

CoNCRA: A Convolutional Neural Networks Code Retrieval Approach

Marcelo de Rezende Martins
rezende.martins@gmail.com
IPT - Institute for Technological Research
Sao Paulo, Sao Paulo, Brazil

Marco Aurélio Gerosa
marco.gerosa@nau.edu
Northern Arizona University (NAU)
Flagstaff, Arizona, United States

ABSTRACT

Software developers routinely search for code using general-purpose search engines. However, these search engines cannot find code semantically unless it has an accompanying description. We propose a technique for semantic code search: A Convolutional Neural Network approach to code retrieval (CoNCRA). Our technique aims to find the code snippet that most closely matches the developer's intent, expressed in natural language. We evaluated our approach's efficacy on a dataset composed of questions and code snippets collected from Stack Overflow. Our preliminary results showed that our technique, which prioritizes local interactions (words nearby), improved the state-of-the-art (SOTA) by 5% on average, retrieving the most relevant code snippets in the top 3 (three) positions by almost 80% of the time. Therefore, our technique is promising and can improve the efficacy of semantic code retrieval.

CCS CONCEPTS

- **Computing methodologies** → *Learning latent representations*;
- **Software and its engineering** → *Reusability*; *General programming languages*.

KEYWORDS

code search, neural networks, joint embedding

ACM Reference Format:

Marcelo de Rezende Martins and Marco Aurélio Gerosa. 2020. CoNCRA: A Convolutional Neural Networks Code Retrieval Approach. In *34th Brazilian Symposium on Software Engineering (SBES '20)*, October 21–23, 2020, Natal, Brazil. ACM, New York, NY, USA, 6 pages. <https://doi.org/10.1145/3422392.3422462>

1 INTRODUCTION

The advent of open-source software and question and answering websites contributed for improving the way developers produce code. Nowadays, code search permeates the development activities. Developers can spend 15% of their time searching online for how a piece of code works, how to fix a bug, and how to use an API [18]. According to Sadowski et al. [15], at Google, developers search for

code 12 times a day, clicking on 2 to 3 results in average per search session.

Most developers use general-purpose search engines (GPSE) to look for code (e.g., Google Search), which uses page rank and other indexes tactics that are not optimized for searching code. Then, general-purpose search engines do not adequately find code snippets unless they have accompanying descriptions. According to Rahman et al. [13], developers spend more time, visit more pages, and change queries more often when they are doing code-related searches. In particular, newcomers to a project can greatly benefit from semantic search since they face a variety of entrance barriers [16].

GitHub, a popular source code hosting platform, has attempted to build a semantic code search. They extracted millions of lines of code from its repositories and matched each code snippet to a docstring. The final results were not satisfactory as the tool could find a relevant code snippet only if the user provided a query that matched the docstring description [8]. According to Cambronero et al. [3], users' intents were better matched to questions collected from question-answering sites related to programming, e.g., Stack Overflow. Those sites allow users to ask a question and approve the best answer for it. Other users vote for the most helpful answer and mark the wrong or not helpful ones. Those collective actions curate and organize information.

Initial code search studies were based on deductive-logic rules and manually extracted features [1]. The recent success of artificial neural networks has shifted recent works to a machine learning-based approach. Cambronero et al. [3] coined a name, neural code search, i.e., code search based on neural networks.

Recent works applied neural networks to summarize and retrieve code snippets. Cambronero et al. [3] proposed a neural network with attention mechanism and Gu et al. [7] presented a recurrent neural network. Our novel approach is based on Convolutional Neural Networks (CNNs). For the best of our knowledge, CNNs have not yet been used to search for code, but have achieved good results in selecting answers [5, 17]. CNNs prioritize local interactions (words nearby) and its translation invariant, which are important traits for our task.

In our study, we answer the following research questions:

- What is the efficacy of the CONCRA technique?
- How does it compare to the baseline and SOTA methods?

2 BACKGROUND

According to Cambronero et al. [3], the main goal of code retrieval is to retrieve code snippets from a code corpus that most closely match a developer's intent, which is expressed in natural language.

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.

SBES '20, October 21–23, 2020, Natal, Brazil

© 2020 Association for Computing Machinery.

ACM ISBN 978-1-4503-8753-8/20/09...\$15.00

<https://doi.org/10.1145/3422392.3422462>

The first studies for code search were based on deductive-logic rules and manually extracted features. The Deductive-logic approach, e.g., boolean model, finds a code that precisely matches the keywords expressed in the query. According to Yan et al. [19], these approaches are good at finding API calls and error messages, but struggle to find reusable code and examples that do not have an exact match between the code and query.

Since neural networks showed good results at translation, question-answering, and classifications tasks in natural language processing, recent works adopted the neural networks for code search. These studies aim to discriminate relevant code snippets from non-relevant ones based on the user's intent. In pursuance of that, code retrieval is reduced to a ranking problem, in which neural networks aim to place code snippets that closely match the developer's intent in the top positions. The most common strategy is *joint embedding*. Joint embedding maps heterogeneous data into a common vector space, in which the distance between embedded input reflects the similarity between the underlying items [11] (see Figure 1).




Figure 1: Illustration of the *joint embedding* technique for code retrieval. Two neural networks map a question and a code snippet into a common vector space. The distance between the vectors reflects the relevance of a code snippet to a question. Adapted from [3].

To apply joint embedding, one needs to consider word, sentence, and joint embedding. Embedding refers to a continuous vector in a lower dimensional vector space. A function that maps an input to a continuous vector is called *encoder*. Given an input set X , an encoder function F can be defined as [3]:

$$F : X \rightarrow E \quad (1)$$

For code retrieval, X can be a set of questions or code snippets and E is a set of continuous vectors or embeddings, such that $E \subset \mathbb{R}^d$, where d is the dimension. The main goal is to learn two encoders F and G that map a question and a code snippet, respectively, into a common vector space, so that the distance between the vectors reflects the relevance of a code snippet to a question.

3 METHODOLOGY

In this work, we propose the use of convolutional neural networks to learn sentence embedding, i.e., convolutional neural networks will encode the question and code snippet into a continuous vector in a shared vector space. In the following, we explain how words are embedded and what objective function we use, as the objective

function tells how neural networks should approximate questions and code snippets.

3.1 Word embedding

The words and terms of a question and code snippet must be encoded into a numeric vector. The most common encoder for words is *word2vec*, which embeds a word into a continuous vector based on the distributional hypothesis. The distributional hypothesis says that two words are similar if they appear together frequently in different contexts [6]. Context can be a sentence, paragraph, or document in NLP tasks. In our case, the context is questions and code snippets.

Word2vec has two strategies: continuous-bag-of-words (CBoW) and skip-gram. The main difference is that CBoW predicts a target word given context words, and skip-gram predicts the context words given a single word. According to Mikolov et al. [12], CBoW showed good results at syntactic tasks, e.g., finding a superlative of a word or identifying an adverb. At the same time, skip-gram presented good performance for semantic tasks, e.g., finding a state's capital or grouping feminine and masculine words.

In our work, we opted for skip-gram as a semantic trait is preferable. Semantic trait can help the neural networks to discriminate conditional clauses (e.g., *if*, *elif*, *elseif*) and loop iteration (e.g., *for*, *while*), for example. Figure 2 shows an application of word2vec in a Python related corpus. We can see the similarities between *file*, *write*, and *open*, and *set*, *list*, and *dict* based on the distance between them.




Figure 2: 2D picture of continuous vectors of the 66 most frequent words from a Python corpus V . The illustration was generated by t-SNE, which allows 2D visualization from high-dimensional data. We applied word2vec with skip-gram and a parameter window 5.

3.2 Sentence embedding

We can combine the word embeddings to obtain a sentence embedding. We combine word embeddings by using convolutional neural networks. Convolutional neural networks showed good results at answer selection tasks in NLP—given a question and a set of answers, the model ranks the best answers. Convolutional networks prioritize local interactions (e.g., words nearby) and cannot capture long-range dependencies (e.g., distant words in a sentence). However, this issue is mitigated for code retrieval since most questions and code snippets are short in length.

Given an sentence $\mathbf{x} = \{\mathbf{x}(0), \mathbf{x}(1), \dots, \mathbf{x}(n-1)\}$, such that $\mathbf{x}(i) \in \mathbb{R}^d$ is a continuous vector that represents the i^{th} word of the sentence. Convolutional neural networks combine the elements of vector \mathbf{x} by applying 2 basic operations: Convolution operation and Maxpool.

A convolution operation uses a filter $F = [F(0), \dots, F(m-1)]$, such that $F \in \mathbb{R}^{m \times d}$. The operation applies the filter in m words (window size) to produce a new vector. Suppose $\mathbf{x}(i, i+j)$ refers to a concatenation of the vectors $\mathbf{x}(i), \mathbf{x}(i+1), \dots, \mathbf{x}(i+j)$. If we apply F to $\mathbf{x}(i, i+m-1)$, then we can calculate a new vector $\mathbf{c}(i)$ by:

$$\mathbf{c}(i) = \tanh \left[\left(\sum_{j=0}^{m-1} \mathbf{x}(i+j)^T F(j) \right) + b \right] \quad (2)$$

In the equation 2, F and b are learnable weights and bias, respectively. The convolution operation slides the filter F across the height of input \mathbf{x} and computes the dot product between the entries of the filter and the input [10]. The operation returns a feature map \mathbf{c} .

$$\mathbf{c} = \{\mathbf{c}(0), \mathbf{c}(1), \dots, \mathbf{c}(n-m)\} \quad (3)$$

The feature map (or activation map) contains the latent and most important features of a sentence. A convolutional neural network may contain thousands of filters, each one extracting specific *m-gram* features (e.g., a filter of m size 2 extracts bigram features). The quantity of feature maps is $|F|$, i.e., the number of filters. After the convolution operation, a pooling layer operates independently on every feature map resizing it spatially, using a max operation [10]. The max operation is applied along the *axis = 0* to produce the final vector \mathbf{o} :

$$\mathbf{o} = \max \left([c_1, c_2, \dots, c_{|F|}], \text{axis} = 0 \right) \quad (4)$$

The max-pooling helps the convolutional neural networks to be translation invariant. Regardless of the word's position shift, the max-pooling selects the most relevant features and inserts it into the final vector. Figure 3 shows the results of each operation. We use the vector \mathbf{o} as our sentence embedding, then \mathbf{o} represents a question or a code snippet, in our case.

3.3 Joint embedding

We treat code retrieval as a ranking problem, in which a model should rank relevant code in the top positions based on the developer's intent. To do so, we adopted an objective function that prioritizes the relative preference of the code snippets. In our case, the objective function helps the neural networks to separate the correct answers from incorrect ones during the training phase.

Given a question and code snippet set, \mathbb{Q} and \mathbb{C} , our training input comprises a triple $\langle \mathbf{q}, \mathbf{c}^+, \mathbf{c}^- \rangle$, where $\mathbf{c}^+ \in \mathbb{C}$ indicates a correct code snippet for a question $\mathbf{q} \in \mathbb{Q}$ and $\mathbf{c}^- \in \mathbb{C}$ an incorrect one sampled from the training data. We used the hinge loss as our objective function. Formally, for a triple $\langle \mathbf{q}, \mathbf{c}^+, \mathbf{c}^- \rangle$, the definition of hinge loss is:

$$J = \max(0, m - h_\theta(\mathbf{q}, \mathbf{c}^+) + h_\theta(\mathbf{q}, \mathbf{c}^-)) \quad (5)$$

m is a margin and h_θ is a similarity function (e.g., *cosine*). During the training phase, the goal is to minimize the cost function J .




Figure 3: Schematic drawing of our Convolutional Neural Networks (CNN) operations. Our example shows 4 filters $F \in \mathbb{R}^{m \times d \times f}$ with window size $m = 2$. We slide each filter across the height of the input $\mathbf{x} \in \mathbb{R}^{n \times d}$, where $n = 6$ and $d = 5$, and computes the dot product between the entries of the filter and the input [10]. It returns a feature map \mathbf{c} . In the end, a max pooling layer resizes every feature map spatially, obtaining the final vector \mathbf{o} . We use the vector \mathbf{o} as our question and code snippet embedding. Adapted from [21].

To obtain that, the model aims to satisfy the following condition: $h_\theta(\mathbf{q}, \mathbf{c}^+) - h_\theta(\mathbf{q}, \mathbf{c}^-) \geq m$. Then, the hinge loss function induces our model to score \mathbf{c}^+ higher than \mathbf{c}^- for a given margin m . For the similarity function h_θ , we used *cosine*.

4 EXPERIMENTS

To verify the efficacy of our approach and how it compares to other approaches, we trained and evaluated all models in the same environment, following Yao et al. [20] and Iyer et al. [9] experiment protocol. We compared the following approaches:

- **CoNCRA**: our proposed approach, described in Section 3. We tried two variations of our approach, the Convolutional Neural Networks (CNN) and Shared CNN. Their difference is that Shared CNN shares the weights by learning the question and code snippet embeddings, while CNN learns different weights for each one.
- **Embedding**: it is our baseline model. It is a simple architecture that applies a max-pooling layer to the word embeddings.
- **Unif**: it is the solution proposed by Cambronero et al. [3]. They used two distinct layers for the question and code snippet embedding. They applied an average pooling layer to word embedding in order to learn the question embeddings. For the code snippet, they used an attention mechanism, which applies a weighted average to each word embedding, "giving attention" to the most relevant word of the code.

We evaluated the models on the StaQC dataset, a systematically mined question-code dataset from Stack Overflow [20]. The main difference of StaQC to other datasets is that it is composed of "how-to-do-it" questions, as most of the answers to those types of questions are straightforward. StaQC contains SQL and Python questions, but, in our experiments, we used Python questions only.

Code snippet	Question	
	Python	SQL
N_1 : Single code snippet in the answer description	85.294	75.637
N_2 : Automatically annotated code snippets	60.083	41.826
N_3 : Manually annotated code snippets	2.169	2.056
Total	147.546	119.519

Table 1: Summary of StaQC dataset [20]. Questions from N_2 sample ("Automatically annotated code snippets") may contain more than one code snippet per answer description. Some code snippets may not be a solution to the question. So, the authors proposed a framework to annotate the code snippets automatically and it could achieve an F1 score of 0,916 and an accuracy of 0,911.

Sample	Quantity of pairs $\langle q_i, c_i^+ \rangle$
N_2 = Training	60.083
$N_3 \supset$ DEV	1.085
$N_3 \supset$ EVAL	1.084
Total	62.252

Table 2: Summary of our training and evaluation samples. The samples are composed of a pair $\langle q_i, c_i^+ \rangle$, where q_i is a question and c_i^+ is a code snippet annotated as solution. We split the manually annotated dataset into two parts: DEV and EVAL, according to [9] procedure.

All models were trained on sample N_2 (see Table 1 and Table 2), because 27% of the questions contain more than one annotated answer, leading to more variance in our training dataset. The training and evaluation follow Iyer et al. [9]'s procedure, in which the model is evaluated on a manually annotated dataset each epoch according to a Mean Reciprocal rank (MRR). The MRR tells if a model ranked the annotated answer in higher positions, i.e., higher values for MRR indicate that the accepted answers were ranked in the top positions.

For the training phase, we split 70% of N_2 sample for training and 30% for validation. We run the models for 500 epochs and stop early if the training loss (J) is less than 0.0001 or the validation loss does not improve after 25 consecutive epochs. For the final evaluation, we choose the best model of the training phase according to MRR and run it for 20 iterations in the N_3 sample (see Table 2). The final result is the average of the MRR for each pair $\langle q_i, c_i^+ \rangle$ of N_3 and other 49 distractors c_j , which were selected randomly from the training sample, such that $c_i^+ \neq c_j$.

We provide the source code for our preliminary experiments in the following repository: <https://github.com/mrezende/concra>. The repository contains our proposed model, the baseline ones, training, and evaluation source code. We also provide the original and pre-processed dataset. The source code is written in Python, version 3.6.9, and we used the libraries Keras (version 2.2.4-tf) and TensorFlow (1.15.2). The tests were all conducted in the Colab

platform¹. In Figure 4, we provide examples of the output of our tool.

Q1: finding the max of a column in an array	Q2: Check if all elements of one array is in another array
A 1.1 <pre>import numpy as np aa= np.random.random([4,3]) print aa print np.max(aa,axis=0) print np.max(aa,axis=1)</pre>	A 2.1 <pre>np.argmaxwhere([B[None,0,1]] == A[:,1,2]).all(-1))</pre>
A 1.2 <pre>chapters = ['one', 'two', 'three',] for x in chapters: print x if x != chapters[-1]: print 'next'</pre>	A 2.2 <pre>def is_slice_in_list(s,l): len_s = len(s) return any(s == l[i:len_s+i] for i in xrange(len(l) - len_s+1))</pre>
A 1.3 <pre>import numpy as np a = np.ones((1,2,2,2)) print(a.shape) # (1, 2, 2, 2)</pre>	A 2.3 <pre>sum(all(a == b for a, b in izip_longest(target, my_list[(i * n):((i + 1) * n)])) for i in range(len(my_list) // n))</pre>

Figure 4: Example of questions and that our model (CoN-CRA) gave. Our model answered the first question (Q1) correctly, selecting an answer (A 1.1) based on the Numpy library, which adds support for multi-dimensional and large arrays in Python. The second answer (A 1.2) for the first question returns the last element of an array, which is incorrect, but seems interesting. The third one (A 1.3) shows the array's dimension using the Numpy library, so our model found a correlation between Numpy and array operations. In the other example (Q2), the first answer is incorrect, as it checks the presence of elements in a matrix, not an array. Again, our model showed a Numpy answer. The second one (A 2.2) is the correct answer, and the third one (A 2.3) checks how many times a sequence repeats in a data frame.

4.1 Results

Table 3 compiles the final results after 20 runs on the N_3 sample. Shared CNN with 4000 filters got the best results (row D3 and F3). Our proposed architecture achieved an MRR score 5% higher on average than the best result obtained by Unif (row B1), which can be considered a state-of-the-art approach. Our MRR result is 11% higher than the baseline model (row A1).

As presented in Table 3, an increase in the number of filters resulted in a better performance for CNN, as the model capacity and the number of extracted features grow (row D3 and F3). We tried different margin loss (see Section 3.3). CNN got the best results with 0.05, while Unif and Embedding got better results with 0.2 and 0.1, respectively.

We verified that shared weights models (row D and F) got better results than independently weights ones (row C and E). One reason is that the optimizer of independent weights architecture has to learn the double of parameters, increasing the learning difficulty [5]. We also tried a batch normalization technique to avoid overfitting

¹<https://colab.research.google.com/>

		Results	
Models		MRR	TOP-1
A1	Embedding ($m = 0.1$)	0.637	0.493 ± 0.009
B1	Unif ($m = 0.2$)	0.675 ± 0.006	0.539 ± 0.009
C1	CNN / $F = 1000$	0.669 ± 0.006	0.527 ± 0.012
C2	CNN / $F = 2000$	0.673 ± 0.007	0.531 ± 0.012
C3	CNN / $F = 4000$	0.687 ± 0.006	0.553 ± 0.011
D1	Shared CNN / $F = 1000$	0.678 ± 0.007	0.548 ± 0.012
D2	Shared CNN / $F = 2000$	0.694 ± 0.008	0.565 ± 0.012
D3	Shared CNN / $F = 4000$	0.700 ± 0.004	0.569 ± 0.009
E1	CNN with BN / $F = 1000$	0.682 ± 0.007	0.543 ± 0.012
E2	CNN with BN / $F = 2000$	0.689 ± 0.006	0.553 ± 0.011
E3	CNN with BN / $F = 4000$	0.688 ± 0.006	0.553 ± 0.011
F1	Shared CNN with BN / $F = 1000$	0.690 ± 0.008	0.553 ± 0.015
F2	Shared CNN with BN / $F = 2000$	0.700 ± 0.007	0.573 ± 0.012
F3	Shared CNN with BN / $F = 4000$	0.701 ± 0.008	0.577 ± 0.015

Table 3: The experimental results of EVAL sample for Embedding, Unif, and our two CNN variations: CNN and Shared CNN. m refers to the margin loss of the hinge loss function (lines A1 and B1). F indicates the number of filters. BN is an acronym of Batch Normalization. Our CNN architecture used a margin loss $m = 0.05$ and a window size of 2.

and help our models learn quickly and a more stable way, but only CNN and Shared CNN got better results. For Unif and Embedding, it did not improve the performance at all.

Figure 5 illustrates the MRR (Mean Reciprocal Rank), showing the first position of the annotated code snippet during the final evaluation.




Figure 5: Histogram of the first positions observed for the annotated code snippet during the final evaluation. The labels *shared-cnn-with-bn*, *unif*, and *embedding* refer to lines F3, A1, and B1, respectively, in Table 3.

Both CNN and Unif ranked the code snippets among the first three positions in 75% of cases. We got a TOP-1 accuracy of 60%, i.e., we ranked the relevant code snippet in the first place in 60% of cases. Unif and Embedding got a TOP-1 accuracy of almost 50%. However, it is worth noticing that MRR considers only the annotated code snippet position. If the model shows up another code snippet, which

correctly answers the question, our metric does not consider it and the model is penalized.

4.2 Threats to validity

We trained the models on the StaQC dataset, a systematically mined Stack Overflow dataset. The authors used neural networks to annotate the code snippets trained on a manual dataset. In our case, we trained the models on the automatically annotated corpus and evaluated on the manual ones. To mitigate bias, we adopted Iyer et al. [9]’s procedure for training and evaluation.

5 RELATED WORK

We summarize the difference between our work and related work in Table 4. Most works differ on how they combine the word embeddings to obtain a sentence embedding. Recent works (row E and F) used Skip-gram (see Section 3.1) and adopted simpler architectures for sentence embedding than previous work (row C and D).

Previous work (rows D to F) that used GitHub data extracted the methods from the source code and matched them to docstring descriptions. Works that used Stack Overflow data (rows A to C) paired question titles to the accepted answers’ code snippet. For the search corpus, some works (rows D to F) adopted a GitHub corpus with millions of pieces of code, while others (rows A to C) retrieved code snippets from a small sample of 50 randomly selected code snippets. These studies (rows A to G) are commonly evaluated using questions collected from Stack Overflow.

Although our architecture is more complex, since it requires more parameters and time to train, than Cambronero et al. [3]’s (row F) architecture (SOTA), we can train our model offline. We matched question title to code snippets collected from Stack Overflow following Allamanis et al. [2] and Iyer et al. [9] (rows A and B) work.

	Work	Feature	Word Embedding	Sentence Embedding
A	Allamanis et al. [2]	Token / Parse Tree	Probabilistic Model	Average / Context matrix
B	Iyer et al. [9]	Token	One-hot encoding	LSTM with attention mechanism
C	Chen and Zhou [4]	Token	Bi-VAE	Bi-VAE, Average and MLP
D	Gu et al. [7]	Token / Method name, API invocation and Token	bi-LSTM	Max-pooling / Max-pooling and MLP
E	Sachdev et al. [14]	Token	Skip-gram	Average / TF-IDF
F	Cambronero et al. [3]	Token	Skip-gram	Average / Attention
G	Our work	Token	Skip-gram	CNN

Table 4: Summary of related joint embedding work. The column *Feature* refers to question and code representation. Different approaches for a question and code snippet are separated by slash ("/"), showing the question method first followed by the code snippet technique. Adapted from [19].

6 CONCLUSION

Our model, CoNCRA, achieved an MRR score 5% higher on average than Unif, a state-of-the-art technique. We could rank the most relevant code snippet among the first 3 (three) positions in 78% of cases. Our technique achieved a TOP-1 accuracy of 60%, while the other techniques achieved 50%.

The results seem promising, and we plan further investigation to check if our model is invariant to other datasets. We will also investigate transfer learning, e.g., checking if a model trained on a Stack Overflow dataset (curated corpus) can find relevant code snippets in a GitHub corpus (non-curated one). Our approach is based on an NLP technique proposed by Feng et al. [5], and future work may use transformers and autoencoders, as those techniques showed good results in many NLP tasks.

7 ACKNOWLEDGEMENT

This work was partially supported by the National Science Foundation (grant 1815503).

REFERENCES

- [1] Miltiadis Allamanis, Earl T. Barr, Premkumar Devanbu, and Charles Sutton. 2018. A Survey of Machine Learning for Big Code and Naturalness. *ACM Comput. Surv.*, Article 81 (July 2018), 37 pages.
- [2] Miltiadis Allamanis, Daniel Tarlow, Andrew D. Gordon, and Yi Wei. 2015. Bimodal Modelling of Source Code and Natural Language. In *Proceedings of the 32Nd International Conference on International Conference on Machine Learning - Volume 37*. JMLR.org, 2123–2132.
- [3] Jose Cambronero, Hongyu Li, Seohyun Kim, Koushik Sen, and Satish Chandra. 2019. When Deep Learning Met Code Search. In *Proceedings of the 2019 27th ACM Joint Meeting on European Software Engineering Conference and Symposium on the Foundations of Software Engineering*, 964–974.
- [4] Qingying Chen and Minghui Zhou. 2018. A Neural Framework for Retrieval and Summarization of Source Code. In *Proceedings of the 33rd ACM/IEEE International Conference on Automated Software Engineering*, 826–831.
- [5] M. Feng, B. Xiang, M. R. Glass, L. Wang, and B. Zhou. 2015. Applying deep learning to answer selection: A study and an open task. In *2015 IEEE Workshop on Automatic Speech Recognition and Understanding (ASRU)*, 813–820.
- [6] Ian Goodfellow, Yoshua Bengio, and Aaron Courville. 2016. *Deep Learning*. MIT Press. <http://www.deeplearningbook.org>.
- [7] Xiaodong Gu, Hongyu Zhang, and Sunghun Kim. 2018. Deep Code Search. In *Proceedings of the 40th International Conference on Software Engineering*, 933–944.
- [8] Hamel Husain and Ho-Hsiang Wu. 2018. *How To Create Natural Language Semantic Search For Arbitrary Objects With Deep Learning*. <https://towardsdatascience.com/semantic-code-search-3cd6d244a39c>
- [9] Srinivasan Iyer, Ioannis Konstas, Alvin Cheung, and Luke Zettlemoyer. 2016. Summarizing Source Code using a Neural Attention Model. In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, 2073–2083.
- [10] Fei-Fei Li, Andrej Karpathy, and Justin Johnson. 2016. *CS231n: Convolutional Neural Networks for Visual Recognition 2016*. <http://cs231n.stanford.edu/>
- [11] Yangyan Li, Hao Su, Charles Ruizhongtai Qi, Noa Fish, Daniel Cohen-Or, and Leonidas J. Guibas. 2015. Joint Embeddings of Shapes and Images via CNN Image Purification. *ACM Trans. Graph.* 34, 6, Article 234 (Oct. 2015), 12 pages.
- [12] Tomas Mikolov, Kai Chen, Greg S. Corrado, and Jeffrey Dean. 2013. Efficient Estimation of Word Representations in Vector Space. <http://arxiv.org/abs/1301.3781>
- [13] Md Masudur Rahman, Jed Barson, Sydney Paul, Joshua Kayani, Federico Andrés Lois, Sebastián Fernandez Quezada, Christopher Parnin, Kathryn T. Stolee, and Baishakhi Ray. 2018. Evaluating How Developers Use General-Purpose Web-Search for Code Retrieval. In *Proceedings of the 15th International Conference on Mining Software Repositories*, 465–475.
- [14] Saksham Sachdev, Hongyu Li, Sifei Luan, Seohyun Kim, Koushik Sen, and Satish Chandra. 2018. Retrieval on Source Code: A Neural Code Search. In *Proceedings of the 2Nd ACM SIGPLAN International Workshop on Machine Learning and Programming Languages*, 31–41.
- [15] Caitlin Sadowski, Kathryn T. Stolee, and Sebastian Elbaum. 2015. How Developers Search for Code: A Case Study. In *Proceedings of the 2015 10th Joint Meeting on Foundations of Software Engineering*, 191–201.
- [16] Igor Steinmacher, Tayana Conte, Marco Aurélio Gerosa, and David Redmiles. 2015. Social barriers faced by newcomers placing their first contribution in open source software projects. In *Proceedings of the 18th ACM conference on Computer supported cooperative work & social computing*, 1379–1392.
- [17] Jiahui Wen, Hongkui Tu, Xiaohui Cheng, Renquan Xie, and Wei Yin. 2019. Joint modeling of users, questions and answers for answer selection in CQA. *Expert Systems with Applications* 118 (2019), 563 – 572.
- [18] Xin Xia, Lingfeng Bao, David Lo, Pavneet Singh Kochhar, Ahmed E. Hassan, and Zhenchang Xing. 2017. What Do Developers Search for on the Web? *Empirical Softw. Engg.* 22, 6 (Dec. 2017), 3149–3185.
- [19] S. Yan, H. Yu, Y. Chen, B. Shen, and L. Jiang. 2020. Are the Code Snippets What We Are Searching for? A Benchmark and an Empirical Study on Code Search with Natural-Language Queries. In *2020 IEEE 27th International Conference on Software Analysis, Evolution and Reengineering (SANER)*, 344–354.
- [20] Ziyu Yao, Daniel S. Weld, Wei-Peng Chen, and Huan Sun. 2018. StaQC: A Systematically Mined Question-Code Dataset from Stack Overflow. In *Proceedings of the 2018 World Wide Web Conference*, 1693–1703.
- [21] Ye Zhang and Byron C. Wallace. 2015. A Sensitivity Analysis of (and Practitioners' Guide to) Convolutional Neural Networks for Sentence Classification. *CoRR abs/1510.03820* (2015). <http://arxiv.org/abs/1510.03820>