

Intelligent Automation for Accelerating the Repair of Software Build Failures

Gengyi Sun

Cheriton School of Computer Science

University of Waterloo

Waterloo, Canada

gengyi.sun@uwaterloo.ca

Abstract—Society has an insatiable hunger for software. It keeps our planes in the air, our cars on the road, and even guides surgical procedures. Yet as software enriches more and more aspects of our lives, its complexity (and that of its maintenance) presents an ever-growing challenge. To manage the development of complex software, build systems are widely adopted to perform routine checks after code submissions. The build system lies at the core of the software delivery process, responsible for transforming source code (pseudo-English machine instructions) into release-ready software that users can install or interact with. While build systems provide numerous benefits, the rapid pace of modern software development generates heavy workloads for them to process. Executing builds requires substantial computing resources and energy, and when a build fails, the consequences ripple throughout the development process. Failures not only block others from validating their work, but also necessitate repeated executions, incurring more resource consumption. In a case study of a software organization, 18% of builds failed, with an average of 56 minutes spent resolving each failure. Such inefficiencies contribute to wasted computing resources and energy and hinder productivity, emphasizing the need for more cost-effective solutions.

This proposal aims to develop automated methods for repairing build failures by addressing four key areas: (1) compilation errors, (2) dependency-induced errors, (3) test execution errors, and (4) a comprehensive solution that integrates these aspects. The approach involves parsing and analyzing build action traces to identify root causes, cataloging and exploiting patterns of reusable build fixes to enable rapid resolution, and leveraging machine learning approaches, such as the fine-tuning and prompting of large language models, to assist developers in re-implementing failed test cases. These innovations will streamline the build repair process, reducing delays and improving overall development efficiency.

Index Terms—build system, automated program repair

I. INTRODUCTION

Modern software is multi-faceted. As modern software provides more functionalities, software systems are becoming larger and more complex. They often incorporate a wide range of artifacts, from source code in different languages to various types of non-code artifacts. Our recent work also shows that 97% of a video game repository with over a million files are non-code files [18], highlighting the scale and complexity of modern software. To weave these artifacts into a cohesive whole, the build system is responsible for compiling application code into binaries, managing dependencies, executing

test cases, and creating deliverables. These benefits come at the cost of a heavy workload that incurs the consumption of expensive computing resources. Despite the growing size and complexity of software systems, there is also an increasing demand for faster delivery. Our recent work has shown that the average build time for a video game project at Ubisoft is 20 minutes and 16% of the builds fail [18]. Failures not only block others from validating their work but also necessitate repeated executions, incurring more resource consumption. In addition to the computing resources, energy, and time spent executing the build job, fixing build failures is also time-consuming [5]–[7], [20]. Currently, build failures are typically resolved manually. In a case study at a large organization, 18% of builds failed, with an average of 56 minutes spent fixing these failures [11]. In accelerating the build process, we aim to reduce inefficiencies that contribute to wasted computing resources and energy and hinder productivity, highlighting the need for automated solutions to address build failures in the software industry.

Prior studies reveal that the top three reasons for build failures in the open-source community are compilation errors, dependency issues, and test execution errors [21]. Extensive work has shown that manually addressing the build failure can be challenging and time-consuming [5]–[7], [20].

Although various methods for fixing compilation errors, dependency issues, and test execution errors have been individually studied, there is limited research on an integrated approach to handling build failures as a whole. Existing work typically evaluates fixes at one stage of the build process, but changes at one stage can introduce new errors in later stages due to the inter-dependencies within the build pipeline. For example, updating dependencies may resolve a dependency-induced failure, but also introduce a new failure due to a version incompatibility. Current approaches fall short of addressing these cross-stage issues.

II. RESEARCH OBJECTIVES

To develop a comprehensive automatic build repair system, we propose the following research objectives (ROs):

RO1) Formulate automatic build repair strategies for compilation errors

In open-source projects, 9.1% of build failures are caused by compilation errors [21] which tend to be syntactical in nature and follow recurrent patterns [15]. Previous work has explored learning-based approaches to address this issue [1], [3], [12], [15], [17], [22]. However, existing work does not validate the generated fixes for compilation errors in the subsequent build stages. A compilable program that behaves in an unexpected way may fail the build at the subsequent testing stage, or worse, allow defective software to be released. Thus, repairing compilation errors alone does not guarantee a successful build. Therefore, when evaluating the required accuracy, we will also consider whether the fix generates new failures in subsequent build stages.

To address RO1, we will leverage natural language processing techniques with fine-tuning to accommodate the software engineering domain knowledge to consider test specifications when generating fixes for compilation errors.

RO2) Formulate automatic build repair strategies for dependency errors

Dependency errors account for 7.1% of build failures in open-source projects [21]. Prior efforts have focused on automating the repair of dependency-related failures by injecting or updating third-party libraries [14]. However, these methods can only successfully repair 46% to 54% of the studied broken builds, indicating that there is potential room for further improvement in this area.

To address RO2, we will propose an ensemble approach for repairing dependency errors by incorporating machine learning approaches to complement the existing heuristic approach.

RO3) Formulate automatic build repair strategies for test execution errors

As the most frequent cause of failing builds, test execution failures have been the focus of extensive research. Test execution failures are the leading cause of build failures, responsible for 41.3% of failures in open-source communities [21]. Unlike syntactically incorrect compilation errors, test cases verify if the code behaves as expected (*i.e.*, semantically correct). Additionally, fixing test execution errors can generate fixes that introduce compilation errors [8], [23], [24]. Existing work employs an iterative approach that compiles and tests these candidate fixes until a true fix is found [24].

To address RO3, in addition to the iterative validations, we will develop grammar-constrained methods that take both syntax and semantics into account to reduce the generation of candidate fixes that could not be compiled.

RO4) Formulate integrated strategies for an automated build repair solution

Prior work has focused on fixing compilation errors [1], [3], [12], [15], [17], [22], test cases [8], [9], [13], or specific types of defects in isolation (*e.g.*, build scripts [4], [14]). At the build level, researchers have proposed approaches that provide

hints to help developers to fix build failures [19], [20]. To the best of our knowledge, there have been no attempts to integrate these efforts to repair the entire build holistically. To further automate the build-repairing process, a comprehensive approach should address all aspects of build failures and restore the build to a functional state in full automation.

To address RO4, our approach will fix build errors sequentially, addressing issues in the order they occur (*i.e.*, compilation, dependency specifications, test executions) while iterating through the stages to ensure a complete repair.

III. EXPECTED CONTRIBUTIONS

We will address build failures with learning-based and template-based approaches. With large language models demonstrating promising performance, research increasingly leverages learning-based techniques for automatic program repair [1], [3], [8], [9], [12], [13], [15], [17], [22], whereas in resource-constrained scenarios, template-based methods can be adopted to balance performance and resource demands.

We will construct datasets that contain all types of build failures to train and evaluate the integrated system. By mining the wealth of open-source repositories hosted on social coding platforms (*e.g.*, GitHub), we can reproduce a large sample of build invocations to construct a dataset containing compilation-related, dependency-induced, and test-related build failures and their corresponding fixes. Off-the-shelf datasets such as DeepFix [2] for compilation errors, and Defects4J [10] for test execution errors can be used to train and test our task-specific components. However, there are limited existing benchmarks for dependency-induced failures and fixes. To address this shortfall of data, we will start studying the dependency-induced build failures by constructing a public dataset to fill in the gap.

Using the collected data, we will conduct an empirical analysis to understand developer-generated fixes and propose models for each type of build failure. These models will consider the interdependencies between different stages of the build process. In resource-constrained scenarios, we can propose template-based fixing approaches that leverage the empirical analysis knowledge.

IV. EVALUATION CRITERIA

Fixes can be validated in several ways. Some approaches compare generated fixes to the ground truth, while others argue that multiple valid fixes exist for a single failure [3]. At the test level, executable datasets such as RunBugRun [16] can be adopted to validate the diverse candidate fixes. At the build level, the ultimate validation is whether the build succeeds after applying the fix.

In addition to the existing effectiveness metrics in measuring repair strategies, we will propose new metrics at the build level. These include the number of builds fully fixed, builds fixed at each stage, and the median number of test cases repaired in partially fixed builds. Efficiency will be measured by the time taken to fix the entire build and the time taken to resolve failures at each stage.

REFERENCES

- [1] Kumar Abhinav, Vijaya Sharvani, Alpana Dubey, Meenakshi D'Souza, Nitish Bhardwaj, Sakshi Jain, and Veenu Arora. *RepairNet: Contextual Sequence-to-Sequence Network for Automated Program Repair*, volume 12748 of *Lecture Notes in Computer Science*, page 3–15. Springer International Publishing, Cham, 2021.
- [2] Rahul Gupta, Soham Pal, Aditya Kanade, and Shirish Shevade. Deepfix: fixing common c language errors by deep learning. In *Proceedings of the Thirty-First AAAI Conference on Artificial Intelligence, AAAI'17*, page 1345–1351. AAAI Press, 2017.
- [3] Hossein Hajipour, Apratim Bhattacharyya, and Mario Fritz. Samplefix: Learning to correct programs by sampling diverse fixes. *CoRR*, abs/1906.10502, 2019.
- [4] Foyzul Hassan and Xiaoyin Wang. Hirebuild: an automatic approach to history-driven repair of build scripts. In *Proceedings of the 40th International Conference on Software Engineering, ICSE '18*, page 1078–1089, New York, NY, USA, 2018. Association for Computing Machinery.
- [5] Michael Hilton, Nicholas Nelson, Timothy Tunnell, Darko Marinov, and Danny Dig. Trade-offs in continuous integration: assurance, security, and flexibility. In *Proceedings of the 2017 11th Joint Meeting on Foundations of Software Engineering, ESEC/FSE 2017*, page 197–207, New York, NY, USA, 2017. Association for Computing Machinery.
- [6] Michael Hilton, Timothy Tunnell, Kai Huang, Darko Marinov, and Danny Dig. Usage, costs, and benefits of continuous integration in open-source projects. In *Proceedings of the 31st IEEE/ACM International Conference on Automated Software Engineering, ASE '16*, page 426–437, New York, NY, USA, 2016. Association for Computing Machinery.
- [7] Yang Hong, Chakkrit Tantithamthavorn, Jirat Pasuksmit, Patanamon Thongtanunam, Arik Friedman, Xing Zhao, and Anton Krasikov. Practitioners' challenges and perceptions of ci build failure predictions at atlassian, 2024.
- [8] Li Huang, Bertrand Meyer, Ilgiz Mustafin, and Manuel Oriol. Execution-free program repair. (arXiv:2405.01309), May 2024. arXiv:2405.01309 [cs].
- [9] Nan Jiang, Thibaud Lutellier, and Lin Tan. Cure: Code-aware neural machine translation for automatic program repair. In *2021 IEEE/ACM 43rd International Conference on Software Engineering (ICSE)*, page 1161–1173, Madrid, ES, May 2021. IEEE.
- [10] René Just, Darioush Jalali, and Michael D. Ernst. Defects4j: a database of existing faults to enable controlled testing studies for java programs. In *Proceedings of the 2014 International Symposium on Software Testing and Analysis, ISSTA 2014*, page 437–440, New York, NY, USA, 2014. Association for Computing Machinery.
- [11] Noureddine Kerzazi, Foutse Khomh, and Bram Adams. Why do automated builds break? an empirical study. In *2014 IEEE International Conference on Software Maintenance and Evolution*, pages 41–50, 2014.
- [12] Xueyang Li, Shangqing Liu, Ruitao Feng, Guozhu Meng, Xiaofei Xie, Kai Chen, and Yang Liu. Transrepair: Context-aware program repair for compilation errors. In *Proceedings of the 37th IEEE/ACM International Conference on Automated Software Engineering, ASE '22*, New York, NY, USA, 2023. Association for Computing Machinery.
- [13] Yi Li, Shaohua Wang, and Tien N. Nguyen. Dear: A novel deep learning-based approach for automated program repair. (arXiv:2205.01859), May 2022. arXiv:2205.01859 [cs].
- [14] Christian Macho, Shane McIntosh, and Martin Pinzger. Automatically Repairing Dependency-Related Build Breakage. In *Proc. of the International Conference on Software Analysis, Evolution, and Reengineering (SANER)*, page 106–117, 2018.
- [15] Ali Mesbah, Andrew Rice, Emily Johnston, Nick Glorioso, and Edward Aftandilian. DeepDelta: Learning to repair compilation errors. In *Proceedings of the 2019 27th ACM Joint Meeting on European Software Engineering Conference and Symposium on the Foundations of Software Engineering*, pages 925–936, Tallinn Estonia, August 2019. ACM.
- [16] Julian Aron Prenner and Romain Robbes. Runbugrun – an executable dataset for automated program repair, 2023.
- [17] HyeonTae Seo, Yo-Sub Han, and Sang-Ki Ko. Multifix: Learning to repair multiple errors by optimal alignment learning. In Marie-Francine Moens, Xuanjing Huang, Lucia Specia, and Scott Wen-tau Yih, editors, *Findings of the Association for Computational Linguistics: EMNLP 2021*, page 4850–4855, Punta Cana, Dominican Republic, November 2021. Association for Computational Linguistics.
- [18] Gengyi Sun, Mehran Meidani, Sarra Habchi, Mathieu Nayrolles, and Shane McIntosh. Code Impact Beyond Disciplinary Boundaries: Constructing a Multidisciplinary Dependency Graph and Analyzing Cross-Boundary Impact. In *Proc. of the International Conference on Software Engineering (ICSE)*, page To appear, 2024.
- [19] Carmine Vassallo, Sebastian Proksch, Timothy Zemp, and Harald C. Gall. Un-break my build: assisting developers with build repair hints. In *Proceedings of the 26th Conference on Program Comprehension, ICPC '18*, page 41–51, New York, NY, USA, 2018. Association for Computing Machinery.
- [20] Carmine Vassallo, Sebastian Proksch, Timothy Zemp, and Harald C. Gall. Every build you break: developer-oriented assistance for build failure resolution. *Empirical Software Engineering*, 25(3):2218–2257, May 2020.
- [21] Carmine Vassallo, Gerald Schermann, Fiorella Zampetti, Daniele Romano, Philipp Leitner, Andy Zaidman, Massimiliano Di Penta, and Sebastiano Panichella. A tale of ci build failures: An open source and a financial organization perspective. In *2017 IEEE International Conference on Software Maintenance and Evolution (ICSME)*, pages 183–193, 2017.
- [22] Zhenyu Wu, Deheng Yang, Yan Lei, Huan Xie, Minghua Tang, and Maojin Li. Labelrepair: Sequence Labelling for Compilation Errors Repair. In *2024 IEEE International Conference on Software Analysis, Evolution and Reengineering (SANER)*, pages 860–871, March 2024.
- [23] He Ye, Matias Martinez, Xiapu Luo, Tao Zhang, and Martin Monperrus. Selfapr: Self-supervised program repair with test execution diagnostics. In *Proceedings of the 37th IEEE/ACM International Conference on Automated Software Engineering, ASE '22*. ACM, October 2022.
- [24] He Ye and Martin Monperrus. Iter: Iterative neural repair for multi-location patches. In *Proceedings of the IEEE/ACM 46th International Conference on Software Engineering, ICSE '24*, New York, NY, USA, 2024. Association for Computing Machinery.