

# Automatic Classification of Error Types in Solutions to Programming Assignments at Online Learning Platform

Artyom Lobanov<sup>1,2</sup>, Timofey Bryksin<sup>1,3</sup>, and Alexey Shpilman<sup>1,2</sup>

<sup>1</sup> JetBrains Research, Saint Petersburg, Russia  
`alexey@shpilman.com`

<sup>2</sup> Higher School of Economics, Saint Petersburg, Russia  
`avlobanov@edu.hse.ru`

<sup>3</sup> Saint Petersburg State University, Saint Petersburg, Russia  
`t.bryksin@spbu.ru`

**Abstract.** Online programming courses are becoming more and more popular, but they still have significant drawbacks when compared to the traditional education system, e.g., the lack of feedback. In this study, we apply machine learning methods to improve the feedback of automated verification systems for programming assignments. We propose an approach that provides an insight on how to fix the code for a given incorrect submission. To achieve this, we detect frequent error types by clustering previously submitted incorrect solutions, label these clusters and use this labeled dataset to identify the type of an error in a new submission. We examine and compare several approaches to the detection of frequent error types and to the assignment of clusters to new submissions. The proposed method is evaluated on a dataset provided by a popular online learning platform.

**Keywords:** MOOC, automatic evaluation, clustering, classification, programming

## 1 Introduction

Recently more and more people get additional education through massive online open courses (MOOC), including programming courses. They are very convenient for students, but you get less feedback on what you are doing wrong since the solutions are usually checked using an automated verification system. In our study, we propose an automatic data-driven method for error type classification that can be used to provide hints for students, rather than just inform them on whether or not their submission has passed all the necessary tests.

The main idea of the proposed approach is to automatically identify and recognize the most common errors through identifying **edit scripts** and analyze these edits through clustering. We use expert evaluation to assign error types to clusters. The general pipeline for the process can be seen in Figure 1.




Fig. 1. General pipeline of the proposed approach.

## 2 Related work

Several papers have investigated topics related to our task, namely representing code changes in a vector form, their clustering, and classification.

Falleri et al. [1] introduced an approach and a tool they called GumTree to generate **edit scripts**: sequences of atomic abstract syntax tree (AST) modifications that turn a source tree into a target tree. Today, GumTree is considered to be a state-of-the-art tool for generating edit scripts and is widely used in different research tasks.

The task of clustering source code changes based on edit scripts was studied in [5]. One of the considered approaches is similar to the one used in our work: edit scripts were generated using ChangeDistiller [2] (a predecessor of GumTree) and a similarity-based clustering algorithm was applied to them.

Another related task is the classification of code changes. In [8], the authors try to detect code changes that are likely to be specific types of refactorings. They define heuristic rules to define each refactoring type. This kind of approach could not be applied to our case because the number of possible errors and their types are not known beforehand. In [4], the authors used an SVM classifier and features such as commit’s metadata, complexity metrics and bag-of-words of the changed code to identify commits that are likely to introduce new bugs. The authors treated code as text using bag-of-words models, whereas working with an AST usually gives more useful information. The somewhat similar idea was implemented in [3], where features of AST changes were used for bug prediction. Their experiments also confirmed that using AST features rather than text-based ones yields better results.

One recent study [9] provides an alternative way to represent edit scripts. The authors employ a deep learning approach to generate a vector of features (embedding) for the edit scripts.

## 3 Overview of the approach

### 3.1 Dataset

The dataset is provided by Stepik<sup>4</sup> and consists of submitted solutions in Java with their metadata, including verification result. We follow the assumption that

<sup>4</sup> Stepik MOOC platform: <http://stepik.org/>

between the first correct solution and the previous (incorrect) one the user fixed a mistake, therefore changes between these versions contain information about a correctable error.

The dataset consists of  $\{\textit{incorrect submission}, \textit{correct submission}\}$  pairs for 2 tasks: 1472 pairs for the problem A (a Java Stream API problem) and 8294 pairs for the problem B (checking double values for equality). The dataset was divided into train/validation/test subsets: 1176/148/148 pairs for the problem A and 6588/200/200 pairs for the problem B respectively.

### 3.2 The pipeline

The pipeline is divided into two stages: training and application. At the training stage, we try to find the most common errors in the incorrect solutions database. To achieve this goal, we cluster edit scripts for incorrect solutions. Edit scripts for the same error are expected to fall into the same cluster. At the next stage, labeled clusters are used to create a classifier. This classifier outputs a type for a new error if it falls into one of the identified clusters, and labels this error as “unknown” otherwise.

To generate edit scripts we used the GumTree library. Since users can fix errors differently, at the clustering step we calculate edit scripts between an incorrect solution and all the correct solutions in the dataset and select the shortest edit scripts. The same procedure is used when we try to identify edit scripts for new incorrect solutions at the classification step.

In this paper, we used classification and clustering algorithms that require only a distance function defined between data points. We considered the following distance metrics for edit scripts:

1. several modifications of the Jaccard similarity coefficient depending on the definition of equality of atomic changes in edit scripts;
2. cosine similarity for the bag-of-words model;
3. cosine similarity for the autoencoder embeddings of the edit scripts.


We use Hierarchical Agglomerative Clustering (HAC) [10] to cluster all solutions edit scripts using one of the distance functions described above. After the clustering is complete, clusters smaller than a certain threshold are removed and others are presented to experts, who label them according to the error type.

When classifying a new incorrect solution, we find the nearest correct one and classify the obtained edit script. The easiest way to choose the right cluster is to find the nearest one. As an alternative, we used the k-nearest neighbors (kNN) [6] method with weighted voting. Since the new object may not belong to any cluster, we provide a user with a hint only if we are sure of the classification accuracy.

## 4 Evaluation

To evaluate various clustering and classification algorithms and their parameters we used the area under the precision-recall curve (PR-AUC) since classification should be tuned according to the particular goals.

In all our experiments we used the same clustering algorithm, but, depending on the values of the hyperparameters, we obtained 96 different clustering patterns for each problem. We don't have a proper way to evaluate the quality of clusters themselves, so we compare the quality of the final classification. The validation dataset was used to compare the quality of approaches and find the best one. All in all, we evaluated 27648 combinations of different parameters. Then, best configurations were applied to the test dataset to get an independent assessment. Approaches based on the cosine similarity of the bag-of-words model and the autoencoder embeddings of the edit scripts demonstrated the best results on our dataset for problems A and B respectively. PR-curves for these classifiers are shown in Figure 2.



**Fig. 2.** Precision-recall curve for best classifiers for test problems A and B.

## 5 Conclusion

In this paper, we present a method for the automatic classification of error types in solutions to programming assignments at an online learning platform. It is based on a notion of an edit script: we cluster these edit scripts at the training stage and classify newly submitted incorrect solutions according to these clusters. Manual labeling of these clusters allows us to provide users with a hint containing an error description. We provide an extensive evaluation of the proposed approach using various clustering methods, edit scripts representations and distance metrics. The evaluation shows that this approach could be successfully implemented in online programming courses at scale.

For the future work, we consider more advanced techniques for embedding atomic changes and edit scripts, e.g., RNN[7]. It is also worthwhile to study the proposed method on a larger number of different problems and to identify characteristics of the dataset that could improve the quality of this approach.

## References

1. Falleri, J.R., Morandat, F., Blanc, X., Martinez, M., Monperrus, M.: Fine-grained and accurate source code differencing. In: ASE (2014)
2. Gall, H.C., Fluri, B., Pinzger, M.: Change analysis with evolizer and changedistiller. *IEEE Software* **26** (2009)
3. Giger, E., Pinzger, M., Gall, H.C.: Comparing fine-grained source code changes and code churn for bug prediction. In: Proceedings of the 8th Working Conference on Mining Software Repositories. pp. 83–92. MSR '11, ACM, New York, NY, USA (2011). <https://doi.org/10.1145/1985441.1985456>, <http://doi.acm.org/10.1145/1985441.1985456>
4. Kim, S., James Whitehead, E., Zhang, Y.: Classifying software changes: Clean or buggy? *Software Engineering, IEEE Transactions on* **34**, 181–196 (04 2008). <https://doi.org/10.1109/TSE.2007.70773>
5. Kreutzer, P., Dotzler, G., Ring, M., Eskofier, B.M., Philippsen, M.: Automatic clustering of code changes. 2016 IEEE/ACM 13th Working Conference on Mining Software Repositories (MSR) pp. 61–72 (2016)
6. Larose, D.T., Larose, C.D.: *Discovering knowledge in data: an introduction to data mining*. John Wiley & Sons (2014)
7. Lipton, Z.C., Berkowitz, J., Elkan, C.: A critical review of recurrent neural networks for sequence learning. arXiv preprint arXiv:1506.00019 (2015)
8. Weißgerber, P., Diehl, S.: Identifying refactorings from source-code changes. 21st IEEE/ACM International Conference on Automated Software Engineering (ASE'06) pp. 231–240 (2006)
9. Yin, P., Neubig, G., Allamanis, M., Brockschmidt, M., Gaunt, A.L.: Learning to represent edits. In: International Conference on Learning Representations (2019), <https://openreview.net/forum?id=BJl6AjC5F7>
10. Zaki, M.J., Meira, W.: *Data mining and analysis: Fundamental concepts and algorithms* (2014)