

Singapore Management University

Institutional Knowledge at Singapore Management University

Research Collection School Of Computing and
Information Systems

School of Computing and Information Systems

4-2023

ASDF: A Differential testing framework for automatic speech recognition systems

Daniel Hao Xian YUEN

Andrew Yong Chen PANG

Zhou YANG

Chun Yong CHONG

Mei Kuan LIM

See next page for additional authors

Follow this and additional works at: https://ink.library.smu.edu.sg/sis_research



Part of the [Databases and Information Systems Commons](#)

Citation

YUEN, Daniel Hao Xian; PANG, Andrew Yong Chen; YANG, Zhou; CHONG, Chun Yong; LIM, Mei Kuan; and LO, David. ASDF: A Differential testing framework for automatic speech recognition systems. (2023). *Proceedings of the 16th IEEE International Conference on Software Testing, Verification and Validation, Dublin, Ireland, 2023 April 16-20*. 461-463.

Available at: https://ink.library.smu.edu.sg/sis_research/8565

This Conference Proceeding Article is brought to you for free and open access by the School of Computing and Information Systems at Institutional Knowledge at Singapore Management University. It has been accepted for inclusion in Research Collection School Of Computing and Information Systems by an authorized administrator of Institutional Knowledge at Singapore Management University. For more information, please email cherylids@smu.edu.sg.

Author

Daniel Hao Xian YUEN, Andrew Yong Chen PANG, Zhou YANG, Chun Yong CHONG, Mei Kuan LIM, and David LO

ASDF: A Differential Testing Framework for Automatic Speech Recognition Systems

Daniel Hao Xian Yuen,^{*} Andrew Yong Chen Pang,^{*} Zhou Yang,[†] Chun Yong Chong,^{*} Mei Kuan Lim,^{*} David Lo[†]

^{*}School of Information Technology, Monash University Malaysia
{dyue0003, apan0027, chong.chunyong, lim.meikuan}@monash.edu

[†]School of Computing and Information Systems, Singapore Management University
{zyang, davidlo}@smu.edu.sg

Abstract—Recent years have witnessed wider adoption of Automated Speech Recognition (ASR) techniques in various domains. Consequently, evaluating and enhancing the quality of ASR systems is of great importance. This paper proposes ASDF, an Automated Speech Recognition Differential Testing Framework to test ASR systems. ASDF extends an existing ASR testing tool, the CrossASR++, which synthesizes test cases from a text corpus. However, CrossASR++ fails to make use of the text corpus efficiently and provides limited information on how the failed test cases can improve ASR systems. To address these limitations, our tool incorporates two novel features: (1) a *text transformation module* to boost the number of generated test cases and uncover more errors in ASR systems, and (2) a *phonetic analysis module* to identify phonemes that the ASR systems tend to transcribe incorrectly. ASDF generates more high-quality test cases by applying various text transformation methods (e.g., changing tense) to the input text in a failed test case. By doing so, ASDF can utilize a small text corpus to generate a large number of audio test cases, something which CrossASR++ is not capable of. In addition, ASDF implements more metrics to evaluate the performance of ASR systems from multiple perspectives. ASDF performs phonetic analysis on the identified failed test cases to identify the phonemes that ASR systems tend to transcribe incorrectly, providing useful information for developers to improve ASR systems. The demonstration video of our tool is made online at <https://www.youtube.com/watch?v=DzVwfc3h9As>. The implementation is available at <https://github.com/danielyuenhx/asdf-differential-testing>.

I. INTRODUCTION

The growing presence of Automated Speech Recognition (ASR) systems in modern society [1], [2], [3] motivates the need to properly test ASR systems. Researchers have proposed a series of methods to test various artificial intelligence systems (e.g., image classification [4], [5], autonomous driving [6], [7], etc.) from various perspectives (e.g., fairness [8], robustness [9]). Recently, the rise of automated audio test case synthesis [10], [11], [12], [13], [14] has significantly reduced human involvement in the ASR testing process. For example, CrossASR++ [11] is a tool that leverages Text-to-Speech (TTS) services to automatically generate audio files from texts and uses them to test ASR systems. Intuitively, the transcription produced by an ASR system should be equivalent to the text used to generate the audio. Otherwise, a failed test case for an ASR system is uncovered.

Although CrossASR++ demonstrates the capability of uncovering failed test cases successfully, it still has a few limitations. First, it fails to make efficient use of the text corpus; it requires taking as input a large number of texts to find sufficient failed test cases. The variation and volume of generated test cases depend solely on the quality and quantity of the provided text corpus. Such exhaustive selection employed in CrossASR++ is time-consuming and inefficient. Second, CrossASR++ only reports the number of failed test cases uncovered, while more fine-grained information is needed to help developers to improve ASR systems.

To tackle the aforementioned limitations, we propose ASDF, a new Automated Speech Recognition Differential Testing Framework to test ASR systems. Our tool has the following features. It employs a *text transformation module* and leverages known errors in the ASR systems to synthesize one or more audio test cases from a single text. After collecting a small initial set of failed test cases, ASDF transforms the failed texts using various text transformation methods to further generate more test cases. For example, changing the tense of a sentence or substituting error-inducing terms with other words that have similar phonemes. Our experiments show that utilizing the text transformation module can boost the number of failed texts by an average of 22.3%. Details of our experiment results can be found in our GitHub repository [15]. ASDF also conducts phonetic analysis to identify the phonemes that are more challenging for ASR systems to transcribe. The phoneme information can provide useful information for developers to further improve ASR systems and their robustness [16].

II. TOOL DESIGN

ASDF extends the differential testing workflow used in CrossASR++ [17]. Figure 1 illustrates the pipeline used by ASDF to generate test cases. The pipeline involves two iterations. In the first iteration, each text in the corpus is processed as follows. A Text-To-Speech (TTS) service is used to generate audio from the text, which is then fed into the ASR systems to obtain their transcriptions. Then, we perform cross-referencing by comparing the outputs from various ASR systems. If there exists at least one system that can correctly transcribe the audio while a specific ASR system fails to do so, we say that a *failed test case* is uncovered for this specific ASR system.




Fig. 1. The overview of the pipeline to generate test cases in ASDF.

We refer interested readers to our previous paper [11] for a more detailed description of this pipeline. ASDF can be easily configured to test multiple ASR systems. The installation, configuration, and usage instructions for ASDF are made publicly available in our GitHub repository [15].

The failed test cases (i.e., original texts and their corresponding transcriptions) uncovered in the first iteration are collected and analyzed in the second iteration, which includes a text transformation module. In this iteration, a text transformer will mutate the texts of the failed test cases, effectively generating one or more transformed test cases to be appended to the corpus. This stage is crucial in extending the number of test cases available in the original corpus.

Lastly, phonetic analysis is performed on the failed test cases from both iterations. The phonetics of the error-inducing terms found in the test cases will be analyzed to determine the most common phonemes that appear in the failed test cases. Analysis of highly occurring phonemes in error-inducing terms can be utilized to identify the phonemes that are more challenging for ASR systems to transcribe. These error-prone phonemes can be targeted for further improvements or research in the context of ASR transcription.

III. CORE FUNCTIONALITIES

ASDF requires 5 inputs: 1) a text corpus, 2) the output directory of results, 3) the number of texts to be processed, 4) the ASR systems under test, and 5) a text transformation method. The input text corpus should be in the format of a `.txt` file. Each line of text in the file is converted into an audio file by gTTS [18]. Users can specify the number of texts to be processed before the *Text Transformation Stage* begins. The ASR systems currently available in ASDF are DeepSpeech [19], wav2letter [20], and wav2vec2 [21]. Other ASR systems can be easily added and tested by inheriting the ASR abstract class and adhering to the interface rules [17].

Cross-referencing is performed amongst the selected ASR systems to ensure that the test cases are valid and can be determined by at least one ASR service. If specified, a text transformer transforms the original texts of the failed test cases found in the first iteration. This step generates new test cases to be appended to the original corpus, which is used to test ASR systems in the second iteration. Lastly, the test results are outputted to the path as specified.

Currently, ASDF supports multiple transformation methods. The *Homophone Transformation* method first identifies the homophone of the error-inducing term of the failed test case. Subsequently, a new example sentence that contains that homophone is obtained through the WordHoard Python library [22]. This new sentence will be used as the new test case. Details and examples of the other available text transformation methods such as Augmentation, Adjacent Deletion, Plurality Transformation, and Tense Transformation can be found in our replication package [15].

IV. USAGE

Ten metrics are used to evaluate the performance of the ASR systems. A bar chart is also plotted to show the phonemes and their frequency of appearance in error-inducing terms. A `.csv` file is produced, which can be used to perform analyses on the performance of each individual ASR system. Details of the metrics can be found in the GitHub repository [15].

A failed text is defined as an audio file that is incorrectly transcribed by at least one ASR system. It is important to note that cross-referencing is used to filter the low-quality, unrealistic texts resulting from text transformation. A transformed text input must be transcribed correctly by at least one ASR service to be deemed valid (or determinable); otherwise, it is deemed indeterminable and discarded.

The percentage of transformed failed text is the ratio of failed transformed text inputs to the total transformed text inputs. The higher the percentage of transformed failed text, the better the quality of the transformation method, as this indicates that the transformed texts are more challenging for ASR systems to transcribe correctly. A failed test case is defined as a specific text output from an individual ASR service that does not match its corresponding input text.

The percentage of transformed failed cases is the ratio of transformed output from the ASR systems that does not match its corresponding transformed input to the total number of transformed outputs from the ASR systems. The higher the percentage of failed transformed cases, the better the coverage of the transformed texts generated across different ASR systems.

V. CHALLENGES AND FUTURE WORK

ASDF generates an audio test suite using a TTS library. This assumes that the generated audio is interchangeable with the real-world speech of humans, which may not be the case in real-life situations. There would still be differences between human speech and computer-generated audio, namely pronunciation, accent, and tone. As a result, if human speech were to be used instead of computer-generated audio, such inconsistencies may yield different results for our differential testing. Furthermore, more sophisticated text transformation strategies can be incorporated into differential testing to test their effectiveness in revealing erroneous words. Such techniques include a change in grammar, a change in words, and a modification of sentence structure.

REFERENCES

- [1] V. Kėpuska and G. Bohouta, "Next-generation of virtual personal assistants (microsoft cortana, apple siri, amazon alexa and google home)," in *2018 IEEE 8th Annual Computing and Communication Workshop and Conference (CCWC)*, 2018, pp. 99–103.
- [2] N. Yoshimura, H. Yoshida, F. Matulic, and T. Igarashi, "Extending discrete verbal commands with continuous speech for flexible robot control," in *Extended Abstracts of the 2019 CHI Conference on Human Factors in Computing Systems*, ser. CHI EA '19. New York, NY, USA: Association for Computing Machinery, 2019, p. 1–6. [Online]. Available: <https://doi.org/10.1145/3290607.3312791>
- [3] A. B. Kocaballi, J. C. Quiroz, L. Laranjo, D. Rezadegan, R. Kocielnik, L. Clark, Q. V. Liao, S. Y. Park, R. J. Moore, and A. Miner, "Conversational agents for health and wellbeing," in *Extended Abstracts of the 2020 CHI Conference on Human Factors in Computing Systems*, ser. CHI EA '20. New York, NY, USA: Association for Computing Machinery, 2020, p. 1–8. [Online]. Available: <https://doi.org/10.1145/3334480.3375154>
- [4] Z. Yang, J. Shi, M. Asyrofi, and D. Lo, "Revisiting neuron coverage metrics and quality of deep neural networks," in *2022 IEEE International Conference on Software Analysis, Evolution and Reengineering (SANER)*. Los Alamitos, CA, USA: IEEE Computer Society, mar 2022, pp. 408–419. [Online]. Available: <https://doi.ieeecomputersociety.org/10.1109/SANER53432.2022.00056>
- [5] X. Xie, L. Ma, F. Juefei-Xu, M. Xue, H. Chen, Y. Liu, J. Zhao, B. Li, J. Yin, and S. See, "Deephunter: A coverage-guided fuzz testing framework for deep neural networks," in *Proceedings of the 28th ACM SIGSOFT International Symposium on Software Testing and Analysis*, ser. ISSTA 2019. New York, NY, USA: Association for Computing Machinery, 2019, p. 146–157. [Online]. Available: <https://doi.org/10.1145/3293882.3330579>
- [6] C. Gong, Z. Yang, Y. Bai, J. Shi, A. Sinha, B. Xu, D. Lo, X. Hou, and G. Fan, "Curiosity-driven and victim-aware adversarial policies," in *Proceedings of the 38th Annual Computer Security Applications Conference*, ser. ACSAC '22. New York, NY, USA: Association for Computing Machinery, 2022, p. 186–200. [Online]. Available: <https://doi.org/10.1145/3564625.3564636>
- [7] M. Zhang, Y. Zhang, L. Zhang, C. Liu, and S. Khurshid, "Deeproad: Gan-based metamorphic testing and input validation framework for autonomous driving systems," in *Proceedings of the 33rd ACM/IEEE International Conference on Automated Software Engineering*, ser. ASE '18. New York, NY, USA: Association for Computing Machinery, 2018, p. 132–142. [Online]. Available: <https://doi.org/10.1145/3238147.3238187>
- [8] M. H. Asyrofi, Z. Yang, I. N. B. Yusuf, H. J. Kang, F. Thung, and D. Lo, "Biasfinder: Metamorphic test generation to uncover bias for sentiment analysis systems," *IEEE Transactions on Software Engineering*, vol. 48, no. 12, pp. 5087–5101, 2022.
- [9] Z. Yang, J. Shi, J. He, and D. Lo, "Natural attack for pre-trained models of code," in *Proceedings of the 44th International Conference on Software Engineering*, ser. ICSE '22. New York, NY, USA: Association for Computing Machinery, 2022, p. 1482–1493. [Online]. Available: <https://doi.org/10.1145/3510003.3510146>
- [10] P. Ji, Y. Feng, J. Liu, Z. Zhao, and Z. Chen, "Asrtest: Automated testing for deep-neural-network-driven speech recognition systems," in *Proceedings of the 31st ACM SIGSOFT International Symposium on Software Testing and Analysis*, ser. ISSTA 2022. New York, NY, USA: Association for Computing Machinery, 2022, p. 189–201. [Online]. Available: <https://doi.org/10.1145/3533767.3534391>
- [11] M. H. Asyrofi, Z. Yang, and D. Lo, "Crossasr++: A modular differential testing framework for automatic speech recognition," in *Proceedings of the 29th ACM Joint Meeting on European Software Engineering Conference and Symposium on the Foundations of Software Engineering*, 2021, pp. 1575–1579.
- [12] M. H. Asyrofi, Z. Yang, J. Shi, C. W. Quan, and D. Lo, "Can differential testing improve automatic speech recognition systems?" in *2021 IEEE International Conference on Software Maintenance and Evolution (ICSME)*, 2021, pp. 674–678.
- [13] Z. Yang, J. Shi, M. H. Asyrofi, B. Xu, X. Zhou, D. Han, and D. Lo, "Prioritizing speech test cases," 2023. [Online]. Available: <https://arxiv.org/abs/2302.00330>
- [14] S. S. Rajan, S. Udeshi, and S. Chattopadhyay, "Aequivox: Automated fairness testing of speech recognition systems," in *Fundamental Approaches to Software Engineering*, E. B. Johnsen and M. Wimmer, Eds. Springer International Publishing, 2022, pp. 245–267.
- [15] (2023, Jan.) Asr differential testing framework (asdf). Online. GitHub. [Online]. Available: <https://github.com/danielyuenhx/asdf-differential-testing>
- [16] A. Fang, S. Filice, N. Limsopatham, and O. Rokhlenko, "Using phoneme representations to build predictive models robust to asr errors," in *Proceedings of the 43rd International ACM SIGIR Conference on Research and Development in Information Retrieval (SIGIR '20)*, 07 2020, pp. 699–708.
- [17] (2021, Dec.) Crossasr++. Online. GitHub. [Online]. Available: <https://github.com/soarsmu/CrossASRplus>
- [18] (2022, Dec.) gtts. Online. GitHub. [Online]. Available: <https://github.com/pndurette/gTTS>
- [19] (2021, Nov.) Deepspeech. Online. GitHub. [Online]. Available: <https://github.com/mozilla/DeepSpeech>
- [20] (2022, Dec.) wav2letter. Online. GitHub. [Online]. Available: <https://github.com/flashlight/wav2letter>
- [21] wav2vec2. Online. HuggingFace. [Online]. Available: https://huggingface.co/docs/transformers/model_doc/wav2vec2
- [22] J. Bumgarner. (2022, May) Wordhoard. Online. GitHub. [Online]. Available: <https://github.com/johnbumgarner/wordhoard>