it also supports all datasets available from `ir_datasets` in its pipelines:

```
import capreolus as cap
collection, benchmark = cap.get_irds("pmc/v2/trec-cds-2016",
    fields=["abstract"], query_type="summary")
index = cap.AnseriniIndex({"stemmer": None}, collection)
index.create_index()
benchmark.qrels
benchmark.queries
```

PyTerrier [58] is a Python interface to the Terrier search engine [63] that enables the creation of flexible retrieval pipelines. It has a native dataset API, but it now also automatically adds all datasets from `ir_datasets`, expanding the number of available

datasets. They can be accessed via the dataset ID with an `irds:` prefix, and then used seamlessly with the rest of PyTerrier:

```
import pyterrier as pt
pt.init()
ds = pt.get_dataset('irds:cord19/trec-covid')
indexer = pt.index.IterDictIndexer('./cord19')
indexer.index(ds.get_corpus_iter(), fields=('abstract',))
topics = ds.get_topics(variant="description")
qrels = ds.get_qrels()
```

OpenNIR [55] provides a command line neural reranking pipeline for several standard IR benchmarks. OpenNIR supports `ir_datasets` for its training, validation, and testing dataset components. Queries and qrels are trivially fed into the training and validation processes. Documents are automatically indexed with Anserini for first-stage retrieval, and document lookups are used to fetch the text when training and scoring. Here is an example testing on the TREC COVID dataset:

```
$ scripts/pipeline.sh test_ds=irds test_ds.ds=cord19/trec-covid
```

Anserini [91], and its Python-wrapper counterpart **Pyserini** [53] focus on reproducibility in IR. They provide a wrapper and suite of tools around a Lucene index. As such, operations on datasets in this tool are tightly coupled with the Lucene and Anserini packages. Though it has support for a wide variety of query and relevance assessments (distributed with the package), the support for document content is sparse, since only a few collections have automatically-downloadable indices. The remainder rely on manual instructions. Queries and qrels from `ir_datasets` can be used with Anserini by using the export CLI (as TSV or TREC format). The CLI can also efficiently output documents in a format it can index in parallel:

```
$ ir_datasets doc_fifos medline/2017
# To index with Anserini, run:
# IndexCollection -collection JsonCollection -input
# /tmp/tmp6sop5gr -threads 23 -index <your_index_path>
```

DiffIR [46] is a tool that enables the visualization and qualitative comparison of search results. Using `ir_datasets`, it shows the textual content of the top results for queries and highlights model-specific impactful text spans.

4 COMMUNITY CONTRIBUTIONS

We welcome (and encourage) community contributions. Extending `ir_datasets` as a separate package is straightforward,⁵ and we also welcome pull requests to the main package.

To maintain quality in `ir_datasets`, we require considerations of ease-of-use, efficiency, data integrity, and documentation. We request that issues are opened before implementation to ensure proper consideration of these aspects. `ir_datasets` provides tools for handling typical data formats (e.g., TREC, TSV, CSV), making the process relatively straightforward. Atypical formats likely require special processing. There are plenty of examples to help guide the contributor.

5 FUTURE APPLICATIONS

We envision `ir_datasets` enabling a variety of useful applications.

Training/evaluation in private settings. This tool could facilitate experiments and tasks that involve keeping data private.

⁵example: <https://github.com/seanmacavaney/dummy-irds-ext>

This is a realistic setting in several circumstances. For instance, a shared task involving searching through clinical notes would likely face challenges distributing this collection due to patient privacy concerns. Or a company may want to offer a shared task using a proprietary document collection or query log. In both these cases, a version of `ir_datasets` could be built that provides this data that is only available in a secure environment (e.g., one where networking is disabled). Participants could feel confident that their code is processing the data correctly, given that it supports the `ir_datasets` API; their code can switch to this dataset simply by using the dataset ID of the private dataset.

Dataset exploration GUI. Performing ad-hoc data analysis using `ir_datasets` is an improvement over prior approaches. The user experience could be further improved through a graphical user interface that facilitate common dataset exploration tasks. For instance, this tool could graphically present the list of queries and link to the text of judged documents. Though this functionality is easy through the Python and command line interfaces, a graphical interface would further reduce friction and ease exploration.

6 CONCLUSION

We presented `ir_datasets`, a tool that provides access to a variety of datasets and benchmarks for search engines. The tool automatically downloads and verifies content when possible, to aid in reproducibility. Through Python and command-line interfaces, users can iterate over documents, queries, and relevance judgments, and perform lookups of documents by ID. The utility of these functionalities are demonstrated through integration with several tools for performing IR experiments. The `ir_datasets` catalog can help users discover datasets and acts as a hub of information with links and citations to relevant literature. We hope that `ir_datasets` reduces researcher burden, helps reduce redundant copies of datasets across toolkits, and enables the creation of new tools.

ACKNOWLEDGMENTS

We thank those who contributed to issues or discussions about `ir_datasets`. This was funded in part by the ARCS Foundation.

REFERENCES

- [1] [n.d.]. *Cranfield collection*. http://ir.dcs.gla.ac.uk/resources/test_collections/cran/
- [2] [n.d.]. *NPL collection*. http://ir.dcs.gla.ac.uk/resources/test_collections/npl/
- [3] [n.d.]. *TensorFlow Datasets, A collection of ready-to-use datasets*. <https://www.tensorflow.org/datasets>
- [4] Mustafa Abualsaud, Christina Lioma, Maria Maistro, Mark D. Smucker, and Guido Zuccon. 2019. Overview of the TREC 2019 Decision Track. In *TREC*.
- [5] James Allan, Javed A. Aslam, Ben Carterette, Virgil Pavlu, and Evangelos Kanoulas. 2008. Million Query Track 2008 Overview. In *TREC*.
- [6] James Allan, Ben Carterette, Javed A. Aslam, Virgil Pavlu, Blagovest Dachev, and Evangelos Kanoulas. 2007. Million Query Track 2007 Overview. In *TREC*.
- [7] James Allan, Donna Harman, Evangelos Kanoulas, Dan Li, Christophe Van Gysel, and Ellen Voorhees. 2017. TREC 2017 Common Core Track Overview. In *TREC*.
- [8] Alexander Bondarenko, Maik Fröbe, Meriem Beloucif, Lukas Gienapp, Yamen Ajjour, Alexander Panchenko, Christian Biemann, Benno Stein, Henning Wachsmuth, Martin Potthast, and Matthias Hagen. 2020. Overview of Touché 2020: Argument Retrieval. In *CLEF*.
- [9] Vera Boteva, Demian Gholipour, Artem Sokolov, and Stefan Riezler. 2016. A Full-Text Learning to Rank Dataset for Medical Information Retrieval. In *Proceedings of the European Conference on Information Retrieval (ECIR)* (Padova, Italy). Springer.
- [10] Stefan Büttcher, Charles L. A. Clarke, and Ian Soboroff. 2006. The TREC 2006 Terabyte Track. In *TREC*.
- [11] Ben Carterette, Virgil Pavlu, Hui Fang, and Evangelos Kanoulas. 2009. Million Query Track 2009 Overview. In *TREC*.

- [12] Charles L. A. Clarke, Falk Scholer, and Ian Soboroff. 2005. The TREC 2005 Terabyte Track. In *TREC*.
- [13] Charles Clarke, Nick Craswell, and Ian Soboroff. 2004. Overview of the TREC 2004 Terabyte Track. In *TREC*.
- [14] Charles L. A. Clarke, Nick Craswell, and Ian Soboroff. 2009. Overview of the TREC 2009 Web Track. In *TREC*.
- [15] Charles L. A. Clarke, Nick Craswell, Ian Soboroff, and Gordon V. Cormack. 2010. Overview of the TREC 2010 Web Track. In *TREC*.
- [16] Charles L. A. Clarke, Nick Craswell, Ian Soboroff, and Ellen M. Voorhees. 2011. Overview of the TREC 2011 Web Track. In *TREC*.
- [17] Charles L. A. Clarke, Nick Craswell, and Ellen M. Voorhees. 2012. Overview of the TREC 2012 Web Track. In *TREC*.
- [18] Arman Cohan, Sergey Feldman, Iz Beltagy, Doug Downey, and Daniel Weld. 2020. SPECTER: Document-level Representation Learning using Citation-informed Transformers. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*. Association for Computational Linguistics, Online, 2270–2282. <https://doi.org/10.18653/v1/2020.acl-main.207>
- [19] Keayn Collins-Thompson, Paul Bennett, Fernando Diaz, Charles L. A. Clarke, and Ellen M. Voorhees. 2013. TREC 2013 Web Track Overview. In *TREC*.
- [20] Keayn Collins-Thompson, Craig Macdonald, Paul Bennett, Fernando Diaz, and Ellen M. Voorhees. 2014. TREC 2014 Web Track Overview. In *TREC*.
- [21] Nick Craswell, Daniel Campos, Bhaskar Mitra, Emine Yilmaz, and Bodo Billerbeck. 2020. ORCAS: 18 Million Clicked Query-Document Pairs for Analyzing Search. *arXiv preprint arXiv:2006.05324* (2020).
- [22] Nick Craswell and David Hawking. 2002. Overview of the TREC-2002 Web Track. In *TREC*.
- [23] Nick Craswell and David Hawking. 2004. Overview of the TREC-2004 Web Track. In *TREC*.
- [24] Nick Craswell, David Hawking, Ross Wilkinson, and Mingfang Wu. 2003. Overview of the TREC 2003 Web Track. In *TREC*.
- [25] Nick Craswell, Bhaskar Mitra, Emine Yilmaz, and Daniel Campos. 2020. Overview of the TREC 2020 deep learning track. In *TREC*.
- [26] Nick Craswell, Bhaskar Mitra, Emine Yilmaz, Daniel Campos, and Ellen Voorhees. 2019. Overview of the TREC 2019 deep learning track. In *TREC 2019*.
- [27] Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding. In *NAACL-HLT*.
- [28] Laura Dietz and Ben Gamari. 2017. TREC CAR: A Data Set for Complex Answer Retrieval. (2017). <http://trec-cars.cs.umh.edu> Version 1.5.
- [29] Laura Dietz, Manisha Verma, Filip Radlinski, and Nick Craswell. 2017. TREC Complex Answer Retrieval Overview. In *TREC*.
- [30] T. Diggelmann, Jordan L. Boyd-Graber, Jannis Bulian, Massimiliano Ciaramita, and Markus Leippold. 2020. CLIMATE-FEVER: A Dataset for Verification of Real-World Climate Claims. *ArXiv abs/2012.00614* (2020).
- [31] Jibril Frej, Didier Schwab, and Jean-Pierre Chevallet. 2020. MLWIKIR: A Python Toolkit for Building Large-scale Wikipedia-based Information Retrieval Datasets in Chinese, English, French, Italian, Japanese, Spanish and More. In *CIRCLE*.
- [32] Jibril Frej, Didier Schwab, and Jean-Pierre Chevallet. 2020. WIKIR: A Python toolkit for building a large-scale Wikipedia-based English Information Retrieval Dataset. In *LREC*.
- [33] Fredric Gey and Douglas Oard. 2001. The TREC-2001 Cross-Language Information Retrieval Track: Searching Arabic using English, French or Arabic Queries. In *TREC*.
- [34] Fredric Gey and Douglas Oard. 2002. The TREC-2002 Arabic/English CLIR Track. In *TREC*.
- [35] David Graff and Kevin Walker. 2001. Arabic Newswire Part 1 LDC2001T55. <https://catalog.ldc.upenn.edu/LDC2001T55>
- [36] Donna Harman. 1994. Overview of the Third Text REtrieval Conference (TREC-3). In *TREC*.
- [37] Donna Harman. 1995. Overview of the Fourth Text REtrieval Conference (TREC-4). In *TREC*.
- [38] Helia Hashemi, Mohammad Aliannejadi, Hamed Zamani, and Bruce Croft. 2020. ANTIQUE: A Non-Factoid Question Answering Benchmark. In *ECIR*.
- [39] Faegheh Hasibi, Fedor Nikolaev, Chenyan Xiong, K. Balog, S. E. Bratsberg, Alexander Kotov, and J. Callan. 2017. DBpedia-Entity v2: A Test Collection for Entity Search. *Proceedings of the 40th International ACM SIGIR Conference on Research and Development in Information Retrieval* (2017).
- [40] William Hersh, Aaron Cohen, Lynn Ruslen, and Phoebe Roberts. 2007. TREC 2007 Genomics Track Overview. In *TREC*.
- [41] William Hersh, Aaron Cohen, Jianji Yang, Ravi Teja Bhupatiraju, Phoebe Roberts, and Marti Hearst. 2007. TREC 2005 Genomics Track Overview. In *TREC*.
- [42] William Hersh, Aaron M. Cohen, Phoebe Roberts, and Hari Krishna Rekapalli. 2006. TREC 2006 Genomics Track Overview. In *TREC*.
- [43] William R. Hersh, Ravi Teja Bhupatiraju, Laura Ross, Phoebe Johnson, Aaron M. Cohen, and Dale F. Kraemer. 2004. TREC 2004 Genomics Track Overview. In *TREC*.
- [44] D. Hoogveen, Karin M. Verspoor, and Timothy Baldwin. 2015. CQADupStack: A Benchmark Data Set for Community Question-Answering Research. *Proceedings of the 20th Australasian Document Computing Symposium* (2015).
- [45] Hamel Husain, Ho-Hsiang Wu, Tiferet Gazit, Miltiadis Allamanis, and Marc Brockschmidt. 2019. CodeSearchNet Challenge: Evaluating the State of Semantic Code Search. *ArXiv* (2019).
- [46] Kevin Martin Jose, Thong Nguyen, Sean MacAvaney, Jeff Dalton, and Andrew Yates. 2021. DiffIR: Exploring Differences in Ranking Models' Behavior. In *Proceedings of the 44th International ACM SIGIR Conference on Research and Development in Information Retrieval*.
- [47] Mandar Joshi, Eunsol Choi, Daniel S. Weld, and Luke Zettlemoyer. 2017. TriviaQA: A Large Scale Distantly Supervised Challenge Dataset for Reading Comprehension. In *ACL*.
- [48] Vladimir Karpukhin, Barlas Oğuz, Sewon Min, Patrick Lewis, Ledell Wu, Sergey Edunov, Danqi Chen, and Wen tau Yih. 2020. Dense Passage Retrieval for Open-Domain Question Answering. *arXiv:2004.04906* [cs.CL]
- [49] Omar Khatib and Matei Zaharia. 2020. CoBERT: Efficient and Effective Passage Search via Contextualized Late Interaction over BERT. *SIGIR* (2020).
- [50] Tom Kwiatkowski, Jennimaria Palomaki, Olivia Redfield, Michael Collins, Ankur Parikh, Chris Alberti, Danielle Epstein, Illia Polosukhin, Matthew Kelcey, Jacob Devlin, Kenton Lee, Kristina N. Toutanova, Llion Jones, Ming-Wei Chang, Andrew Dai, Jakob Uszkoreit, Quoc Le, and Slav Petrov. 2019. Natural Questions: a Benchmark for Question Answering Research. *TACL* (2019).
- [51] Jimmy Lin and Miles Efron. 2013. Overview of the TREC-2013 Microblog Track. In *TREC*.
- [52] Jimmy Lin, Miles Efron, Yulu Wang, and Garrick Sherman. 2014. Overview of the TREC-2014 Microblog Track. In *TREC*.
- [53] Jimmy Lin, Xueguang Ma, Sheng-Chieh Lin, Jheng-Hong Yang, Ronak Pradeep, and Rodrigo Nogueira. 2021. Pyserini: An Easy-to-Use Python Toolkit to Support Replicable IR Research with Sparse and Dense Representations. *ArXiv abs/2102.10073* (2021).
- [54] Cheng Luo, Tetsuya Sakai, Yiqun Liu, Zhicheng Dou, Chenyan Xiong, and Jingfang Xu. 2017. Overview of the NTCIR-13 We Want Web Task. In *NTCIR*.
- [55] Sean MacAvaney. 2020. OpenNIR: A Complete Neural Ad-Hoc Ranking Pipeline. In *Proceedings of the Thirteenth ACM International Conference on Web Search and Data Mining*. 845–848. <https://doi.org/10.1145/3336191.3371864>
- [56] 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*. 49–58. <https://doi.org/10.1145/3397271.3401093>
- [57] Sean MacAvaney, Andrew Yates, Kai Hui, and Ophir Frieder. 2019. Content-Based Web Supervision for Ad-Hoc Re-Ranking. In *SIGIR*.
- [58] Craig Macdonald and Nicola Tonello. 2020. Declarative Experimentation in Information Retrieval using PyTerrier. In *Proceedings of ICTIR 2020*.
- [59] Iain Mackie, Jeffrey Dalton, and Andrew Yates. 2021. How Deep is your Learning: the DL-HARD Annotated Deep Learning Dataset. In *Proceedings of the 44th International ACM SIGIR Conference on Research and Development in Information Retrieval*.
- [60] Macedo Maia, S. Handschuh, A. Freitas, Brian Davis, R. McDermott, M. Zarrouk, and A. Balahur. 2018. WWW'18 Open Challenge: Financial Opinion Mining and Question Answering. *Companion Proceedings of the The Web Conference 2018* (2018).
- [61] Antonio Mallia, Michal Siedlaczek, Joel Mackenzie, and Torsten Suel. 2019. PISA: Performant Indexes and Search for Academia. In *Proceedings of the Open-Source IR Replicability Challenge co-located with 42nd International ACM SIGIR Conference on Research and Development in Information Retrieval, OSIRRC@SIGIR 2019, Paris, France, July 25, 2019*. 50–56. <http://ceur-ws.org/Vol-2409/docker08.pdf>
- [62] Jiaxin Mao, Tetsuya Sakai, Cheng Luo, Peng Xiao, Yiqun Liu, and Zhicheng Dou. 2018. Overview of the NTCIR-14 We Want Web Task. In *NTCIR*.
- [63] Iadh Ounis, Gianni Amati, Vassilis Plachouras, Ben He, Craig Macdonald, and Douglas Johnson. 2005. Terrier information retrieval platform. In *ECIR*. Springer, 517–519.
- [64] Joao Palotti, Guido Zuccon, Jimmy, Pavel Pecina, Mihai Lupu, Lorraine Goeuriot, Liadh Kelly, and Allan Hanbury. 2017. CLEF 2017 Task Overview: The IR Task at the eHealth Evaluation Lab - Evaluating Retrieval Methods for Consumer Health Search. In *CLEF*.
- [65] The pandas development team. 2020. *pandas-dev/pandas: Pandas*. <https://doi.org/10.5281/zenodo.3509134>
- [66] Nick Craswell Li Deng Jianfeng Gao Xiaodong Liu Rangan Majumder Andrew McNamara Bhaskar Mitra Tri Nguyen Mir Rosenberg Xia Song Alina Stoica Saurabh Tiwary Tong Wang Payal Bajaj, Daniel Campos. 2016. MS MARCO: A Human Generated Machine Reading Comprehension Dataset. In *InCoCo@NIPS*.
- [67] Navid Rekabsaz, Oleg Lesota, Markus Schedl, Jon Brassey, and Carsten Eickhoff. 2021. TripClick: The Log Files of a Large Health Web Search Engine. In *SIGIR*.
- [68] Kirk Roberts, Dina Demner-Fushman, Ellen Voorhees, William R. Hersh, Steven Bedrick, and Alexander J. Lazar. 2018. Overview of the TREC 2018 Precision Medicine Track. In *TREC*.
- [69] Kirk Roberts, Dina Demner-Fushman, Ellen Voorhees, William R. Hersh, Steven Bedrick, Alexander J. Lazar, and Shubham Pant. 2017. Overview of the TREC

- 2017 Precision Medicine Track. In *TREC*.
- [70] Kirk Roberts, Dina Demner-Fushman, Ellen Voorhees, William R. Hersh, Steven Bedrick, Alexander J. Lazar, Shubham Pant, and Funda Meric-Bernstam. 2019. Overview of the TREC 2019 Precision Medicine Track. In *TREC*.
- [71] Kirk Roberts, Dina Demner-Fushman, Ellen M. Voorhees, and William R. Hersh. 2016. Overview of the TREC 2016 Clinical Decision Support Track. In *TREC*.
- [72] Kirk Roberts, Matthew S. Simpson, Ellen Voorhees, and William R. Hersh. 2015. Overview of the TREC 2015 Clinical Decision Support Track. In *TREC*.
- [73] Willie Rogers. 2000. TREC Mandarin LDC2000T52. <https://catalog ldc.upenn.edu/LDC2000T52>
- [74] Willie Rogers. 2000. TREC Spanish LDC2000T51. <https://catalog ldc.upenn.edu/LDC2000T51>
- [75] Evan Sandhaus. 2008. The new york times annotated corpus. *Linguistic Data Consortium, Philadelphia* 6, 12 (2008), e26752.
- [76] Royal Sequiera and Jimmy Lin. 2017. Finally, a Downloadable Test Collection of Tweets. In *SIGIR*.
- [77] Matthew S. Simpson, Ellen M. Voorhees, and William Hersh. 2014. Overview of the TREC 2014 Clinical Decision Support Track. In *TREC*.
- [78] Alan Smeaton and Ross Wilkinson. 1996. Spanish and Chinese Document Retrieval in TREC-5. In *TREC*.
- [79] Ian Soboroff, Shudong Huang, and Donna Harman. 2018. TREC 2018 News Track Overview. In *TREC*.
- [80] Ian Soboroff, Shudong Huang, and Donna Harman. 2019. TREC 2019 News Track Overview. In *TREC*.
- [81] Nandan Thakur, Nils Reimers, Andreas Rücklé, Abhishek Srivastava, and Iryna Gurevych. 2021. BEIR: A Heterogenous Benchmark for Zero-shot Evaluation of Information Retrieval Models. *arXiv preprint arXiv:2104.08663* (4 2021). <https://arxiv.org/abs/2104.08663>
- [82] James Thorne, Andreas Vlachos, Christos Christodoulopoulos, and Arpit Mittal. 2018. FEVER: a Large-scale Dataset for Fact Extraction and VERification. In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long Papers)*. Association for Computational Linguistics, New Orleans, Louisiana, 809–819. <https://doi.org/10.18653/v1/N18-1074>
- [83] Ellen Voorhees. 2004. Overview of the TREC 2004 Robust Retrieval Track. In *TREC*.
- [84] E. Voorhees, Tasmee Alam, Steven Bedrick, Dina Demner-Fushman, W. Hersh, Kyle Lo, Kirk Roberts, I. Soboroff, and Lucy Lu Wang. 2020. TREC-COVID: Constructing a Pandemic Information Retrieval Test Collection. *ArXiv abs/2005.04474* (2020).
- [85] Ellen M. Voorhees. 2005. Overview of the TREC 2005 Robust Retrieval Track. In *TREC*.
- [86] Henning Wachsmuth, Shahbaz Syed, and Benno Stein. 2018. Retrieval of the Best Counterargument without Prior Topic Knowledge. In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)* (Melbourne, Australia). Association for Computational Linguistics, 241–251. <http://aclweb.org/anthology/P18-1023>
- [87] David Wadden, Shanchuan Lin, Kyle Lo, Lucy Lu Wang, Madeleine van Zuylen, Arman Cohan, and Hannaneh Hajishirzi. 2020. Fact or Fiction: Verifying Scientific Claims. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*. Association for Computational Linguistics, Online, 7534–7550. <https://doi.org/10.18653/v1/2020.emnlp-main.609>
- [88] Lucy Lu Wang, Kyle Lo, Yoganand Chandrasekhar, Russell Reas, Jiangjiang Yang, Darrin Eide, K. Funk, Rodney Michael Kinney, Ziyang Liu, W. Merrill, P. Mooney, D. Murdick, Devvret Rishi, Jerry Sheehan, Zhihong Shen, B. Stilson, A. Wade, K. Wang, Christopher Wilhelm, Boya Xie, D. Raymond, Daniel S. Weld, Oren Etzioni, and Sebastian Kohlmeier. 2020. CORD-19: The Covid-19 Open Research Dataset. *ArXiv* (2020).
- [89] Ross Wilkinson. 1997. Chinese Document Retrieval at TREC-6. In *TREC*.
- [90] Thomas Wolf, Quentin Lhoest, Patrick von Platen, Yacine Jernite, Mariama Drame, Julien Plu, Julien Chaumond, Clement Delangue, Clara Ma, Abhishek Thakur, Suraj Patil, Joe Davison, Teven Le Scao, Victor Sanh, Canwen Xu, Nicolas Patry, Angie McMillan-Major, Simon Brandeis, Sylvain Gugger, François Lagunas, Lysandre Debut, Morgan Funtowicz, Anthony Moi, Sasha Rush, Philipp Schmid, Pierric Cistac, Victor Muštar, Jeff Boudier, and Anna Tordjmann. 2020. Datasets. *GitHub*. Note: <https://github.com/huggingface/datasets> 1 (2020).
- [91] Peilin Yang, Hui Fang, and Jimmy Lin. 2017. Anserini: Enabling the Use of Lucene for Information Retrieval Research. *SIGIR* (2017).
- [92] Zhilin Yang, Peng Qi, Saizheng Zhang, Yoshua Bengio, William Cohen, Ruslan Salakhutdinov, and Christopher D. Manning. 2018. HotpotQA: A Dataset for Diverse, Explainable Multi-hop Question Answering. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*. Association for Computational Linguistics, Brussels, Belgium, 2369–2380. <https://doi.org/10.18653/v1/D18-1259>
- [93] Andrew Yates, Kevin Martin Jose, Xinyu Zhang, and Jimmy Lin. 2020. Flexible IR pipelines with Capreolus. In *Proceedings of the 29th ACM International Conference on Information & Knowledge Management*. 3181–3188.
- [94] Guido Zuccon, Joao Palotti, Lorraine Goeuriot, Liadh Kelly, Mihai Lupu, Pavel Pecina, Henning Müller, Julie Budaher, and Anthony Deacon. 2016. The IR Task at the CLEF eHealth Evaluation Lab 2016: User-centred Health Information Retrieval. In *CLEF*.