

# Neural Code Search Evaluation Dataset

Hongyu Li  
Facebook, Inc.  
U.S.A  
hongyul@fb.com

Seohyun Kim  
Facebook, Inc.  
U.S.A  
skim131@fb.com

Satish Chandra  
Facebook, Inc.  
U.S.A  
satch@fb.com

## Abstract

There has been an increase of interest in code search using natural language. Assessing the performance of such code search models can be difficult without a readily available evaluation suite. In this paper, we present an evaluation dataset of natural language query and code snippet pairs for future work. We also provide the results of two code search models ([6] and [1]) from recent work as a benchmark.

## 1 Introduction

In recent years, learning the mapping between natural language and code snippets has been a popular field of research. In particular, [6], [1], [2] have explored finding relevant code snippets given a natural language query, with the models varying from using word embeddings and IR techniques to using sophisticated neural networks. To evaluate the performance of these models, Stack Overflow questions and code answer pairs are prime candidates, as Stack Overflow questions well resemble what a developer may ask. Such an example is "Close/hide the Android Soft Keyboard".<sup>1</sup> One of the first answers<sup>2</sup> on Stack Overflow correctly answers this question. However, collecting these questions can be tedious, and systematically comparing various models can pose a challenge.

To this end, we have constructed an evaluation dataset, which contains natural language queries and relevant code snippet answers from Stack Overflow. It also includes code snippet examples from the search corpus (public repositories from GitHub) that correctly answers each query. We hope that this dataset can be served as a benchmark to evaluate performance across various code search models.

The paper is organized as follows. First we will explain what data we are releasing in the dataset. Then we will describe the process for obtaining this dataset. Finally, we will evaluate two code search models of our own creation, NCS and UNIF, on the evaluation dataset as a benchmark.

## 2 Dataset Contents

In this section, we explain what data we are releasing.

### 2.1 GitHub Repositories

The most popular Android repositories on GitHub (ranked by the number of stars) is used to create the search corpus.

<sup>1</sup> <https://stackoverflow.com/questions/1109022/close-hide-the-android-soft-keyboard>

<sup>2</sup> <https://stackoverflow.com/a/1109108>

```
{
  "id": 4716813,
  "filepath": "Mindgames/VideoStreamServer/playersdk/src
/main/java/com/kaltura/playersdk/
PlayerViewController.java",
  "method_name": "notifyKPlayerEvent",
  "start_line": 506,
  "end_line": 566,
  "url": "https://github.com/Mindgames/VideoStreamServer
/blob/b7c73d2bcd296b3a24f83cf67d6a5998c7a1af6b/
playersdk/src/main/java/com/kaltura/playersdk/
PlayerViewController.java#L506-L566"
}
```

Listing 2. Search corpus example.

For each repository that we indexed, we provide the link, specific to the commit that was used.<sup>3</sup> In total, there are 24,549 repositories.<sup>4</sup> We will release a text file containing the download links for these GitHub repositories. See Listing 1 for an example.

### 2.2 Search Corpus

The search corpus is indexed using all method bodies parsed from the 24,549 GitHub repositories. In total, there are 4,716,814 methods in this corpus. The code search model will find relevant code snippets (i.e. method bodies) from this corpus given a natural language query. In this data release, we will provide the following information for each method in the corpus:

- **id:** Each method in the corpus has a unique numeric identifier. This ID number will also be referenced in our evaluation dataset.
- **filepath:** The file path is in the format of  
`:owner/:repo/relative-file-path-to-the-repo`
- **method\_name**
- **start\_line:** Starting line number of the method in the file.
- **end\_line:** Ending line number of the method in the file.
- **url:** GitHub link to the method body with commit ID and line numbers encoded.

Listing 2 provides an example of a method in the search corpus.

<sup>3</sup> From August 2018

<sup>4</sup> There were originally 26,109 repositories - the difference is due to reasons outside of our control (e.g. repositories getting deleted). Note that not all of the links in this dataset may not always be available in the future for the similar reasons.

```

https://github.com/00-00-00/ably-chat/archive/9bb2e36acc24f1cd684ef5d1b98d837055ba9cc8.zip
https://github.com/01sadra/Detoxiom/archive/c3fffd36989b0cd93bd09cbaa35123b9d605f989.zip
https://github.com/041lameya/MPG_update/archive/27ac5531ca2c2f123e0cb854ebcb4d0441e2bc98.zip
https://github.com/0508994/MinesweeperGO/archive/ba0e0e45d2da21dde2365ce09277aad511de6885.zip
https://github.com/07101994/My-PPT-Presentation/archive/b89b17a962d5c3e5682fa751228a9f9ca593d77b.zip
https://github.com/0912718/ICT-lab/archive/d1d723edb722013cc83761f0f9df252cfd3361c3.zip
https://github.com/0Cubed/ZeroMediaPlayer/archive/d84c675f9dc8b16f823bb252db9ee368fbd5cd8e.zip
...

```

**Listing 1.** GitHub repositories download links example.

```

{
  "stackoverflow_id": 1109022,
  "question": "Close/hide the Android Soft Keyboard",
  "question_url": "https://stackoverflow.com/questions/1109022/close-hide-the-android-soft-keyboard",
  "answer_url": "https://stackoverflow.com/a/1109108",
  "examples": [1841045, 1800067, 1271795],
  "examples_url": [
    "https://github.com/alextselegidis/easyappointments-android-client/blob/39f1e8763eb80e8ca97e070b0e9333c1cd386d90/app/src/main/java/com/easyappointments/LoginActivity.java#L102-L106",
    "https://github.com/zelloptt/zello-android-client-sdk/blob/87b45b6...",
    "https://github.com/systems/conference-android/blob/a67982abf54e0..."
  ]
}

```

**Listing 3.** Evaluation dataset example.

### 2.3 Evaluation Dataset

The evaluation dataset is composed of 287 Stack Overflow question and answer pairs, for which we release the following information:

- **stackoverflow\_id**: StackOverflow post ID.
- **question**: Title of the StackOverflow post.
- **question\_url**: URL of the StackOverflow post.
- **answer\_url**: URL of the StackOverflow answer to the question. The code snippets in the answer were extracted and compared with in our evaluation pipeline.
- **examples**: 3 methods from the search corpus that best answer the question (most similar to the Stack Overflow answer).
- **examples\_url**: GitHub links to the examples.

Note that there may be more acceptable answers to each question. See Listing 3 for a concrete example of an evaluation question in this dataset.

### 2.4 NCS / UNIF Score Sheet

We provide the evaluation results for two code search models of our creation, each with two variations:

- **NCS**: an unsupervised model which uses word embedding derived directly from the search corpus[6].
- **NCS<sub>postrank</sub>**: an extension of the base NCS model that performs a post-pass ranking, as explained in [6].
- **UNIF<sub>android</sub>**, **UNIF<sub>stackoverflow</sub>**: a supervised extension of the NCS model that uses a bag-of-words-based neural network with attention. The supervision is learned

using *GitHub-Android-Train* and *StackOverflow-Android-Train* datasets, respectively, as described in [1].

We provide the rank of the first correct answer (FRank) for each question in our evaluation dataset. The score sheet is saved in a comma-delimited csv file as illustrated in Listing 4.

## 3 How we Obtained the Dataset

In this section, we describe the procedure for how we obtained the data.

**GitHub repositories.** We obtained the information of the GitHub repositories with the GitHub REST API [3], and the source files were downloaded using publicly available links.

**Search corpus.** The search corpus was obtained by dividing each file in the GitHub repositories by method-level granularity.

**Evaluation dataset.** The benchmark questions were collected from a data dump publicly released by Stack Exchange [4]. To select the set of Stack Overflow question and answer pairs, we created a heuristics-based filtering pipeline where we discarded open-ended, discussion-style questions. We first obtained the most popular 17,000 questions on Stack Overflow with “Android” and “Java” tags. The dataset is further filtered with the following criteria: 1) there exists an upvoted code answer, 2) the ground truth code snippet has at least one match in the search corpus. From this pipeline, we were able to obtain 518 questions. Finally, we manually went through these questions and filtered out questions with vague queries and/or code answers. The final dataset contains 287 Stack Overflow question and answers pairs.

**NCS / UNIF score sheet.** To judge whether a method body correctly answers the query, we compare how similar it is to the Stack Overflow answer - we do this systematically using a code-to-code similarity tool, called Aroma [5]. Aroma gives a similarity score between two code snippets; if this score is above a certain threshold (0.25 in our case), we count it as success. This similarity score, aims to mimic manually assessing the correctness of search results in an automatic and reproducible fashion, while leaving out human judgment in the process. More details on how we chose this threshold can be found in [1].

---

```
No.,StackOverflow ID,Question,NCS FRank,NCS_postrank FRank,UNIF_android FRank,UNIF_stackoverflow FRank
1,1109022,Close/hide the Android Soft Keyboard,NF,1,1,1
2,4616095,How to get the build/version number of your Android application?,17,1,31,19
3,3004515,Sending an Intent to browser to open specific URL,2,1,5,2
4,1560788,How to check internet access on Android? InetAddress never times out,1,4,5,1
5,3423754,Retrieving Android API version programmatically,5,1,22,10
6,1397361,How do I restart an Android Activity,NF,3,2,1
```

---

**Listing 4.** Score sheet example. "NF" stands for correct answer not found in the top 50 returned results.

**Table 1.** Number of questions answered in the top 1, 5, 10 and MRR for NCS, NCS<sub>postrank</sub>, UNIF<sub>android</sub> and UNIF<sub>stackoverflow</sub>.

Model	Answered@1	Answered@5	Answered@10	MRR
NCS	33	74	98	0.189
NCS <sub>postrank</sub>	85	151	180	0.4
UNIF <sub>android</sub>	25	74	110	0.178
UNIF <sub>stackoverflow</sub>	<b>104</b>	<b>164</b>	<b>188</b>	<b>0.465</b>

## 4 Evaluation

We provide the results for four models: NCS, NCS<sub>postrank</sub>, UNIF<sub>android</sub>, and UNIF<sub>stackoverflow</sub>.

Table 1 reports the number of questions answered within the top\_n returned code snippet, where n = 1, 5, and 10 (Answered@1, 5, 10 in Table 1), as well as the Mean Reciprocal Rank (MRR).

## References

- [1] Jose Cambrero, Hongyu Li, Seohyun Kim, Koushik Sen, and Satish Chandra. When deep learning met code search. *CoRR*, abs/1905.03813, 2019. URL: <https://arxiv.org/abs/1905.03813>, arXiv:1905.03813.
- [2] Xiaodong Gu, Hongyu Zhang, and Sunghun Kim. Deep code search. In *Proceedings of the 40th International Conference on Software Engineering*, pages 933–944. ACM, 2018.
- [3] GitHub Inc. Github rest api v3. URL: <https://developer.github.com/v3/search/>.
- [4] Stack Exchange Inc. datastack exchange data dump, 2018. URL: <https://archive.org/details/stackexchange>.
- [5] Sifei Luan, Di Yang, Celeste Barnaby, Koushik Sen, and Satish Chandra. Aroma: Code recommendation via structural code search. *CoRR*, abs/1812.01158, 2018. URL: <http://arxiv.org/abs/1812.01158>, arXiv:1812.01158.
- [6] Saksham Sachdev, Hongyu Li, Sifei Luan, Seohyun Kim, Koushik Sen, and Satish Chandra. Retrieval on source code: a neural code search. In *Proceedings of the 2nd ACM SIGPLAN International Workshop on Machine Learning and Programming Languages*, pages 31–41. ACM, 2018.