

MODEL-BASED APPROACH FOR MEASURING THE FAIRNESS IN ASR

Zhe Liu, Irina-Elena Veliche, Fuchun Peng

Facebook AI, Menlo Park, CA, USA

ABSTRACT

The issue of fairness arises when the automatic speech recognition (ASR) systems do not perform equally well for all subgroups of the population. In any fairness measurement studies for ASR, the open questions of how to control the confounding factors, how to handle unobserved heterogeneity across speakers, and how to trace the source of any word error rate (WER) gap among different subgroups are especially important - if not appropriately accounted for, incorrect conclusions will be drawn. In this paper, we introduce mixed-effects Poisson regression to better measure and interpret any WER difference among subgroups of interest. Particularly, the presented method can effectively address the three problems raised above and is very flexible to use in practical disparity analyses. We demonstrate the validity of proposed model-based approach on both synthetic and real-world speech data.

Index Terms— Automatic speech recognition, fairness, Poisson regression, random effect

1. INTRODUCTION

Automatic speech recognition (ASR) systems are getting better with the advent of new technologies. However, the issue of *fairness* arises when these tools do not perform equally well for all subgroups of the population [1, 2, 3, 4, 5]. The concern of fairness is not limited to speech recognition tasks, but also comes to light in other machine learning applications, including facial recognition [6, 7], natural language processing [8, 9], and healthcare [10].

The fairness issue was highlighted most recently by authors in [4] who found that five state-of-the-art ASR systems showed substantial racial disparities, having an average word error rate (WER) of 0.35 for black speakers compared with 0.19 for white speakers. Here, WER serves as the most widely used metric for measuring the performance of an ASR system, which is derived from the Levenshtein distance [11] working at the word level:

$$\text{WER} = \frac{\sum_{s=1}^n e_s}{\sum_{s=1}^n m_s} \quad (1)$$

where m_s is the number of words in the s th sentence (i.e. reference text of audio) of the evaluation dataset, and e_s represents the sum of insertion, deletion, and substitution errors computed from the dynamic string alignment of the recognized word sequence with the reference word sequence.

In most of previous research studies on measuring fairness in ASR, WER was computed per each subgroup (e.g. black speakers versus white speakers) and from the comparison of these WER numbers, conclusions were drawn on whether significant disparities exist among these subgroups of interest. Although this is a very simple way for fairness measurement, there are several open questions that have not yet been properly addressed in these analyses.

First, how to effectively control the *confounding* factors which may affect the measured results but are not of primary interest? For example, we need to deal with any unbalanced gender or age distribution of speakers in a racial disparity study; otherwise, it would be difficult to tell whether any WER gap among different racial groups is due to the race factor or any confounding factor of gender or age. Propensity-score matching was utilized in [4] to select a subset of audio snippets of white and black speakers with similar distributions of age and gender. However, this matching procedure suffers from discarding “unmatched” but informative samples from the analysis and in some scenarios, a good matching might not even exist.

Secondly, how to appropriately take into account speaker-level effect on measured WERs and handle unobserved *heterogeneity* across different speakers? Speech recognition accuracy on utterances from the same speaker could be highly correlated and a neglect of such dependency structure in speech data might lead to underestimated variance and wrong conclusions. Moreover, it is typically reasonable to consider speakers were being randomly selected from the subgroups of interest, and we are not particularly interested in these specific speakers, but in the population that they represent.

Third, how to efficiently trace the source of any WER gap among different subgroups, that is, does such disparity mainly come from phonetic, phonological, prosodic characteristics, or grammatical, lexical, semantic characteristics, or both? Perplexity was used in existing literature to evaluate the grammatical, lexical, semantic properties of disparities across different subgroups. However, more advanced approaches are still called for to provide deeper insights and analyses on whether language model or acoustic model should account for the overall disparities in WERs if at all.

In this paper, we present a model-based approach to better measure the fairness issue in ASR and study any performance disparities across different subgroups of our interest. In particular, we introduce *mixed-effects Poisson regression* [12, 13, 14], treating utterance-level word errors as the regression response, logarithm number of words in the reference text as an offset, speaker identification as a random effect, subgroup label of interest and any other explanatory or confounding variables as fixed effects. The presented method can address the three problems that we previously raised, and is very flexible to use. As classical and powerful statistical tools, mixed-effects model and Poisson regression are not new in analyzing real-world scientific problems. But to the best of our knowledge, our work is the first to introduce sophisticated statistical regression-based approach to investigate fairness issues in ASR and illustrate how it helps measure and interpret any WER difference across different subgroups of the population in any disparity study. In particular, our proposed method prevents underestimating the standard errors and avoids drawing false positive conclusions on non-fairness.

The rest of this paper is structured as follows. Section 2 introduces the use of mixed-effects Poisson regression on ASR fairness. Sections 3 and 4 demonstrate the validity of the proposed method on synthetic and real-world speech data. We conclude in Section 5.

2. METHODS

In this section, we present mixed-effects Poisson regression method and illustrate how it helps measure any WER gap between different subgroups in disparity studies.

Suppose we would want to investigate the fairness in ASR with respect to some factor variable of primary interest (e.g. gender of speakers). For the s th utterance in the evaluation dataset, we denote its factor level as $f(s)$ (e.g. male speaker or female speaker), where f is a deterministic function with l as the total number of levels. We aim to test whether the effect of this factor is statistically significant on measured WER results across its different groups of levels.

2.1. Poisson Regression for Measuring Fairness

Poisson regression serves as an appropriate approach to model rate data [15], where the rate is a count of events (e.g. word errors in our use case) divided by some measure of that unit’s exposure (e.g. number of words in the reference). An *offset* variable is needed to scale the modeling of the mean parameter in Poisson regression with a log link. Here, the underlying assumption is that the number of word errors occurred in any utterance is proportional to the number of words in the corresponding reference text.

More specifically, to measure the effect of factor $f(\cdot)$ on WER results across l different subgroups, the vanilla Poisson regression model is described as follows:

$$C_s \stackrel{\text{i.i.d.}}{\sim} \text{Poisson}(\lambda_s) \quad (2)$$

$$\log(\lambda_s) = \log(N_s) + \mu_{f(s)} \quad (3)$$

where C_s is the count of word errors (sum of insertion, deletion, and substitution errors), λ_s is the Poisson (mean) parameter, N_s is the number of words in the reference text for the s th utterance in the evaluation dataset, and $\mu_{f(s)}$ refers to the factor effect corresponding to the subgroup of $f(s)$. The notation of i.i.d in (2) represents independent and identically distributed, where we will revisit this distribution assumption later in this section. Note that any utterance with empty reference text should be removed from the analysis since it does not provide any insight on fairness measurement.

This model can be fitted using maximum likelihood approach. Standard statistical testing, for example, likelihood ratio test (LRT) [16], can be conducted afterwards to compute the p -value of the factor $f(\cdot)$ on measured WER results.

Sometimes, it is possible to analyze rate data using a binomial response model. However, in our application, number of word errors occurred in some utterance could be larger than the total number of words in the reference, which limits the use of binomial regression here. If the rate is relatively small, the Poisson approximation to the binomial is effective.

One of the key features of Poisson distribution is that the variance equals the mean. In certain circumstances, it is found that the empirical variance is greater than the mean, known as *overdispersion* [17, 18]. A common reason is the omission of relevant explanatory variables, or the present of dependent samples, which we will explore more in the next two subsections.

2.2. Poisson Regression with Explanatory Covariates

It is natural and flexible to extend the vanilla Poisson regression model (2) (3) to include additional explanatory or confounding covariates, which can be utilized to capture effects of confounding variables on WERs among different subgroups:

$$\log(\lambda_s) = \log(N_s) + \mu_{f(s)} + \theta^T x_s \quad (4)$$

Here, x_s represents the vector of any explanatory variables in the regression model and θ refers to the coefficient parameter vector that shall be learned. For example, in a racial disparity analysis, we would want to add the gender or age information of speakers to the regression model in order to control any confounding effects.

In particular, we can include any representative vector [19], for example, sentence embedding, of the true reference text per each utterance as extra explanatory variables, which would help us understand the source of any performance gap between different subgroups of interest. For instance, after controlling the effect of sentence embedding covariates that account for grammatical, lexical, or semantic characteristics, if the factor effect of interest is still statistically significant, we can tell that phonetic, phonological, or prosodic characteristics substantially contribute to the overall disparities among different subgroups of the factor $f(\cdot)$. Thus this can provide insights on whether language model or acoustic model should be responsible for the overall disparities in WERs if at all.

2.3. Mixed-Effects Poisson Regression

Block-structured evaluation data arises naturally in any real-world speech recognition applications. In particular, utterances from the same speaker could share common correlated features (e.g. accent of speaker), and thus analyses that assume independence of these observations will be inappropriate. The use of random effect [13, 14] is one usual and convenient way to model such structure.

Suppose we want to investigate the effect of race on speech recognition accuracy across a sample of speakers. Typically, we would treat the racial effect as fixed in the regression. On the other hand, it makes more sense to treat the speaker effect as random. It is reasonable to consider these speakers as being randomly selected from a larger collection of speakers whose characteristics we would like to estimate. We are not particularly interested in these specific speakers, but in the whole population. Generally, blocking factors can often be viewed as random effects.

A mixed-effects Poisson regression is a model containing both fixed effect and random effect. Regarding the fairness measurement of speech recognition accuracy among different subgroups of the factor $f(\cdot)$, we describe the model in detail as follows:

$$r_i \stackrel{\text{i.i.d.}}{\sim} \mathcal{N}(0, \sigma^2) \quad (5)$$

$$C_{ij} | \lambda_{ij} \stackrel{\text{i.i.d.}}{\sim} \text{Poisson}(\lambda_{ij}) \quad (6)$$

$$\log(\lambda_{ij}) = \log(N_{ij}) + \mu_{f(i)} + r_i + \theta^T x_{ij} \quad (7)$$

where the utterance-level index of subscription notation ij represents the j th utterance from the i th speaker, r_i denotes the speaker-level random effect that is independently sampled from a Gaussian distribution with mean 0 and variance σ^2 which is learnable. Note that any C_{ij} and $C_{i'j'}$ are no longer independent for $j \neq j'$ since they are observed from the same speaker i , while any C_{ij} and $C_{i'j'}$ are still independent for $i \neq i'$ since they are observed from different speakers. Also, we use $\mu_{f(i)}$ to denote the fixed effect for the factor $f(\cdot)$ of primary interest, since typically it is at speaker level.

This mixed-effects model can be fitted via maximum likelihood and the expression for its likelihood is an integral over the random effect, which must be approximated, for example, via adaptive Gauss-Hermite quadrature [20]. Again, LRT can be performed to calculate the p -value of the factor $f(\cdot)$ on measured WER results. In practice, it would be particularly interesting to extract the conditional modes of the speaker-level random effect for subsequent analysis and assumption verification.

3. SIMULATION EXPERIMENTS

In this section, we conduct simulation experiments to show that the proposed mixed-effects Poisson regression could properly address the problems of confounding factor and speaker effect in ASR fairness measurements.

3.1. Experiment on Confounding Factor

We generate synthetic data to investigate the effect of confounding factor on ASR fairness measurements over *case* group and *control* group, defined by some primary factor of interest.

Under the scenario that recognition errors from different utterances are independent from each other, the number of errors on the s th utterance is randomly sampled from a Poisson distribution with the mean parameter written as

$$\lambda_s = N_s \cdot \exp(\mu_{f(s)} + \theta_s \cdot \text{Bernoulli}(p_{f(s)})) \quad (8)$$

where $f(s) \in \{\text{case}, \text{control}\}$ indicates which group the utterance comes from, N_s denotes the number of words in the reference, $\mu_{f(s)}$ refers to the group effect, θ_s represents the effect of confounding factor, and $p_{f(s)}$ is the mean parameter of a Bernoulli distribution which controls the frequency that the confounding effect is present in the corresponding group.

In our experiment, we set $N_s = 10$, $\mu_{f(s)} = \log(0.05)$, $\theta_s = 0.1$ for every s , and $p_{f(s)}$ is varied at 50%, 60%, 70%, 90% for the case group, and 50%, 40%, 30%, 10% for the control group, respectively. For each of case or control group, we generate 5,000 utterances independently.

Here, we would like to evaluate the ratio of WERs between the case and control groups, and in particular, conduct statistical testing to determine whether significant difference on WERs exists between the two groups. Based on our setup, the ground truth WER ratio is 1.0, that is, in theory there is no WER difference between the two groups. Notice that the presence of confounding factor could introduce confounding and mislead the results since it raises up the mean number of errors by $\exp(0.1) - 1 \approx 11\%$ at utterance level.

In this study, we consider the **baseline** measurement method as the one that computes the ratio of empirical WER of case group over the one of control group. The *bootstrap* method [21, 22] is applied to compute the 95% confidence interval (CI) of the ratio. Then if the CI does not cover the point of 1.0, we claim the WER gap between the two groups is statistically significant. Regarding **model-based** approach, we fit a Poisson regression according to (2) (4) which linearly incorporates the confounding factor, and then compute the 95% CI associated with the group effect ratio.

For each approach, we repeat the simulation for 1,000 times and compare the average of estimated WER ratios as well as the false positive rate, that is, the frequency of times that the statistical significance on WER ratio is falsely claimed by the method. Strictly speaking, a 95% CI means that if we were able to have 100 different datasets from the same distribution of the original data and compute a 95% CI based on each of these datasets, then approximately 95 of these 100 CIs will contain the true value of the statistic of interest [23, 24, 25]. Thus in theory, we expect 5% false positive rate if the method works correctly and generates valid CIs.

The result is shown in Table 1, where we can see that the mean ratio and false positive rate of the baseline method increase dramatically when the confounding rate differs more and more between case and control groups. This is expected since the baseline method does not take into account the information of confounding factor which does harm to the inference. For model-based approach, we observe

Table 1. Simulation result on confounding factor experiment with various confounding rates p_{case} and p_{control} across groups.

Confounding Rate		Baseline		Model-Based	
within Case	within Control	Mean Ratio	% False Positive	Mean Ratio	% False Positive
50%	50%	1.000	4.9%	1.000	4.7%
60%	40%	1.021	12.1%	1.001	5.8%
70%	30%	1.041	29.8%	1.000	5.4%
90%	10%	1.084	83.3%	1.001	5.1%

the mean ratios are around 1.0 and the false positive rates are around 5% for all the setups, which demonstrates that it can successfully address the confounding effect and result in valid estimates of WER ratios and corresponding CIs.

3.2. Experiment on Speaker Effect

In this experiment, we generate synthetic data to study the impact of speaker effect on ASR fairness measurements of the two groups.

For any of case or control group, assume there are I distinct speakers and each speaker has equal number of utterances. For the i th speaker and j th utterance from the speaker, the number of errors is sampled from a Poisson distribution with the mean parameter written as

$$\lambda_{ij} = N_{ij} \cdot \exp(\mu_{f(i)} + r_i), \quad r_i \stackrel{\text{i.i.d.}}{\sim} \mathcal{N}(0, \sigma^2) \quad (9)$$

where $f(i)$ indicates which group the speaker is from, N_{ij} denotes the number of words in the reference, $\mu_{f(i)}$ refers to the group effect, and r_i represents the speaker effect drawn from a Gaussian distribution with mean 0 and standard deviation σ .

In our experiment, we set $N_{ij} = 10$, $\mu_{f(i)} = \log(0.05)$ for every i, j , number of speakers I is varied at 100, 500 and standard deviation σ is varied at 0.2, 0.4. For any of case or control group, we generate 5,000 utterances.

Again, we want to evaluate the ratio of WERs and the ground truth shall be 1.0. The baseline method is the same with the one used for confounding factor experiment while for the model-based approach, we fit a mixed-effects Poisson regression according to (5) (6) (7) which treats the speaker identification as a random effect with learnable σ .

The result is shown in Table 2, where we can see that the mean ratios of both methods are around 1.0. However, for the baseline method, we observe high false positive rates and in particular, the higher the standard deviation σ or the smaller the number of speakers, the larger the false positive rate. Instead, the model-based approach always results in approximate 5% false positive rate. This demonstrates that it can successfully deal with speaker effect and is superior than the traditional baseline method.

4. REAL DATA EXPERIMENTS

In this section, we apply the proposed mixed-effects Poisson regression on real-world speech datasets for fairness investigation.

4.1. Datasets and Setup

We consider the following two ASR datasets in the experiments:

Table 2. Simulation result on speaker effect experiment with various numbers of speakers and values of standard deviation σ .

Speaker Effect		Baseline		Model-Based	
Num of Speakers	Standard Deviation	Mean Ratio	% False Positive	Mean Ratio	% False Positive
500	0.2	1.000	8.0%	1.000	4.8%
500	0.4	1.001	14.9%	1.001	4.5%
100	0.2	1.000	16.6%	1.000	5.0%
100	0.4	0.999	42.6%	0.999	5.2%

Table 3. Summary of LibriSpeech and Voice Command evaluation datasets in the experiments of real-world data analysis.

Feature	Evaluation Dataset		
	LibriSpeech Test-Clean	LibriSpeech Test-Other	Voice Command
# of Utterances	2,620	2,939	17,783
# of Speakers	40	33	95
# of Male Speakers	20	16	41

- *LibriSpeech* [26]. A widely used voice dataset which consists of 960 hours transcribed training utterances. The evaluation dataset has the splits of *Test-Clean* from 40 speakers and *Test-Other* from 33 speakers.
- *Voice Command*. This is a de-identified dataset collected using mobile devices through crowd-sourcing from a data supplier for ASR. No personally identifiable information (PII) is contained in this dataset. The participants are instructed to say voice commands on the topics of calling friends, playing music, etc. It consists of 2,440 hours transcribed training utterances. The evaluation set contains around 18K utterances from 95 speakers.

Table 3 shows details of the two evaluation datasets on number of utterances and number of speakers.

The ASR system in this investigation is an RNN-T model with Emformer encoder [27], LSTM predictor, and a joiner, having approximately 80 million parameters in total. For each of *LibriSpeech* or *Voice Command* data, the ASR model is trained from scratch using the corresponding training utterances.

4.2. Evaluation Results

For the *LibriSpeech* data, we study the ASR fairness on *gender*, that is, we would like to test whether there exists statistical significance on the WER ratio between male speakers and female speakers.

The **baseline** approach, which is widely used in practice, computes the ratio of empirical WER from male speakers group over the empirical WER from female speakers group. The bootstrap method is applied to compute the 95% CI of the ratio. For **model-based** approach, we fit a mixed-effects Poisson regression based on (5) (6) (7) with *gender* as the fixed effect and speaker label as a random effect.

Result is shown in Table 4. We can see that the baseline method leads to statistically significance claims on both *Test-Clean* and *Test-Other* sets, and interestingly, their conclusions are actually opposite. Specifically, on *Test-Clean* split of evaluation dataset, the baseline method shows that male speakers group has significant lower WER compared to female speakers group, while on *Test-Other* split, male

Table 4. Real-world analysis result on LibriSpeech dataset.

LibriSpeech Dataset	Baseline		Model-Based	
	WER Ratio	Confidence Interval	WER Ratio	Confidence Interval
<i>Test-Clean</i>	0.86	(0.76, 0.97)	0.88	(0.67, 1.14)
<i>Test-Other</i>	1.34	(1.23, 1.46)	1.28	(0.93, 1.76)

Table 5. Real-world analysis result on LibriSpeech dataset with sentence embedding as explanatory variables.

LibriSpeech Dataset	Model-Based (Embed)	
	WER Ratio	Confidence Interval
<i>Test-Clean</i>	1.01	(0.76, 1.33)
<i>Test-Other</i>	1.19	(0.87, 1.64)

speakers group has significant higher WER compared to the group of female speakers. On the other hand, the model-based approach does not claim any significant results on both splits. This makes sense since numbers of speakers in both splits are quite small, which results in high variance estimation that does not lead to statistically significant results. Thus utterances from more speakers are needed to reduce the standard errors and draw a more sound conclusion.

To further trace the source of WER gap, Table 5 shows the result of mixed-effect Poisson regression with sentence embedding of the true reference text as extra explanatory variables. We use pre-trained fastText word embeddings [28] with 300 dimensions and take their average to obtain the representation at sentence level. From the result, after excluding the effect from grammatical, lexical, or semantic characteristics, the WER gap between the two groups become smaller. Although it is not statistically significant, acoustic characteristics appear to contribute to the WER disparity on *Test-Other*.

Table 6. Real-world analysis result on Voice Command dataset.

Voice Command Dataset	Baseline		Model-Based	
	WER Ratio	Confidence Interval	WER Ratio	Confidence Interval
<i>Test</i>	1.08	(0.99, 1.20)	1.15	(0.78, 1.72)

We also investigate ASR fairness on *gender* for *Voice Command* dataset. The baseline and model-based methods are the same with the ones applied for *LibriSpeech*. Result is shown in Table 6. The baseline method does not claim that the WER on male speakers group is statistically significantly higher than the WER of female speakers group, but it is very close. The model-based method clearly does not lead to significant result, due to the relatively small number of speakers in each group.

5. CONCLUSIONS

In this paper, we introduce mixed-effects Poisson regression to better measure and interpret any WER difference among subgroups of interest. The presented method is very flexible and can effectively address the open problems of how to control the confounding factors, how to handle unobserved heterogeneity across speakers, and how to trace the source of any WER gap among different subgroups.

6. REFERENCES

- [1] Rachael Tatman and Conner Kasten, “Effects of talker dialect, gender & race on accuracy of Bing speech and YouTube automatic captions,” in *INTERSPEECH*, 2017, pp. 934–938.
- [2] Rachael Tatman, “Gender and dialect bias in YouTube’s automatic captions,” in *Proceedings of the First ACL Workshop on Ethics in Natural Language Processing*, 2017, pp. 53–59.
- [3] Ninareh Mehrabi, Fred Morstatter, Nripsuta Saxena, Kristina Lerman, and Aram Galstyan, “A survey on bias and fairness in machine learning,” *ACM Computing Surveys*, vol. 54, no. 6, pp. 1–35, 2021.
- [4] Allison Koenecke, Andrew Nam, Emily Lake, Joe Nudell, Minnie Quartey, Zion Mengesha, Connor Toups, John R Rickford, Dan Jurafsky, and Sharad Goel, “Racial disparities in automated speech recognition,” *Proceedings of the National Academy of Sciences*, vol. 117, no. 14, pp. 7684–7689, 2020.
- [5] Joshua L Martin and Kevin Tang, “Understanding racial disparities in automatic speech recognition: the case of habitual “be”,” in *INTERSPEECH*, 2020, pp. 626–630.
- [6] Clare Garvie and Jonathan Frankle, “Facial-recognition software might have a racial bias problem,” *The Atlantic*, vol. 7, 2016.
- [7] Tian Xu, Jennifer White, Sinan Kalkan, and Hatice Gunes, “Investigating bias and fairness in facial expression recognition,” in *European Conference on Computer Vision*, 2020, pp. 506–523.
- [8] Su Lin Blodgett, Lisa Green, and Brendan O’Connor, “Demographic dialectal variation in social media: A case study of African-American English,” in *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing*, 2016, pp. 1119–1130.
- [9] Aylin Caliskan, Joanna J Bryson, and Arvind Narayanan, “Semantics derived automatically from language corpora contain human-like biases,” *Science*, vol. 356, no. 6334, pp. 183–186, 2017.
- [10] Ziad Obermeyer, Brian Powers, Christine Vogeli, and Sendhil Mullainathan, “Dissecting racial bias in an algorithm used to manage the health of populations,” *Science*, vol. 366, no. 6464, pp. 447–453, 2019.
- [11] Gonzalo Navarro, “A guided tour to approximate string matching,” *ACM Computing Surveys*, vol. 33, no. 1, pp. 31–88, 2001.
- [12] P McCullagh and John A Nelder, *Generalized linear models*, vol. 37, CRC Press, 1989.
- [13] Badi Baltagi, *Econometric analysis of panel data*, John Wiley & Sons, 2008.
- [14] Julian J Faraway, *Extending the linear model with R: generalized linear, mixed effects and nonparametric regression models*, CRC press, 2016.
- [15] A Colin Cameron and Pravin K Trivedi, *Regression analysis of count data*, vol. 53, Cambridge university press, 2013.
- [16] Gary King, *Unifying political methodology: The likelihood theory of statistical inference*, Cambridge University Press, 1989.
- [17] Richard Berk and John M MacDonald, “Overdispersion and Poisson regression,” *Journal of Quantitative Criminology*, vol. 24, no. 3, pp. 269–284, 2008.
- [18] Jay M Ver Hoef and Peter L Boveng, “Quasi-Poisson vs. Negative Binomial regression: how should we model overdispersed count data?,” *Ecology*, vol. 88, no. 11, pp. 2766–2772, 2007.
- [19] Tomas Mikolov, Kai Chen, Greg Corrado, and Jeffrey Dean, “Efficient estimation of word representations in vector space,” *arXiv preprint arXiv:1301.3781*, 2013.
- [20] Milton Abramowitz and Irene A Stegun, *Handbook of mathematical functions with formulas, graphs, and mathematical tables*, vol. 55, US Government printing office, 1964.
- [21] Bradley Efron and Robert J Tibshirani, *An introduction to the bootstrap*, CRC Press, 1994.
- [22] Bradley Efron, “Second thoughts on the bootstrap,” *Statistical Science*, vol. 18, no. 2, pp. 135–140, 2003.
- [23] Jerzy Neyman, “X—outline of a theory of statistical estimation based on the classical theory of probability,” *Philosophical Transactions of the Royal Society of London. Series A, Mathematical and Physical Sciences*, vol. 236, no. 767, pp. 333–380, 1937.
- [24] Alan Stuart and Maurice G Kendall, *The advanced theory of statistics*, Griffin, 1963.
- [25] David Roxbee Cox and David Victor Hinkley, *Theoretical statistics*, CRC Press, 1979.
- [26] Vassil Panayotov, Guoguo Chen, Daniel Povey, and Sanjeev Khudanpur, “Librispeech: an ASR corpus based on public domain audio books,” in *Proc. ICASSP*, 2015.
- [27] Yangyang Shi, Yongqiang Wang, Chunyang Wu, Ching-Feng Yeh, Julian Chan, Frank Zhang, Duc Le, and Mike Seltzer, “Emformer: Efficient memory transformer based acoustic model for low latency streaming speech recognition,” in *Proc. ICASSP*, 2021.
- [28] Piotr Bojanowski, Edouard Grave, Armand Joulin, and Tomas Mikolov, “Enriching word vectors with subword information,” *Transactions of the Association for Computational Linguistics*, vol. 5, pp. 135–146, 2017.