

Assessing Robustness of ML-Based Program Analysis Tools using Metamorphic Program Transformations

Applis, L.H.; Panichella, A.; van Deursen, A.

Publication date

2021

Document Version

Accepted author manuscript

Published in

IEEE/ACM International Conference on Automated Software Engineering - NIER Track

Citation (APA)

Applis, L. H., Panichella, A., & van Deursen, A. (Accepted/In press). Assessing Robustness of ML-Based Program Analysis Tools using Metamorphic Program Transformations. In *IEEE/ACM International Conference on Automated Software Engineering - NIER Track* IEEE / ACM.

Important note

To cite this publication, please use the final published version (if applicable). Please check the document version above.

Copyright

Other than for strictly personal use, it is not permitted to download, forward or distribute the text or part of it, without the consent of the author(s) and/or copyright holder(s), unless the work is under an open content license such as Creative Commons.

Takedown policy

Please contact us and provide details if you believe this document breaches copyrights. We will remove access to the work immediately and investigate your claim.

Assessing Robustness of ML-Based Program Analysis Tools using Metamorphic Program Transformations

Leonhard Applis
TU Delft
Delft, Netherlands
L.H.Applis@Tudelft.nl

Annibale Panichella
TU Delft
Delft, Netherlands
A.Panichella@Tudelft.nl

Arie van Deursen
TU Delft
Delft, Netherlands
A.Deursen@Tudelft.nl

ABSTRACT

Metamorphic testing is a well-established testing technique that has been successfully applied in various domains, including testing deep learning models to assess their robustness against data noise or malicious input. Currently, metamorphic testing approaches for machine learning (ML) models focused on image processing and object recognition tasks. Hence, these approaches cannot be applied to ML targeting program analysis tasks. In this paper, we extend metamorphic testing approaches for ML models targeting software programs. We present LAMPION, a novel testing framework that applies (semantics preserving) metamorphic transformations on the test datasets. LAMPION produces new code snippets equivalent to the original test set but different in their identifiers or syntactic structure. We evaluate LAMPION against CodeBERT, a state-of-the-art ML model for Code-To-Text tasks that creates Javadoc summaries for given Java methods. Our results show that simple transformations significantly impact the target model behavior, providing additional information on the models reasoning apart from the classic performance metric.

CCS CONCEPTS

• **Software and its engineering** → **Software testing and debugging; Documentation**; • **Computing methodologies** → **Natural language generation; Neural networks**.

KEYWORDS

metamorphic testing, machine learning in software engineering, documentation-generation

ACM Reference Format:

Leonhard Applis, Annibale Panichella, and Arie van Deursen. 2021. Assessing Robustness of ML-Based Program Analysis Tools using Metamorphic Program Transformations. In *Proceedings of The 36th IEEE/ACM International Conference on Automated Software Engineering (ASE 2021)*. ACM, New York, NY, USA, 5 pages. <https://doi.org/10.1145/1122445.1122456>

1 INTRODUCTION

Artificial Intelligence (AI) has been applied to software engineering (SE) to address many tasks, such as fault localization [13], test-case

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.
ASE 2021, 15 - 19 November, 2021, Melbourne, Australia

© 2021 Association for Computing Machinery.
ACM ISBN 978-x-xxxx-xxxx-x/YY/MM... \$15.00
<https://doi.org/10.1145/1122445.1122456>

generation [2], fuzzing [12] or optimizing meta-parameters [20]. Recently, modern sequence-to-sequence deep learning models have shown promising results sparking new types of applications. Among them is the creation of code from verbatim description (*tex-to-code*) [6], or generation of documentation for source-code of previously unseen methods (*code-to-text*) [9, 11]. Yet, we argue that *it is not clear the extent to which these models truly behave as intended*, apart from their reported accuracy. Hence, applying testing strategies for ML-based program analysis solutions is critical.

In recent years, there has been great interest in Testing ML, where the goal is indeed to go beyond assessing accuracy (see the survey by Zhang et al. [28]). Many of the approaches have been taken from *classic* software testing and have been adapted for ML. One example is metamorphic testing, which is a well-established technique that is considered a powerful approach as it addresses the *Oracle Problem* [3] in test generation. Metamorphic testing has been successfully used in ML [15, 26] for image processing and object recognition. For example, image rotation is an information-preserving transformation as it alters the pixels in the image without changing its label (oracle). In computer vision, a *robust* ML model must not provide different predictions for the image altered with metamorphic transformations. Hence, quantifying the number of transformed images on which an ML model provides different answers quantifies its *robustness* against different transformations.

While extensive research has been conducted on metamorphic testing for vision computing tasks [15, 22, 26], the existing metamorphic transformations are domain-specific. Consequently, they cannot be applied and do not hold for different domains and types of data. In this paper, we *extend the concept of metamorphic testing to machine learning models trained on and targeting source code*.

We define a set of transformations that alter features of code but yield the effectively equal program, such as introducing `if(true)`-conditions or `+0` behind integer expressions. Using those, we modify the test-datapoints (programs) in order to detect differences in the models' predictions and metrics. We expect that the models are robust towards some transformations while others affect the metrics (negatively). The information gained could help to evaluate existing models, compare them to each other and provide suggestions and warnings for end-users and researchers alike. With our research and tool, we contribute on the following points:

- (1) Create a systematic approach, namely LAMPION, to quantify the robustness of a source-code-based model
- (2) Enable researchers to compare the *robustness* of models similarly to existing quality metrics
- (3) Groundwork for data augmentation in the field of ML4SE.
- (4) Empirically show the importance of robustness and testing when referring to ML-based program analysis.

To the best of our knowledge, we are the first to propose the use of metamorphic transformations for assessing ML-based program analysis tools. Our initial experiments on CodeBERT [6], a state-of-the-art ML model widely used in the SE literature [10, 14, 17, 23], demonstrate the feasibility of the approach, and the type of lessons that can be learned from applying LAMPION.

2 BACKGROUND AND RELATED WORK

Metamorphic testing is a technique based upon the concept of *metamorphic relations*, which is a property-based technique that exploits known equality of certain output values. Prominent examples are programs that implement mathematical functions; The sine function has a well-known metamorphic relation: $\forall x \in \mathbb{R} : \sin(x) = \sin(x + 2\pi)$. Testers can easily create new test cases based on this relation and assess the program correctness. A broad view of metamorphic testing studies and applications can be found in the survey by Segura et al. [19]. While metamorphic testing has not been applied to ML models for SE, metamorphic transformations and relations are known in software engineering and are tightly coupled to *refactoring*, program *optimization*, and linting. Metamorphic transformations are also used for compiler optimization to create more efficient code, using techniques like loop unrolling or function inlining [5].

Metamorphic Testing for ML. Metamorphic testing has been applied recently to machine learning, especially to image-based object-detection tasks [26][15]. A metamorphic transformation on images performs *information-preserving* alternations on an image. For example, the image of a cat might be mirrored, yet a classifier should still be able to recognize it as such. Blurring or saturating of images [24] change the data significantly; nevertheless, they are still easily classifiable by humans. These transformed images can be used to access robustness by generating more datapoints in the test set [26]. It can also be applied to generate more training data, which can result in a more robust or precise model [22].

The existing literature focuses on MTs that are specific to images and pixels. In this paper, we transplant the testing methodology to a new domain, namely ML models designed for program analysis. This requires defining new metamorphic relations and transformations for our domain, which we describe in Section 3.

Adversarial attacks. Related work stems from Compton et al. [4] that introduces randomization of variable-names in the training dataset of a code2vec model for training data augmentation. Their study shows that the model trained on the augmented training dataset achieves slightly better accuracy than the model trained on the original dataset which motivates to systematically investigate for overfitting. Similarly, Yefet et al. [27] prove that they can generate adversarial attacks on Code2Vec-based classifiers by changing variable names or introducing new variables. As this existing research motivates to inspect identifier names, we include them into our approach in addition to other transformations.

3 OUR FRAMEWORK: LAMPION

Overview. Figure 1 depicts the metamorphic testing approach, we named LAMPION, and designed for testing ML models trained on source-code programs. LAMPION relies on the MTs defined in the subsections below. Our approach consists of three main steps. First,

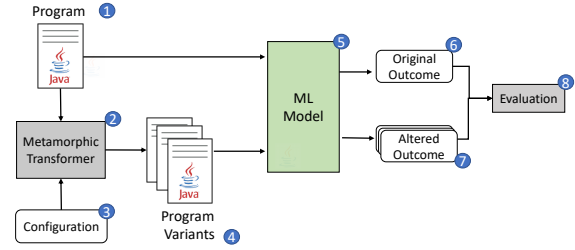


Figure 1: LAMPION— Metamorphic Testing Framework for ML-based Program Analysis

LAMPION takes as inputs a pre-trained model and a program not used during the training process (items ⑤ and ① in Figure 1). It generates program variants (item ④) by using the MTs (item ②) and based on a given configuration file (item ③). The configuration specifies the type of transformation applied and the number of repetitions (order). Then, the original program and its equivalent variants are fed to the pre-trained model. Finally, LAMPION compares the outcome produced by the pre-trained model for the original program (item ⑥) and its metamorphic variants (item ⑦). If there is no difference in the outcome, it means that the model is *robust* to the MT. Otherwise, we found a weakness in the model.

Metamorphic Relation for Programs. The first step is to identify *metamorphic relations* for software programs, which are the data points for ML-based program analysis. *Metamorphic relations* (MRs) relate multiple programs that differ in their structures (e.g., AST) but that are effectively equivalent. As such, ML models should provide the very same output (e.g., same label) for programs that are related to one other according to an MR. Therefore, given a program P , we use MRs to generate equivalent yet different programs P'_1, \dots, P'_k to test a given ML model under analysis.

In ML applications, the oracle function corresponds to the labels that humans provide for a given program P . The type of label for each program (data point) is task-dependent. For example, in ML-based program documentation, the label (oracle) is the natural language description developers write for the program P .

We identify two types of metamorphic relations for programs and that are useful to test ML models for program analysis:

MR-1: Addition of uninformative code elements. Such a code element (e.g., comments, un-used variables, etc.) does not change the behavior of the target program P . As such, the label (oracle) for P and its variants with MR-1 relation remains the same.

MR-2: Replace a code element with another equivalent element. Equivalent program elements (e.g., different variable names) do not change the AST of the programs but the labels of the nodes within the AST. Using different yet equivalent elements does not change the behaviors of a program P either.

Metamorphic Transformations. Given the two MRs defined above, we can define a set of *metamorphic transformations* that satisfy our MRs. A *metamorphic transformation* (MT) is a procedure that generates new programs P'_1, \dots, P'_k (*follow-up* programs) starting from an input program P and using a metamorphic relation. We have two constraints for MTs: First, the oracle function must give the same output for the initial program P and the transformed

Table 1: Overview of metamorphic transformations for programs

Transformation	Short	Description	Estimated Effect	Variations
if-true	MT-IF	Wrapping a random expression in an "if(true)" statement	Structural Changes, introduction of conditions, introduction of keywords	if-false-else
add-unused-variable	MT-UV	Add a random unused variable	Introduction of names, introduction of types	Full random and pseudo random names (Postfix R & P), names looked up from a dictionary or the program under test
rename-entity	MT-RE	Rename a class, method or variable	Introduction of names, removal of known names	For variables, classes and member-types separate
lambda-identity	MT-ID	Wrap an expression in an identity-lambda function (including function call)	Introduction of complex structure, introduction of operators	-
delegation-method	MT-DM	extract an expression to a function, invoke the function instead of the method	structural changes, change of scope for information, introduction of names	same as MT-UV
comment-alternation	MT-CO	Add,remove or move comments	Introduction or removal of natural language	Full or pseudo random comments generated
parameter-introduction	MT-PI	Introduce an unused parameter	Change of method signature, introduction of names, introduction of types	same as MT-UV
whitespace-alternation	MT-WS	Add or remove whitespace	Change of code-layout	-
add-neutral-element	MT-NE	Add the neutral element to a primitively typed expression	Change of structure, introduction of tokens	Complex equivalent transformations (e.g. replacing <i>true</i> with $0^1 == 1$)

program $f(P)$. Second, P is a valid input for the ML model, then $f(P)$ must be valid input for the model too.

A summary of MTs is presented in Table 1. They target various features of the code, such as structure, tokens, and identifier names. Different models are known to have constraints by their design. For example, Code2Vec defines an AST-depth; hence, the model is known to break when introducing many redundant structure elements. Other models – especially deep learning models like CodeBERT – do not specify the features they target and were not previously inspected. Further constraints come from the dataset - if the imports of a program are not known, explicit names cannot be resolved. Similarly, most ML-experiments ignore whitespace completely as part of their preprocessing, ruling out any MTs targeting layout. Hence, the scientist must carefully select fitting MTs per experiment or adjust them accordingly.

The presented table can be considered a starting point for metamorphic transformations applied to ML-based program analysis solutions.

4 EMPIRICAL STUDY

We first want to assess whether the proposed MTs impact the performance of machine learning models. In an ideal case, ML models should not be affected by the metamorphic transformations, i.e., the model is not sensitive to changes that do not alter the code behavior. Hence, RQ1 should cover the general impact of applying one single transformation at the time, hereafter referred to as *first-order* MTs:

Research Question 1

To what extent do first-order metamorphic transformations affect the performance of ML models?

We also want to compare the different types of transformations w.r.t the benchmark. We may expect that different transformations have different impacts on ML models. Furthermore, we aim to understand which model features are more robust, e.g., whether name-changes affect the model *more* than structural AST changes.

Research Question 2

To what extent different types of MTs have a different impact on the performance of ML models?

Benchmark. For an initial study, we picked CodeBERT [6], particularly its downstream task of code-summarization. Code-summarization should clarify *what the model understands*, and the output can give clearer insights than *cold metrics*. We trained a CodeBERT-Java Model as described in the official repository by Microsoft [1], using the standard parameters given in the readme. CodeBERT has been trained on 6 programming languages with a total of 8.3M datapoints (code snippets) and achieves state-of-the-art results of an average BLEU4-Score of 17.65 in the CodeSearchNet-challenge [8].

Methodology / Experiment Design. We developed a metamorphic transformer for Java-Programs that works at the source-code level. In addition, we need a (pretrained) model and an existing benchmark that either consists of .java files or provides sufficient pre- and post-processing to transform the datapoints.

To answer RQ1, we apply MTs to all datapoints in the test set, resulting in a set of variant-code-snippets. We then re-calculate the performance metrics for the variant-code snippets (metamorphic test cases) as well as for the original ones. We use the BLEU4-Score [18] as performance metric, which is the standard metric used in code-to-text and text-to-code generation tasks [10, 18]. The BLEU4-score is computed by tokenizing the gold standard and the generated text into n-grams, and comparing the resulting sets of n-grams. The metric value ranges from 1.0 (perfect translation) to 0 (not a single matching word or n-gram). In addition, we use the Jaccard-distance to measure the percentage of words that differ between the two Javadoc-comments generated by CodeBERT before and after applying an MT.

To assess the significance of the differences in BLEU4-score achieved by the model with and without metamorphic transformations, we use the Wilcoxon rank-sum test [25]. We verified beforehand that the achieved results follow a non-normal distribution by applying the Shapiro-Wilkinson test [21]. We answer RQ2 by grouping the existing results by type of MT. On the MT-groupings we use the Friedman test [7] and the post-hoc Nemenyi test [16]. The Friedman test tests for significant differences among the different MTs in terms of their impact on the BLEU4-score. Furthermore, we use the post-hoc Nemenyi test to perform a pairwise comparison. The Nemenyi test measures the difference across the MTs by computing the average rank of each treatment across all datapoints

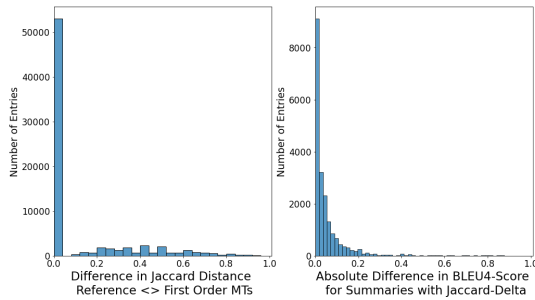


Figure 2: Overview of changes for first-order MTs

in the test set. A lower average rank means that an MT changes the metric more significantly.

5 PRELIMINARY RESULTS

Results for RQ1. Figure 2 shows the histogram of deltas in the BLEU4-Score produced by CodeBERT before and after applying MTs. The Figure also shows the histogram of Jaccard distance between the reference and the post-MT generated JavaDoc. We observe that out of 72,989 produced JavaDoc summaries, for 16,566 datapoints CodeBERT generated summaries with a non-zero delta in BLEU-Score (22.6%). For these 16,566 datapoints with changes, the average difference in BLEU-Score is 0.06. Many summaries change when we apply the MTs, but they perform comparably in terms of BLEU-Scores to the unaltered; For example, both summaries could miss the same number of keywords, just by different tokens. More in detail, there are 52,838 Java methods in the test set out of 72,989 with zero Jaccard distance. This results in 20,151 snippets that do not pass the metamorphic tests (27.6%). Finally, the Wilcoxon rank-sum test revealed that there is a statistically significant (p -value<0.01) difference in the BLEU-Scores achieved by CodeBERT for the code snippets with changes between pre- and post-transformations.

Results for RQ2. We applied the Friedman test and the post-hoc Nemenyi procedure to analyze the impact of the different MTs on the BLEU-Score. With a p -value<0.01, the Friedman test indicates a statistical difference across the different types of MTs. The results of the post-hoc Nemenyi test are reported in Figure 3. From the ranking, we can see that the most impactful transformations are MT-UVR and MT-UVP, while MT-IF and MT-NE are the least impactful on the BLEU-Score. In terms of significance, we can conclude that MT-UVP is statistically more impactful than the other MTs.

6 DISCUSSION

We presented an effective approach for testing the robustness of a model towards metamorphic transformations on source code. According to the empirical results, our approach is capable of producing significant changes in the summaries generated by CodeBERT, highlighting potential weaknesses in the model as it does not satisfy metamorphic relations. In other words, slightly different variants of the same program can lead to significantly different results.

While in this paper we focus on the Code-To-Text tasks of CodeBERT, we expect the found implications to hold true for other down-stream tasks as well. This can be considered a call-to-arms




Figure 3: Results of the Friedman test and Nemenyi post-hoc procedure for different MTs.

for researchers and practitioners to test machine learning models trained on source code using metamorphic testing in addition to the traditional performance metric (e.g., accuracy).

We envision that a handful robustness-criteria are defined for the next generation of ML-based program analysis solutions, documented and tested using metamorphic transformations. We hope to motivate reviewers of future research to perform sanity checks on newly published models and add robustness as a mandatory attribute of being *SOTA*.

LAMPION can also be used to increase the size of the test-set by generating new program variants, without requiring human labeling. This could be potentially beneficial for SE tasks where labeling data is very expensive or few datapoints are available.

7 CONCLUSION

This paper introduces metamorphic relations to test ML-models program analysis solutions. Using this technology, our objective is to gain further information on the model's behavior apart from the performance metric (e.g., accuracy). To achieve this, we presented a generic approach (LAMPION) and applied it in a case study on CodeBERT's Code-To-Text tasks. To evaluate the case study, we perform various statistical tests to prove or disprove changes in the resulting performance metric.

Our approach and framework can empower experts and laymen alike to assess the robustness of their models and provide additional tests on quality. We tried to keep the approach ① lightweight in concept, ② expendable in implementation (due to plug-in MTs), ③ independent of the task (any language and quality metric).

While our initial implementation is in Java, we expect that a re-implementation for any language is an easy task and the statistical analysis can be reused for most experiments.

8 ONLINE RESOURCES

The code for a sample metamorphic transformer, the grid experiment and the evaluation can be found on Github under the Lampion repository¹. The evaluation also contains the results of the experiment as seen in this paper. A reproduction package for the Codebert Experiment and containerization thereof can be found separately on Github². The model, cleaned test-set and post-transformation datasets can be found on SurfDrive³.

¹<https://github.com/ciselab/Lampion>

²<https://github.com/ciselab/CodeBert-CodeToText-Reproduction>

³<https://surfdrive.surf.nl/files/index.php/f/8713322177>

REFERENCES

- [1] 2020. CodeXGLUE: A Benchmark Dataset and Open Challenge for Code Intelligence. (2020).
- [2] Saswat Anand, Edmund K Burke, Tsong Yueh Chen, John Clark, Myra B Cohen, Wolfgang Grieskamp, Mark Harman, Mary Jean Harrold, Phil McMinn, Antonia Bertolino, et al. 2013. An orchestrated survey of methodologies for automated software test case generation. *Journal of Systems and Software* 86, 8 (2013), 1978–2001.
- [3] E. T. Barr, M. Harman, P. McMinn, M. Shahbaz, and S. Yoo. 2015. The Oracle Problem in Software Testing: A Survey. *IEEE Transactions on Software Engineering* 41, 5 (2015), 507–525. <https://doi.org/10.1109/TSE.2014.2372785>
- [4] Rhys Compton, Eibe Frank, Panos Patros, and Abigail Koay. 2020. Embedding Java Classes with code2vec. *Proceedings of the 17th International Conference on Mining Software Repositories* (Jun 2020). <https://doi.org/10.1145/3379597.3387445>
- [5] Keith Cooper and Linda Torczon. 2011. *Engineering a compiler*. Elsevier.
- [6] Zhangyin Feng, Daya Guo, Duyu Tang, Nan Duan, Xiaocheng Feng, Ming Gong, Linjun Shou, Bing Qin, Ting Liu, Daxin Jiang, et al. 2020. Codebert: A pre-trained model for programming and natural languages. *arXiv preprint arXiv:2002.08155* (2020).
- [7] Milton Friedman. 1937. The use of ranks to avoid the assumption of normality implicit in the analysis of variance. *Journal of the american statistical association* 32, 200 (1937), 675–701.
- [8] Hamel Husain, Ho-Hsiang Wu, Tiferet Gazit, Miltiadis Allamanis, and Marc Brockschmidt. 2019. Codesearchnet challenge: Evaluating the state of semantic code search. *arXiv preprint arXiv:1909.09436* (2019).
- [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] Tae-Hwan Jung. 2021. CommitBERT: Commit Message Generation Using Pre-Trained Programming Language Model. *arXiv preprint arXiv:2105.14242* (2021).
- [11] Boao Li, Meng Yan, Xin Xia, Xing Hu, Ge Li, and David Lo. 2020. *DeepCommenter: A Deep Code Comment Generation Tool with Hybrid Lexical and Syntactical Information*. Association for Computing Machinery, New York, NY, USA, 1571–1575. <https://doi.org/10.1145/3368089.3417926>
- [12] Jun Li, Bodong Zhao, and Chao Zhang. 2018. Fuzzing: a survey. *Cybersecurity* 1, 1 (2018), 1–13.
- [13] Xia Li, Wei Li, Yuqun Zhang, and Lingming Zhang. 2019. Deepfl: Integrating multiple fault diagnosis dimensions for deep fault localization. In *Proceedings of the 28th ACM SIGSOFT International Symposium on Software Testing and Analysis*. 169–180.
- [14] Ehsan Mashhadi and Hadi Hemmati. 2021. Applying CodeBERT for Automated Program Repair of Java Simple Bugs. *arXiv preprint arXiv:2103.11626* (2021).
- [15] Christian Murphy, Gail Kaiser, Lifeng Hu, and Leon Wu. 2008. Properties of machine learning applications for use in metamorphic testing. *20th International Conference on Software Engineering and Knowledge Engineering, SEKE 2008* (2008), 867–872.
- [16] Peter Nemenyi. 1963. Distribution-free multiple comparisons PhD thesis Princeton University Princeton.
- [17] Cong Pan, Minyan Lu, and Biao Xu. 2021. An Empirical Study on Software Defect Prediction Using CodeBERT Model. *Applied Sciences* 11, 11 (2021), 4793.
- [18] Kishore Papineni, Salim Roukos, Todd Ward, and Wei-Jing Zhu. 2002. BLEU: a method for automatic evaluation of machine translation. In *Proceedings of the 40th annual meeting of the Association for Computational Linguistics*. 311–318.
- [19] Sergio Segura, Gordon Fraser, Ana B Sanchez, and Antonio Ruiz-Cortés. 2016. A survey on metamorphic testing. *IEEE Transactions on software engineering* 42, 9 (2016), 805–824.
- [20] Saad Shafiq, Atif Mashkoor, Christoph Mayr-Dorn, and Alexander Egyed. 2020. Machine Learning for Software Engineering: A Systematic Mapping. *arXiv preprint arXiv:2005.13299* (2020).
- [21] Samuel Sanford Shapiro and Martin B Wilk. 1965. An analysis of variance test for normality (complete samples). *Biometrika* 52, 3/4 (1965), 591–611.
- [22] Muhammad Sharif, Sajjad Mohsin, Muhammad Younas Javed, and Muhammad Atif Ali. 2012. Single Image Face Recognition Using Laplacian of Gaussian and Discrete Cosine Transforms. *Int. Arab J. Inf. Technol.* 9, 6 (2012), 562–570.
- [23] Michele Tufano, Cody Watson, Gabriele Bavota, Massimiliano Di Penta, Martin White, and Denys Poshyvanyk. 2018. An empirical investigation into learning bug-fixing patches in the wild via neural machine translation. In *Proceedings of the 33rd ACM/IEEE International Conference on Automated Software Engineering*. 832–837.
- [24] C. N. Vasconcelos, A. Paes, and A. Montenegro. 2016. Towards Deep Learning Invariant Pedestrian Detection by Data Enrichment. In *2016 15th IEEE International Conference on Machine Learning and Applications (ICMLA)*. 837–841. <https://doi.org/10.1109/ICMLA.2016.0150>
- [25] Frank Wilcoxon. 1992. Individual comparisons by ranking methods. In *Breakthroughs in statistics*. Springer, 196–202.
- [26] Xiaoyuan Xie, Joshua W.K. Ho, Christian Murphy, Gail Kaiser, Baowen Xu, and Tsong Yueh Chen. 2011. Testing and validating machine learning classifiers by metamorphic testing. *Journal of Systems and Software* 84, 4 (2011), 544 – 558. <https://doi.org/10.1016/j.jss.2010.11.920> The Ninth International Conference on Quality Software.
- [27] Noam Yefet, Uri Alon, and Eran Yahav. 2019. Adversarial Examples for Models of Code. *arXiv preprint arXiv:1910.07517* (2019).
- [28] Jie M Zhang, Mark Harman, Lei Ma, and Yang Liu. 2020. Machine learning testing: Survey, landscapes and horizons. *IEEE Transactions on Software Engineering* (2020).