



Beyond Levenshtein: Leveraging Multiple Algorithms for Robust Word Error Rate Computations And Granular Error Classifications

Korbinian Kuhn^{1,2}, Verena Kersken¹, Gottfried Zimmermann¹

¹ Stuttgart Media University, Germany

² University of Tübingen, Germany

kuhngo@hdm-stuttgart.de, kersken@hdm-stuttgart.de, gzimmermann@acm.org

Abstract

The Word Error Rate (WER) is the common measure of accuracy for Automatic Speech Recognition (ASR). Transcripts are usually pre-processed by substituting specific characters to account for non-semantic differences. As a result of this normalisation, information on the accuracy of punctuation or capitalisation is lost. We present a non-destructive, token-based approach using an extended Levenshtein distance algorithm to compute a robust WER and additional orthographic metrics. Transcription errors are also classified more granularly by existing string similarity and phonetic algorithms. An evaluation on several datasets demonstrates the practical equivalence of our approach compared to common WER computations. We also provide an exemplary analysis of derived use cases, such as a punctuation error rate, and a web application for interactive use and visualisation of our implementation. The code is available open-source.¹

Index Terms: speech recognition, word error rate, computational metrics

1. Introduction

Transcription accuracy is a central performance evaluation of Automatic Speech Recognition (ASR) models, because it directly measures their ability to convert spoken language into accurate text representations. However, the complexity of speech and language poses challenges in determining accuracy and evaluations extend beyond merely counting errors. Evaluating transcription accuracy involves rating the severity of errors, measuring the impact of punctuation or capitalisation discrepancies, and assessing how errors influence the human-perceived understandability of a text. A critical aspect in the evaluation process is the need for automated metrics that can efficiently assess large datasets to demonstrate the robustness of an ASR model under diverse linguistic scenarios and environmental conditions [1]. Metrics may also be tailored towards specific use cases, such as the correct transcription of numbers and addresses [2].

The standard metric to report accuracy in ASR research is the Word Error Rate (WER) [3, 4, 5]. It represents the average number of transcription errors per 100 words. The underlying algorithm is the Levenshtein distance, which calculates the minimum number of operations (insertions, deletions, or substitutions) needed to transform one string into another [6]. The WER is computed by determining the minimum edit distance on a word level to measure the number of modifications between an ASR generated hypothesis transcript and a manually created "error-free" reference. Alternatively, for languages that do not

use spaces between words (e.g. Chinese), the Character Error Rate (CER) is used.

To make the WER more reliable, it is necessary to reduce errors resulting from non-semantic text differences like capitalisation, abbreviations, or numerical notations. Transcripts are commonly pre-processed before the WER is calculated. Standard normalisation removes newline characters, repeated spaces, punctuation, and capitalisation. Unifying abbreviations, common contractions, spelling differences, and written numbers can further reduce non-semantic differences [7, 5]. A drawback of these modifications is that they eliminate criteria affecting the readability of a transcript. Reading ASR-generated text can be challenging and is further complicated by the presence of orthographic errors [8]. While the Federal Communications Commission (FCC)² mandates correct punctuation and capitalisation, the WER is limited in its ability to measure their accuracy.

Especially in the field of accessibility, the WER has been criticised as a valid measure of accuracy because the importance of individual words to the overall understandability of a text is not considered, and all errors are penalised equally. Several studies report a weak correlation between WER and understandability ratings by human readers [9, 10, 11]. Other measures, such as Match Error Rate (MER) and Word Information Lost (WIL) [12] also rely on the Levenshtein distance, and have similar problems as the WER. General machine translation metrics like BLEU[13] or METEOR[14] are less specific for the evaluation of transcription accuracy, but are often reported complementary to the WER, e.g. to reflect the multilingual capabilities of an ASR model. Although there is no established automated method for calculating punctuation or capitalisation errors, the precision of specific aspects can be quantified using F1-scores or the Slot Error Rate (SER) [15].

More recent approaches use artificial intelligence (AI) to classify the importance of words and the severity of errors [16]. However, a machine learning based metric can inherit biases present in its training data, be susceptible to overfitting, or adapt poorly to domain shifts. Automated-Caption Evaluation (ACE)[17] and its successor ACE2[18], for example, are trained on conversational speech and are less suitable for the evaluation of live television captioning [19]. Metrics such as the NER-model[20] and the Weighted Word Error Rate (WWER)³, require manual effort to rate the severity of each transcription error. Although manual evaluations can be helpful for specific use cases, they lack objectivity, cannot be replicated, and are hard to obtain for large datasets.

²Federal Communications Commission. (2014, Feb. 24). Closed Captioning Quality Report and Order, Declaratory Ruling, FNPRM.

³T. Apone, M. Brooks, T. O'Connell, "Caption Accuracy Metrics Project," 2010.

¹<https://github.com/shuffle-project/beyond-levenshtein>




Figure 1: **Processing pipeline:** The lexer transforms the input texts into a list of tokens, which are further normalised by several text pre-processors. An extended variant of the Levenshtein distance algorithm with compound word detection and variable edit costs determines the shortest route of modifications. Substitutions are further classified as punctuation, capitalisation, or word errors (e.g. suffix or homophone). The route is used for calculating metrics like WER, categorising types of errors, and visualising text differences.

While there is an ongoing search for alternative metrics, the WER remains the commonly used one. We propose an extended Levenshtein distance algorithm that allows the computation of a robust WER while preserving punctuation and capitalisation. We further utilise commonly used algorithms in the field of natural language processing (NLP) to classify transcription errors more granularly. We see the following applications of our approach for future research:

- Calculation of individual or combined metrics based on error types (e.g. word, punctuation, capitalisation, number, ...).
- Detection of performance increases and negative side effects of ASR models complementary to the overall WER.
- Visual analysis of transcription errors and text normalisations through a graphical web interface.

2. Error Rate Computation

The process is visualised in Fig. 1 and explained in detail in the following section.

2.1. Lexer

In typical WER computations, the input strings are split into parts at each space or newline character, e.g. with a regular expression. In contrast, a lexer can decide on more complex rules based on the preceding and following characters. We use a lexer that parses a string into categorised tokens such as words, numbers, punctuation or symbols. Each token contains the raw value with all characters and a normalised value used to compare the equality of two tokens. Non-word characters, such as spaces, tabs, newline characters, em-dashes, or preceding quotation marks are stored as a prefix or suffix of a token. A new token is created if a word character appears after a suffix. Punctuation characters also create new tokens, except for a few exceptions, e.g. when a period is at the end of an abbreviation (like Mrs.) or between two digits (like 3.14). Some common symbols, like per cent or currency signs, also create a new token.

2.2. Normalisation

The initial list of tokens is then processed through various normalisers, which can modify, split, or merge tokens. For example, if a common contraction like "won't" is detected, a token is split into two word tokens "will" and "not". If a token's value is changed, the name of the normalisation is added to that token's list of normalisations for a later analysis. All normalisations are non-destructive, and the original (unmodified) value is stored with the token. The normalisers further extend existing solu-

tions of JiWER⁴ and Whisper⁵.

The following normalisations are applied: Common abbreviations (e.g. "Mr." or "etc.") and contractions (e.g. "I'm" or "gonna") are replaced with their long form. Annotations within parenthesis (e.g. <unknown>, (pause), [unintelligible]) and interjections (e.g. "hmm", "um") are removed. Spelling differences between UK- and US-English are unified to the US variant (e.g. "analyse", "colour"). Diacritics are replaced with ASCII letters (e.g. "ä"). Symbols for currencies and per cent are replaced with their literal expression. The notation of numbers and currencies can lead to multiple non-semantic errors without normalisation (e.g. "two thousand dollars" and "\$2000"). As the normalisation from text to numbers is not trivial, we reimplemented the number normaliser from Whisper, even though it is not error-free.

2.3. Extended Levenshtein Distance

We extend the Levenshtein distance to detect compound words and use variable costs per operation depending on the involved token types. Algorithm is defined as follows:

$$d_{a,b}(i, j) = \min \begin{cases} 0 & \text{if } i = j = 0 \\ d_{a,b}(i-1, j-1) & \text{if } a_i = b_j \\ d_{a,b}(i-x, j-y) & \text{if } a_{[i-x:i]} = b_{[j-y:j]} \\ d_{a,b}(i-1, j) + \text{cost}(a_i) & \text{if } i > 0 \\ d_{a,b}(i, j-1) + \text{cost}(b_j) & \text{if } j > 0 \\ d_{a,b}(i-1, j-1) + \text{cost}(a_i, b_j) & \end{cases} \quad (1)$$

where $d_{a,b}$ is the recursively defined distance function of two token lists a and b . i and j index into the respective list. $a_{[i-x:i]}$ represents a slice of a with length x starting at index $i-x$, to capture compound words. The condition $1 \leq x \leq \min(i, K)$ must hold, where K is the maximum number of compounds. cost is defined depending on the number of parameters as:

$$\text{cost}(c_k) = \begin{cases} 0.5 & \text{if } c_k \in \mathbb{P} \\ 1 & \text{otherwise} \end{cases} \quad (2)$$

$$\text{cost}(a_i, b_j) = \begin{cases} 2 & \text{if } a_i \in \mathbb{A} \wedge b_j \in \mathbb{B} \text{ where } \mathbb{A} \cap \mathbb{B} = \emptyset \\ 0.5 & \text{if } a_i \in \mathbb{P} \wedge b_j \in \mathbb{P} \\ 0.5 & \text{if } \lfloor a_i \rfloor = \lfloor b_j \rfloor \\ 1 & \text{otherwise} \end{cases} \quad (3)$$

where $\text{cost}(\cdot)$ is a function that returns the costs of insertion or deletion for one token and $\text{cost}(\cdot, \cdot)$ returns the costs for the

⁴<https://github.com/jitsi/jiwer>

⁵<https://github.com/openai/whisper>

substitution of two tokens depending on the token type. \mathbb{A} and \mathbb{B} are sets that contain all tokens of a type, where \mathbb{P} is the concrete set of all punctuation. We define $\mathbb{P} = \{., !?; : \}$. $[c_k]$ converts all characters of c_k to lowercase.

K can be adjusted to improve computation speed in balance to the morpheme-to-word ratio of the applied language (lower for analytic and higher for synthetic languages). We used $K = \infty$ in our implementation.

	Hypothesis	Icecream	is	not	essential	for	wellbeing	of	every	one
Reference	-	I	I	I	I	I	I	I	I	I
Ice	D	C_E	S	S	S	S	S	S	S	S
cream	D	C_R	S	S	S	S	S	S	S	S
is	D	D	O	I	I	I	I	I	I	I
essential	D	D	D	S	O	I	I	I	I	I
.	D	D	D	D	D	I	I	I	I	I
For	D	D	D	S	S	S	I	I	I	I
the	D	D	D	S	S	D	S	S	S	S
well-being	D	D	D	S	S	D	C_E	S	S	S
of	D	D	D	S	S	D	D	O	I	I
everyone	D	D	D	S	S	D	D	D	C_R	C_H
!	D	D	D	D	D	D	D	D	D	D

Figure 2: *Backtrace Matrix: The shortest route is found by traversing the operations starting from the bottom right of the matrix. The operations and movements are: O=OK (up left), I=Insertion (left), D=Deletion (up), S=Substitution (up left), C_H=Compound hypothesis (left), C_R=Compound reference (up), C_E=Compound end (up left).*

The Levenshtein distance is calculated by iteratively filling a 2D matrix, which stores the optimal sequence of edits to transform one string into another (or, in our case, two lists of tokens). Each cell in the matrix corresponds to a specific position of the reference and the hypothesis. The matrix is of size $(len(ref) * len(hyp))$, where each cell (i, j) represents the cost of transforming the first i tokens of the reference into the first j tokens of the hypothesis. The value in each cell is the minimum cost among the possible operations.

Compound words are detected if one or more consecutive tokens of the hypothesis match one or more consecutive tokens of the reference (ignoring spaces and hyphens), for example, "icecream" and "ice-cream". Compounds will not be recognised if they are incorrectly separated by punctuation characters. In some cases this approach can also lead to false positives, for example, "a long" and "along".

In contrast to the original Levenshtein distance, modification costs depend on the involved token types. Insertions, deletions, and substitutions of words are also penalised with 1, but capitalisation and punctuation errors are considered less critical and are weighted with 0.5. Substitutions between a word and a punctuation token are penalised with 2. Thus, removing a comma and inserting a word is preferred over substituting a comma to a word.

2.4. Route

After the matrix is filled, the best route (shortest path) is traversed by backtracing the operations matrix from the bottom right and following each operation (see Fig. 2). The determined route consists of elements with either a reference token (deletion), a hypothesis token (insertion), or both (ok, substitution, compound). Every substitution is further classified as punctuation or word by comparing the hypothesis and reference token values. Word substitutions are further classified as capitalisation, compound word, and number errors if possible. If the hypothesis value is partly included in the reference word, it is categorised as prefix, suffix, affix or with the Porter-Stemmer[21] algorithm. Homophones are detected using the Double-Metaphone[22] algorithm.

2.5. Analysis

2.5.1. Metrics

Punctuation and capitalisation accuracy is reported using SER and F1-scores, as there is no standard metric for these evaluations [15]. WER and SER are defined as follows:

$$WER = SER = \frac{S + D + I}{C + S + D} \quad (4)$$

where C is the number of correct tokens, S the number of substitutions (mismatching tokens), D the number of deletions (missing tokens), I the number of insertions (incorrect predictions).

Using the same parameters, F1-score is defined as follows:

$$F1 = \frac{2C}{2C + 2S + D + I} \quad (5)$$

2.5.2. Visualisation

An interactive web application is provided for testing the algorithm with pre-transcribed samples of multiple speech datasets or custom user input.⁶ The interface, shown in Fig. 3, visualises text differences, and normalisations. Metrics such as the WER, SER, and F1-scores are displayed in comparison to the "traditional" WER. Multiple graphs accumulate the classified substitution errors and the number of normalisations. All normalisations can be toggled to examine their impact on the calculated metrics.

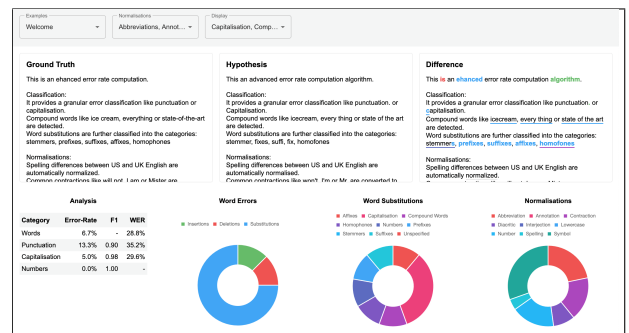


Figure 3: *An interactive web application visualises text differences, error types, and normalisations and calculates several error metrics like WER, SER, and F1-scores.*

⁶<https://shuffle-project.github.io/beyond-levenshtein>

3. Evaluation

We used eight open datasets for long-form English transcriptions covering different scenarios like book readings, colloquial speech, business meetings, and presentations: CORAAL[23], Earnings-21[24], Earnings-22[25], Kincaid46⁷, LibriSpeech-PC[26], Meanwhile[5], Rev16[5], TED-LIUM 3[27]. In total, these datasets contain 745 files and 349 hours of audio.

Two popular ASR frameworks were used to transcribe the datasets: Whisper as an end-to-end (E2E) model (all English and multilingual models) and Vosk as a hybrid model using Kaldi (small-en-us-0.15, en-us-0.22). Punctuation and capitalisation of the Vosk transcripts were created with the vosk-recasepunc-en-0.22 model. In total, 11175 transcripts were generated.

3.1. Word Error Rate Robustness

We evaluated the extended Levenshtein distance algorithm by comparing three different WER computations for all transcriptions: Ours (implemented as described in the previous section), Whisper (using the Whisper text normaliser and JiWER WER computation), JiWER (using the JiWER text normaliser and WER computation).

Table 1 shows the average WER of all transcriptions per dataset. Our implementation reports the lowest WER on average and on the most datasets. While the results are close to Whisper’s text normalisation, they are slightly higher for JiWER, which applies less extensive text pre-processing. The average standard deviation of ours ($SD = 17.5$), Whisper ($SD = 17.3$), and JiWER ($SD = 17.3$) is similar.

Table 1: Average Word Error Rate of all transcriptions for different computation libraries per dataset.

Dataset	Ours	Whisper	JiWER
CORAAL	38.2	38.4	40.2
Earnings-21	20.4	20.5	23.8
Earnings-22	27.6	27.7	31.4
Kincaid46	19.9	19.7	21.3
LibriSpeech-PC	13.7	14.1	14.7
Meanwhile	14.3	14.9	17.1
Rev16	19.1	19.2	20.4
TED-LIUM 3	8.9	9.0	10.2
Average	25.7	26.0	27.9

A one-way ANOVA was conducted to determine the effect of the three different implementations on the resulting WER. There was a significant effect in WER between at least two implementations [$F(2, 33522) = 52.001, p < .001$]. A Tukey’s HSD test showed a significant difference between ours and JiWER ($p < .001$), and between Whisper and JiWER ($p < .001$). There was no significant difference between ours and Whisper ($p = .601$).

Bootstrap resampling with 10000 iterations was conducted to determine 95% confidence intervals for the mean WER difference between the three implementations [28]. The mean difference between ours and Whisper was -0.002 , $SD = 0.007$, 95% CI $[-0.002, -0.002]$, between ours and JiWER was -0.022 , $SD = 0.018$, 95% CI $[-0.022, -0.021]$, and

⁷J. Kincaid, "Which Automatic Transcription Service is the Most Accurate?," 2018.

between Whisper and JiWER was -0.019 , $SD = 0.018$, 95% CI $[-0.020, -0.019]$.

3.2. Punctuation And Capitalisation

As an exemplary use case, table 2 shows the average WER of different ASR models and complementary metrics for punctuation and capitalisation accuracy. Although the WER of the hybrid ASR model vosk is slightly lower than Whisper’s base.en model, the post-processed punctuation and capitalisation are less accurate. A comparison of Whisper’s largest English model with the multilingual large-v3 model shows that the WER is lower for the multilingual model, but the English model performs better on the other metrics.

Table 2: Metrics for word, punctuation, and capitalisation errors of different ASR-models.

Model	Words WER↓	Punctuation SER↓	F1↑	Capitalisation SER↓	F1↑
vosk-en-us-0.22	24.6	80.5	0.51	59.2	0.72
whisper-base.en	25.1	48.8	0.66	35.8	0.82
whisper-medium.en	19.4	43.4	0.68	32.3	0.83
whisper-large-v3	15.9	48.5	0.61	60.6	0.70

4. Conclusion

Despite the ongoing search for alternative measures, WER is the standard metric for reporting the accuracy of ASR. Current ASR-models can achieve very low error rates, which increases the interest in evaluating additional criteria such as punctuation or capitalisation accuracy. However, the WER is limited in its ability to measure their accuracy. We presented an extended Levenshtein distance algorithm to handle punctuation, capitalisation, and compound words. The token-based approach preserves the original text and applies text pre-processing non-destructively.

A statistical analysis using several long-form English datasets showed that our implementation computes similar WER values as commonly used libraries. The results indicate that the calculated WER does not differ significantly between the edit-distance algorithms, but rather between the amount of text pre-processing. Given the small difference between ours and Whisper’s, it is reasonable to conclude that the two implementations demonstrate practical equivalence in computing WER. Thus, the presented algorithm can be used alternatively to calculate a robust WER, whilst allowing the computation of additional metrics such as punctuation accuracy.

An exemplary analysis of these metrics showed that two ASR-models can achieve a similar WER but differ significantly in punctuation and capitalisation accuracy. Further research is required to test the practicability of these additional measures for evaluating and optimising ASR-models and if they can support the definition of accessibility requirements for ASR-generated transcriptions.

5. Acknowledgements

We thank Daniel Griebhaber for his valuable thoughts on the extended Levenshtein distance algorithm. This work was conducted as part of the SHUFFLE Project and funded by "Stiftung Innovation in der Hochschullehre".

6. References

- [1] R. Geirhos, J.-H. Jacobsen, C. Michaelis, R. Zemel, W. Brendel, M. Bethge, and F. A. Wichmann, "Shortcut learning in deep neural networks," *Nature Machine Intelligence*, vol. 2, no. 11, pp. 665–673, 2020.
- [2] A. Aksënova, D. van Esch, J. Flynn, and P. Golik, "How might we create better benchmarks for speech recognition?" in *Proceedings of the 1st Workshop on Benchmarking: Past, Present and Future*. Online: Association for Computational Linguistics, aug 2021, pp. 22–34.
- [3] A. Baeviski, H. Zhou, A. Mohamed, and M. Auli, "Wav2vec 2.0: A framework for self-supervised learning of speech representations," in *Proceedings of the 34th International Conference on Neural Information Processing Systems*, ser. NIPS'20. Red Hook, NY, USA: Curran Associates Inc., 2020.
- [4] Y. Zhang, D. S. Park, W. Han, J. Qin, A. Gulati, J. Shor, A. Jansen, Y. Xu, Y. Huang, S. Wang, Z. Zhou, B. Li, M. Ma, W. Chan, J. Yu, Y. Wang, L. Cao, K. C. Sim, B. Ramabhadran, T. N. Sainath, F. Beaufays, Z. Chen, Q. V. Le, C.-C. Chiu, R. Pang, and Y. Wu, "Bigssl: Exploring the frontier of large-scale semi-supervised learning for automatic speech recognition," *IEEE Journal of Selected Topics in Signal Processing*, vol. 16, no. 6, pp. 1519–1532, 2022.
- [5] A. Radford, J. W. Kim, T. Xu, G. Brockman, C. Mcleavey, and I. Sutskever, "Robust speech recognition via large-scale weak supervision," in *Proceedings of the 40th International Conference on Machine Learning*, ser. Proceedings of Machine Learning Research, A. Krause, E. Brunskill, K. Cho, B. Engelhardt, S. Sabato, and J. Scarlett, Eds., vol. 202. PMLR, 23–29 Jul 2023, pp. 28 492–28 518.
- [6] V. I. Levenshtein, "Binary codes capable of correcting deletions, insertions, and reversals," in *Soviet physics doklady*, vol. 10, no. 8. Soviet Union, 1966, pp. 707–710.
- [7] A. Koenecke, A. Nam, E. Lake, J. Nudell, M. Quartey, Z. Mengesha, C. Touns, J. R. Rickford, D. Jurafsky, and S. Goel, "Racial disparities in automated speech recognition," *Proceedings of the National Academy of Sciences*, vol. 117, no. 14, pp. 7684–7689, 2020.
- [8] J. Butler, "Perspectives of deaf and hard of hearing viewers of captions," *American Annals of the Deaf*, vol. 163, no. 5, pp. 534–553, 2019.
- [9] Y.-Y. Wang, A. Acero, and C. Chelba, "Is word error rate a good indicator for spoken language understanding accuracy," in *2003 IEEE Workshop on Automatic Speech Recognition and Understanding (IEEE Cat. No.03EX721)*. St Thomas, VI, USA: IEEE, 2003, pp. 577–582.
- [10] T. Mishra, A. Ljolje, and M. Gilbert, "Predicting human perceived accuracy of asr systems." in *INTERSPEECH 2011, 12th Annual Conference of the International Speech Communication Association, Florence, Italy, August 27-31, 2011*. Florence, Italy: ISCA, 08 2011, pp. 1945–1948.
- [11] B. Favre, K. Cheung, S. Kazemian, A. Lee, Y. Liu, C. Munteanu, A. Nenkova, D. Ochei, G. Penn, S. Tratz, C. Voss, and F. Zeller, "Automatic human utility evaluation of asr systems: Does wer really predict performance?" in *INTERSPEECH 2013, 14th Annual Conference of the International Speech Communication Association, Lyon, France, August 25-29, 2013*. Lyon, France: ISCA, 01 2013, pp. 3463–3467.
- [12] A. Morris, V. Maier, and P. Green, "From wer and ril to mer and wil: improved evaluation measures for connected speech recognition," in *Eighth International Conference on Spoken Language Processing*. Jeju Island, Korea: ISCA, 10 2004, pp. 2765–2768.
- [13] K. Papineni, S. Roukos, T. Ward, and W.-J. Zhu, "Bleu: a method for automatic evaluation of machine translation," in *Proceedings of the 40th annual meeting of the Association for Computational Linguistics*, 2002, pp. 311–318.
- [14] S. Banerjee and A. Lavie, "Meteor: An automatic metric for mt evaluation with improved correlation with human judgments," in *Proceedings of the acl workshop on intrinsic and extrinsic evaluation measures for machine translation and/or summarization*, 2005, pp. 65–72.
- [15] V. Păiș and D. Tufiş, "Capitalization and punctuation restoration: a survey," *Artificial Intelligence Review*, vol. 55, no. 3, pp. 1681–1722, Mar 2022.
- [16] A. A. Amin, S. Hassan, M. Huenerfauth, and C. O. Alm, "Modeling word importance in conversational transcripts: Toward improved live captioning for deaf and hard of hearing viewers," in *Proceedings of the 20th International Web for All Conference*, ser. W4A '23. New York, NY, USA: Association for Computing Machinery, 2023, p. 79–83.
- [17] S. Kafle and M. Huenerfauth, "Evaluating the usability of automatically generated captions for people who are deaf or hard of hearing," in *Proceedings of the 19th International ACM SIGACCESS Conference on Computers and Accessibility*, ser. ASSETS '17. New York, NY, USA: Association for Computing Machinery, 2017, p. 165–174.
- [18] —, "Predicting the understandability of imperfect english captions for people who are deaf or hard of hearing," *ACM Transactions on Accessible Computing*, vol. 12, no. 2, jun 2019.
- [19] T. Wells, D. Christoffels, C. Vogler, and R. Kushalnagar, "Comparing the accuracy of ace and wer caption metrics when applied to live television captioning," in *Computers Helping People with Special Needs: 18th International Conference, ICCHP-AAATE 2022, Lecco, Italy, July 11–15, 2022, Proceedings, Part I*. Berlin, Heidelberg: Springer-Verlag, 2022, p. 522–528.
- [20] P. Romero-Fresco and J. M. Pérez, *Accuracy Rate in Live Subtitling: The NER Model*. London, UK: Palgrave Macmillan UK, 2015, pp. 28–50.
- [21] M. F. Porter, "An algorithm for suffix stripping," *Program*, vol. 14, no. 3, pp. 130–137, Jan 1980.
- [22] L. Philips, "The double metaphone search algorithm," *C/C++ Users J.*, vol. 18, no. 6, p. 38–43, jun 2000.
- [23] K. Gunter, C. Vaughn, and T. Kendall, "Contextualizing/s/retraction: Sibilant variation and change in washington dc african american language," *Language Variation and Change*, vol. 33, no. 3, pp. 331–357, 2021.
- [24] M. Del Rio, N. Delworth, R. Westerman, M. Huang, N. Bhandari, J. Palakapilly, Q. McNamara, J. Dong, P. Želasko, and M. Jetté, "Earnings-21: A Practical Benchmark for ASR in the Wild," in *Proc. Interspeech 2021*, 2021, pp. 3465–3469.
- [25] M. D. Rio, P. Ha, Q. McNamara, C. Miller, and S. Chandra, "Earnings-22: A practical benchmark for accents in the wild," 2022.
- [26] A. Meister, M. Novikov, N. Karpov, E. Bakhturina, V. Lavrukhin, and B. Ginsburg, "Librispeech-pc: Benchmark for evaluation of punctuation and capitalization capabilities of end-to-end asr models," 2023.
- [27] F. Hernandez, V. Nguyen, S. Ghannay, N. Tomashenko, and Y. Estève, "Ted-lium 3: Twice as much data and corpus repartition for experiments on speaker adaptation," in *Speech and Computer*, A. Karpov, O. Jokisch, and R. Potapova, Eds. Cham: Springer International Publishing, 2018, pp. 198–208.
- [28] B. Efron and R. J. Tibshirani, *An introduction to the bootstrap*. CRC press, 1994.