

Generalizability of Code Clone Detection on CodeBERT

Tim Sonnekalb
tim.sonnekalb@dlr.de

Institute of Data Science
German Aerospace Center (DLR)
Jena, Germany

Clemens-Alexander Brust
clemens-alexander.brust@dlr.de
Institute of Data Science
German Aerospace Center (DLR)
Jena, Germany

Bernd Gruner
bernd.gruner@dlr.de

Institute of Data Science
German Aerospace Center (DLR)
Jena, Germany

Patrick Mäder
patrick.maeder@tu-ilmenau.de
Data-intensive Systems and Visualization Group
Technical University Ilmenau
Ilmenau, Germany

ABSTRACT

Transformer networks such as CodeBERT already achieve outstanding results for code clone detection in benchmark datasets, so one could assume that this task has already been solved. However, code clone detection is not a trivial task. Semantic code clones, in particular, are challenging to detect.

We show that the generalizability of CodeBERT decreases by evaluating two different subsets of Java code clones from BigCloneBench. We observe a significant drop in F1 score when we evaluate different code snippets and functionality IDs than those used for model building.

CCS CONCEPTS

• Software and its engineering → Software maintenance tools.

KEYWORDS

transformer, code clone detection, bigclonebench codebert, codexglue

ACM Reference Format:

Tim Sonnekalb, Bernd Gruner, Clemens-Alexander Brust, and Patrick Mäder. 2022. Generalizability of Code Clone Detection on CodeBERT. In *37th IEEE/ACM International Conference on Automated Software Engineering (ASE '22)*, October 10–14, 2022, Rochester, MI, USA. ACM, New York, NY, USA, 3 pages. <https://doi.org/10.1145/3551349.3561165>

1 INTRODUCTION

Large language models, like transformers, are context-sensitive and should therefore be well suited for clone detection. Microsoft’s benchmark CodeXGlue [4] contains several machine learning for software engineering tasks solved by their CodeBERT model [1]. One can observe the generalization capability of CodeBERT across tasks. One task is code clone detection which CodeXGlue models as binary classification. BigCloneBench (in Java) and POJ-104 (in C) are used as benchmark datasets.

Permission to make digital or hard copies of part or all 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 third-party components of this work must be honored. For all other uses, contact the owner/author(s).
ASE '22, October 10–14, 2022, Rochester, MI, USA
© 2022 Copyright held by the owner/author(s).
ACM ISBN 978-1-4503-9475-8/22/10.
<https://doi.org/10.1145/3551349.3561165>

BigCloneBench is a well-known code clone detection dataset that contains Java code snippets. Several methods [2, 3, 6, 10] evaluate this benchmark dataset, making it a standard in this field. There are mainly 3 versions of BigCloneBench: BCB v1 [7], BCB v2 [9] and CDLH (Clone Detection with Learning to Hash) [10]. The latter is a pre-processed version of BCB v1 and is used in CodeXGlue [4].

Code clone detection, in general, can be divided into four types of code clones. Type 1-3 are syntactic code clones with minor changes in their syntax. Type 4 clones or semantic clones are difficult to detect, as only a manual validation can safely check the different syntax but the same semantics. CodeBERT achieves an F1 score of 96.5% for this task [4] which could be caused by group leakage. We verify this score by evaluating two different subsets of BigCloneBench not used for model building.

2 DATA

Svajlenko et al. [7] create the BigCloneBench dataset (BCB v1) based on a sizable inter-project repository IJaDataset 2.0 of Java projects¹. They develop a semi-automated processing method to mine code clones based on ten given target similarity classes. In a follow-up work [9], the authors mine a larger version of BigCloneBench (BCB v2) with 43 code functionalities and provide the clone pairs in a database.

The CDLH work [11] presents another filtered version of BCB v1. The authors do not provide the pre-processed dataset or scripts. CodeXGlue uses the CDLH dataset and has published it with slightly different properties: instead of 9,134 given code snippets in CDLH, CodeXGlue contains 9,126. Though this dataset was filtered out of BCB v1, we reverse engineer 12 functionalities instead of 10. All properties are listed in Table 1.

After a detailed analysis of the dataset versions, we notice the following problems: The datasets are highly imbalanced in their two classes, *true clones* and *false clones*, in their clone types, and in their functionalities.

The clone class distribution is due to the automatic creation process of the dataset. The true clones are assigned following an automated process, while the false clones are labeled manually by judges. For this reason, there are far fewer false clones than true ones (comp. Tab. 1).

¹The link is not available anymore, but Svajlenko et al. provide a mirror: <https://github.com/clonebench/BigCloneBench>

Table 1: Dataset properties in BigCloneBench versions.

Property	BCB v1	BCB v2	CodeXGlue (CDLH)
Code snippets	59,650	73,319	9,126 (9,134)
Functionalities	10	43	12 (10)
Total clones	6,508,632	8,863,185	1,731,857
True clones	6,246,167	8,584,153	1,170,336
False clones	262,465	279,032	561,521

The clone types are determined by a similarity measure and are assigned to the Types *Type-1*, *Type-2*, *Very Strongly Type-3*, *Strongly Type-3*, *Moderately Type-3*, and *Weakly Type-3/Type 4* by hard-defined thresholds. Weakly Type-3 (syntactic) and Type-4 (semantic) clones cannot be separated with this approach. This class has a share of 98.2% of the whole dataset (comp. with [7]).

The authors [7] create a search heuristic for the functionality classes. A blueprint for every functionality class is defined and compared with the code examples. All 43 functionalities in BCB v2 are described here: [3]. The functionality classes with the most examples have very general names like *Copy* (~4.8M), *Zip Files* (~966k), or *Secure Hash* (~900k). We note that the full functionality of 25k software projects cannot be mapped to 43 classes in sufficient precision. Moreover, a code snippet can be assigned to multiple functionality classes if it matches multiple blueprints. BCB assigns only one functionality per example.

Another problem is that recall is only conditionally suitable as a measure for evaluating clone detection methods trained on previously labeled data. The original authors of BigCloneBench come to the same conclusion [8]. Recall is the proportion of true positives out of all positives and cannot be computed unless all combinations of clone pairs are given. For a more significant number of examples, this cannot be checked for resource reasons, and therefore it cannot be ruled out that there are code clones. An additional qualitative assessment on a small test subset can be helpful to obtain reliable measures.

2.1 Evaluation Datasets

CodeBERT uses only a small subset of ~12% of the entire BigCloneBench for model building. To test whether CodeBERT’s generalization ability on the remaining code snippets is as high as the authors claim, we create two evaluation sets by filtering BCBv2:

- 1. BCBv2 \CDLH separated by clone IDs** We extracted a list with all code snippet IDs from training and validation sets of CDLH. From BCB v2, we removed all clone pairs with a matching id from this code ID list.
- 2. BCBv2 \CDLH separated func IDs** We created a list of clone pairs from the CDLH training and validation sets and reconstructed their functionality IDs. With this functionality ID list, we selected all clone pairs from BCB v2 not included in the CDLH functionality IDs.

The functionality IDs were not given in CodeXGlue, so we reverse engineer the functionality types of each example in the CDLH dataset. We succeed in recovering ~60% of the functionalities from the clone pairs. The clone IDs are independent of the functionalities. They have remained the same across versions of the dataset.

Table 2: Evaluation metrics of CodeBERT in %.

Evaluation Dataset	Precision	Recall	F1 score
Lu et al. [4] (paper)	-	-	95.6
Gou et al. [2] (paper)	94.7	93.4	94.1
Lu et al. [4] (repeated)	93.03	95.82	94.40
1. BCBv2 \CDLH separated by clone IDs	98.18	52.41	68.34
2. BCBv2 \CDLH separated by func IDs	97.77	32.67	48.97

3 EVALUATION

We use the same hyperparameters and base model *codebert-base* as the benchmark to ensure comparability. We reproduce the fine-tuning of the model and run our evaluation on 4 Nvidia A100 GPUs.

The baseline of CodeXGlue was fine-tuned on only 10% of training and 10% of validation data [5]. The train-valid-test fraction of CDLH is 52/24/24, which results in ~90k training and ~41k validation examples for fine-tuning. Table 2 shows the evaluation results of our repeated fine-tuning of *codebert-base* compared to the values given in the original papers and to our two created datasets. The CodeXGlue paper does not provide precision and recall values. Our repetition yielded a slightly worse F1 score by 1.2%.

Evaluating our self-created subsets of BCBv2, we find a significant drop in F1 score of ~27% for dataset 1 and ~47% for dataset 2. Dataset 1 includes different clone IDs but partly code snippets with the same functionalities. Dataset 2 includes completely different functionalities, which could mean that the functionalities used for model building are too few to generalize across functionalities and other clone pairs.

Furthermore, transformer networks are limited by their input sequence length. CodeBERT’s input limit is 800 tokens, which results in a sequence length of 400 per code snippet. More extended code snippets are truncated, which affects ~32% of the CodeXGlue code snippets.

The bimodal CodeBERT model is pre-trained on the CodeSearchNet dataset, which includes pairs of documentation and code snippets from 6 programming languages, the percentage of Java in it is ~24%. Even if the fine-tuning was done for only one language, some tokens might have different cross-language semantics. Moreover, there is a gap in pre-training on combinations of natural language and programming languages and fine-tuning only on programming language pairs.

4 CONCLUSION

We investigated the generalizability of the CodeBERT model by evaluating on unseen data from a larger portion of the same dataset. Although the results of CodeBERT for clone detection look good at first glance, there is still a significant drop in F1 score when evaluated on other code snippets and other functionalities.

We recommend the authors of CodeXGlue to revise their benchmark and use a more extensive training dataset by evaluating on the whole BigCloneBench dataset. Furthermore, BigCloneBench is a highly unbalanced dataset, so this must always be considered when pre-processing the data.

REFERENCES

- [1] Zhangyin Feng, Daya Guo, Duyu Tang, Nan Duan, Xiaocheng Feng, Ming Gong, Linjun Shou, Bing Qin, Ting Liu, Daxin Jiang, and Ming Zhou. 2020. CodeBERT: A Pre-Trained Model for Programming and Natural Languages. *arXiv:2002.08155 [cs]* (Sept. 2020). <http://arxiv.org/abs/2002.08155> arXiv: 2002.08155.
- [2] Daya Guo, Shuo Ren, Shuai Lu, Zhangyin Feng, Duyu Tang, Shujie Liu, Long Zhou, Nan Duan, Jian Yin, Daxin Jiang, and Ming Zhou. 2020. GraphCodeBERT: Pre-training Code Representations with Data Flow. *arXiv:2009.08366 [cs]* (Sept. 2020). <http://arxiv.org/abs/2009.08366> arXiv: 2009.08366.
- [3] Muhammad Hammad, Onder Babur, Hamid Abdul Basit, and Mark van den Brand. 2020. DeepClone: Modeling Clones to Generate Code Predictions. In *Reuse in Emerging Software Engineering Practices: 19th International Conference on Software and Systems Reuse, ICSR 2020, Hammamet, Tunisia, December 2–4, 2020, Proceedings* (Hammamet, Tunisia). Springer-Verlag, Berlin, Heidelberg, 135–151. https://doi.org/10.1007/978-3-030-64694-3_9
- [4] Shuai Lu, Daya Guo, Shuo Ren, Junjie Huang, Alexey Svyatkovskiy, Ambrosio Blanco, Colin Clement, Dawn Drain, Daxin Jiang, Duyu Tang, Ge Li, Lidong Zhou, Linjun Shou, Long Zhou, Michele Tufano, Ming Gong, Ming Zhou, Nan Duan, Neel Sundaresan, Shao Kun Deng, Shengyu Fu, and Shujie Liu. 2021. CodeXGLUE: A Machine Learning Benchmark Dataset for Code Understanding and Generation. <https://doi.org/10.48550/ARXIV.2102.04664>
- [5] Microsoft Research. 2021. CodeXGLUE/Code-Code/Clone-detection-BigCloneBench at main · microsoft/CodeXGLUE · GitHub. Online. Last visited 25.07.2022.. <https://github.com/microsoft/CodeXGLUE/blob/919cf4278dda016f8fcc49ba3e4ef7971137e4fc/README.md>
- [6] Vaibhav Saini, Farima Farmahinifarahani, Yadong Lu, Pierre Baldi, and Cristina V. Lopes. 2018. Oreo: Detection of Clones in the Twilight Zone. In *Proceedings of the 2018 26th ACM Joint Meeting on European Software Engineering Conference and Symposium on the Foundations of Software Engineering* (Lake Buena Vista, FL, USA) (*ESEC/FSE 2018*). Association for Computing Machinery, New York, NY, USA, 354–365. <https://doi.org/10.1145/3236024.3236026>
- [7] Jeffrey Svajlenko, Judith F. Islam, Iman Keivanloo, Chanchal K. Roy, and Mohammad Mamun Mia. 2014. Towards a Big Data Curated Benchmark of Inter-project Code Clones. In *2014 IEEE International Conference on Software Maintenance and Evolution*. 476–480. <https://doi.org/10.1109/ICSME.2014.77>
- [8] Jeffrey Svajlenko and Chanchal K. Roy. 2015. Evaluating clone detection tools with BigCloneBench. In *2015 IEEE International Conference on Software Maintenance and Evolution (ICSME)*. 131–140. <https://doi.org/10.1109/ICSM.2015.7332459>
- [9] Jeffrey Svajlenko and Chanchal K. Roy. 2016. BigCloneEval: A Clone Detection Tool Evaluation Framework with BigCloneBench. In *2016 IEEE International Conference on Software Maintenance and Evolution (ICSME)*. 596–600. <https://doi.org/10.1109/ICSME.2016.62>
- [10] Wenhan Wang, Ge Li, Bo Ma, Xin Xia, and Zhi Jin. 2020. Detecting code clones with graph neural network and flow-augmented abstract syntax tree. In *2020 IEEE 27th International Conference on Software Analysis, Evolution and Reengineering (SANER)*. IEEE, 261–271.
- [11] Huihui Wei and Ming Li. 2017. Supervised Deep Features for Software Functional Clone Detection by Exploiting Lexical and Syntactical Information in Source Code.. In *IJCAI*. 3034–3040.