



Modeling Dependent Structure for Utterances in ASR Evaluation

Zhe Liu, Fuchun Peng

Meta AI, Menlo Park, CA, USA

zheliu@fb.com, fuchunpeng@fb.com

Abstract

The *bootstrap* resampling method has been popular for performing significance analysis on word error rate (WER) in automatic speech recognition (ASR) evaluation. To deal with dependent speech data, the *blockwise bootstrap* approach is also introduced. By dividing utterances into uncorrelated blocks, this approach resamples these blocks instead of original data. However, it is typically nontrivial to uncover the dependent structure among utterances and identify the blocks, which might lead to subjective conclusions in statistical testing. In this paper, we present *graphical lasso* based methods to explicitly model such dependency and estimate uncorrelated blocks of utterances in a rigorous way, after which blockwise bootstrap is applied on top of the inferred blocks. We show the resulting variance estimator of WER in ASR evaluation is statistically consistent under mild conditions. We also demonstrate the validity of proposed approach on the LibriSpeech dataset.

Index Terms: automatic speech recognition, word error rate, statistical hypothesis testing, bootstrap, graphical modeling

1. Introduction

Word error rate (WER) is a common metric for the performance of an automatic speech recognition (ASR) system. Derived from the Levenshtein distance [1], WER can be calculated as

$$W = \frac{\sum_{i=1}^n e_i}{\sum_{i=1}^n m_i}, \quad (1)$$

where m_i is the number of words in the i th reference, and e_i refers to the sum of insertion, deletion, and substitution errors computed from the dynamic string alignment of the recognized sequence with the reference sequence.

When it comes to comparing the transcription quality of two ASR systems on the same evaluation dataset, the absolute or relative WER difference of system B versus system A is widely used:

$$\Delta W_{abs} = W_B - W_A = \frac{\sum_{i=1}^n (e_i^B - e_i^A)}{\sum_{i=1}^n m_i}, \quad (2)$$

$$\Delta W_{rel} = \frac{W_B - W_A}{W_A} = \frac{\sum_{i=1}^n (e_i^B - e_i^A)}{\sum_{i=1}^n e_i^A}. \quad (3)$$

Oftentimes, it is crucial to understand the reliability of computed WER on some evaluation dataset as well as the comparison of WERs between two ASR systems. For instance, we would like to answer if a WER of 0.05 on only 500 evaluation utterances is trustworthy in measuring the overall ASR quality for the target domain; we also want to know if a 2% relative WER improvement comparing ASR system B with ASR system A is due to chance or real.

For these purposes, statistical testing on WER has been explored [2, 3, 4, 5, 6, 7, 8]. The work of [5] presents a *bootstrap* approach for significance analysis on ASR evaluation which makes no distributional approximations and is easy to use.

Specifically, given any sequence of independent and identically distributed random variables $\{U_i\}_{i=1}^n$, we would like to estimate the variance of some statistic $T(U_1, \dots, U_n)$. The bootstrap method [9, 10] resamples data from the empirical distribution of $\{U_i\}_{i=1}^n$ for B^{boot} times, and then re-calculates the statistic T on each of these “bootstrap” samples. The variance of T can be estimated from the sample variance of these computed statistics. In significance analysis for WER, U_i represents the information of the i th utterance, and T refers to the statistic of W as in (1).

However, the vanilla bootstrap fails in dealing with dependent data, which is prevalent in the speech world. More recently, authors in [8] propose the *blockwise bootstrap* approach that addresses this issue by dividing utterances into nonoverlapping and uncorrelated blocks, then resampling these blocks instead of the original data.

One key issue confronting the blockwise bootstrap estimator in [8] is how to determine the block structure over utterances in ASR evaluation such that speech recognition errors in different blocks are uncorrelated. In general, speech recognition errors are considered as dependent if the corresponding utterances share similar acoustic or semantic characteristics. However, it is nontrivial to properly define such dependency in a rigorous way, which could lead to subjective conclusions and thus makes this approach less attractive in practical applications. A more advanced data-driven method is called for to quantify the dependent structure of utterances in ASR evaluation.

In this work, we propose to leverage *graphical modeling* [11] to estimate the dependent structure among speech utterances in ASR evaluation. Undirected graphical models have become a powerful tool of representing complex interactions among high-dimensional random variables. In our applications, each utterance is represented as an embedding vector with its acoustic and semantic information being encoded, and is considered as a random variable in the graph network. The inferred uncorrelated blocks of utterances, will be then used in blockwise bootstrap for performing statistical testing on WER.

Under the assumption of multivariate Gaussian among utterance representation random variables, the problem of dependent structure estimation is reduced to learning a sparse precision matrix and identifying its zero-pattern. The *graphical lasso* [12, 13], a sparse penalized maximum likelihood estimator for the precision matrix, has been studied in recent years. We introduce this method and its variant on relaxing the Gaussian assumption [14] to significance analysis on ASR evaluations.

Conventionally, clustering based methods can be used on

utterance embeddings and partition them into groups. However, there is typically no theories to establish the uncorrelatedness between utterances from different clusters. Compared with these methods, the use of graphical modeling provides more statistical guarantees in the dependent structure inference for utterances. Moreover, our approach is able to estimate the entire conditional independence graph and corresponding correlation matrix, which is more flexible to use.

We mainly pursue three goals in this paper: (1) to present new estimation approaches for utterances' dependent structure with the use of graphical lasso and its variant, and illustrate their applications in significance analysis on ASR quality evaluations; (2) to provide theories on the consistency of the resulting variance estimator; (3) to demonstrate empirical results on LibriSpeech speech dataset using our proposed methods. To the best of our knowledge, our paper is the first to introduce graphical modeling based approaches in estimating underlying dependent structure among evaluation utterances and explicitly addressing dependent speech data in ASR evaluations. Besides the applications to WER significance analysis, the proposed methods could also be useful in error analysis for ASR and correlation-aware speech model training, among others.

2. Methods

Suppose $\{U_i\}_{i=1}^n$ is an evaluation dataset for ASR, consisting of n utterances with the i th utterance denoted by U_i . Two ASR systems A and B are evaluated on $\{U_i\}_{i=1}^n$. Assume we have their evaluation results as follows:

$$\{(m_i, e_i^A, e_i^B)\}_{i=1}^n, \quad (4)$$

where m_i is the number of reference words of the i th utterance, e_i^A and e_i^B represent the numbers of word errors in ASR systems A and B , respectively. The statistics of our interest are the WER of ASR system A denoted as W_A , the absolute WER difference ΔW_{abs} and relative WER difference ΔW_{rel} comparing ASR system B versus ASR system A . In particular, we aim to compute the 95% confidence intervals for these statistics.

2.1. Graph Structure Estimation

We treat U_i as a random variable representing the i th utterance in the evaluation set. Denote u_i as the L -dimensional embedding vector for U_i with $u_i = (u_{i1}, \dots, u_{iL})^T$. Let $\bar{u}_i = \frac{1}{L} \sum_{l=1}^L u_{il}$ be the sample mean. In practice, the representation u_i for the utterance U_i can be obtained from the speaker embedding vector, or sentence embedding vector of the reference, or their concatenation.

Suppose the random vector $(U_1, \dots, U_n)^T$ for utterances has the density P^* . Then an undirected graph G for P^* has n vertices, collected in a set V , one for each variable. We represent the edges as a set E of unordered pairs: $(i, j) \in E$ if and only if there is an edge between U_i and U_j . In a conditional independence graph, an edge between U_i and U_j is absent if U_i and U_j are independent (denoted as $U_i \perp\!\!\!\perp U_j$), given the other variables. With the observations of $\{u_i\}_{i=1}^n$, the goal is to estimate the graph structure G .

Under the assumption that $(U_1, \dots, U_n)^T \sim \mathcal{N}_n(\mu, \Sigma)$, it has been shown that there is no edge between U_i and U_j if and only if $\Theta_{ij} = 0$, where $\Theta = \Sigma^{-1}$ is the precision matrix. Therefore, the graph estimation problem in Gaussian graphical models are equivalent to the estimation of a sparse Θ .

The graphical lasso estimator [12, 13] encourages some of the entries of the estimated precision matrix to be zero. Specif-

ically, it maximizes the ℓ_1 -regularized log-likelihood under the constraint that Θ is positive definite (denoted as $\Theta \succ 0$):

$$\hat{\Theta}_\lambda = \operatorname{argmax}_{\Theta \succ 0} \left\{ \log \det(\Theta) - \operatorname{tr}(\hat{\Sigma} \Theta) - \lambda \sum_{i \neq j} |\Theta_{ij}| \right\}, \quad (5)$$

where $\log \det(\cdot)$ denotes the log-determinant of a matrix, $\operatorname{tr}(\cdot)$ represents the trace of a matrix, λ is a regularization parameter, and $\hat{\Sigma}$ is the empirical covariance matrix:

$$\hat{\Sigma}_{ij} = \frac{1}{L-1} \sum_{l=1}^L (u_{il} - \bar{u}_i)(u_{jl} - \bar{u}_j). \quad (6)$$

The optimization problem in (5) can be solved using the coordinate descent algorithm [15, 16]. The regularization parameter λ can be chosen via cross-validation in practical applications.

While the graphical lasso is useful in many scenarios, a reliance on exact normality can be limiting. The *nonparanormal* [14], a form of Gaussian copula, weakens the Gaussian assumption by imposing normality on the transformed random vector $(f_1(U_1), \dots, f_n(U_n))^T$, with each $f_i(\cdot)$ being a function that meets certain conditions. This allows arbitrary single variable marginal probability distribution in the model.

Specifically, we say $(U_1, \dots, U_n)^T$ follows a nonparanormal distribution if there exist monotonic and differentiable functions $\{f_i(\cdot)\}_{i=1}^n$ such that $(f_1(U_1), \dots, f_n(U_n))^T \sim \mathcal{N}_n(\mu, \Sigma)$. Under this assumption, it was shown in [14] that there is no edge between U_i and U_j if and only if $\Theta_{ij} = 0$. The nonparanormal estimator uses a two-step procedure: (1) for each variable, replace the observations with their normal scores, subject to a Winsorized truncation, and (2) apply the graphical lasso to the transformed data to estimate Θ .

Once we obtain an estimated sparse precision matrix $\hat{\Theta}$ via the graphical lasso or nonparanormal estimator, we can determine the corresponding graph structure by identifying the zero-pattern of $\hat{\Theta}$. Independent blocks of random variables (i.e. utterances) can thus be formed by distinct connected components of the estimated graph. To be more specific, let $\{\hat{S}_k\}_{k=1}^K$ be the set of indices for K connected components in the estimated graph, then under the Gaussian or nonparanormal assumption, we have $U_i \perp\!\!\!\perp U_j$ for any $i \in \hat{S}_{k'}$ and $j \in \hat{S}_{k''}$ where $i \neq j$ and $k' \neq k''$. Here, it is worth noting that uncorrelated and joint normality implies independent random variables. Although in this paper we only utilize the inferred blocks in significance analysis for WER, the estimated covariance matrix $\hat{\Theta}^{-1}$ from the method could be useful in other practical applications.

2.2. Bootstrap with Inferred Blocks

Once the set of independent blocks $\{\hat{S}_k\}_{k=1}^K$ for utterances is estimated, we can apply the blockwise bootstrap [8] to compute the 95% confidence interval for any WER related statistic.

For $b = 1, \dots, B^{boot}$ where B^{boot} is a large number, we randomly sample with replacement K elements $\hat{S}_1^{(b)}, \dots, \hat{S}_K^{(b)}$ from the set of $\{\hat{S}_k\}_{k=1}^K$ to generate a bootstrap sample

$$\{(m_i, e_i^A, e_i^B)\}_{i \in \hat{S}_{k'}^{(b)}, k' = 1, \dots, K}, \quad (7)$$

where each $\hat{S}_{k'}^{(b)} \in \{\hat{S}_k\}_{k=1}^K$. Then for this bootstrap sample, the relative WER difference of system B versus system A can be computed as

$$\Delta W_{rel}^{(b)} = \frac{\sum_{k'=1}^K \sum_{i \in \hat{S}_{k'}^{(b)}} (e_i^B - e_i^A)}{\sum_{k'=1}^K \sum_{i \in \hat{S}_{k'}^{(b)}} e_i^A}. \quad (8)$$

Once we have all $\{\Delta W_{rel}^{(b)}\}_{b=1}^{B^{boot}}$ computed, the 95% confidence interval for the statistic ΔW_{rel} can be calculated by their empirical percentiles at 2.5% and 97.5%:

$$[\text{percentile}(\Delta W_{rel}^b, 2.5\%), \text{percentile}(\Delta W_{rel}^b, 97.5\%)]. \quad (9)$$

Here, a 95% confidence interval means that if we were able to have 100 different datasets from the same distribution of original dataset, and compute a 95% confidence interval based on each of these 100 datasets, then around 95 out of these 100 confidence intervals will contain the true value of the statistic of interest [17, 18, 19].

The confidence intervals for W_A or ΔW_{abs} can be calculated in a similar manner with above.

3. Theoretical Properties

We present theories to show the blockwise variance estimator with inferred blocks from the graphical lasso is statistically consistent for WER estimation under some mild conditions. Specifically, as the number of evaluation utterances increases indefinitely, the resulting sequence of variance estimates converges to the truth variance [20].

For simplicity, suppose all utterances in the evaluation set have the same number of words in the reference, that is, $m_i = m$ for all $i = 1, \dots, n$. Denote $Z_i = e_i/m$. Then the statistic of interest can be written as

$$W_n = \frac{1}{n} \sum_{i=1}^n Z_i. \quad (10)$$

Further, assume the n utterances can be divided into K_n^* nonoverlapping and independent blocks, denoted as $\{S_k\}_{k=1}^{K_n^*}$, with each block having the same number of utterances d_n^* .

Suppose $\{\widehat{S}_k\}_{k=1}^{K_n}$ is the set of estimated independent blocks from the graphical lasso. Following Theorem 1 in [21], under some conditions on the precision matrix of Gaussian distribution and the scaling of triple (L, D, n) , the set of $\{\widehat{S}_k\}_{k=1}^{K_n}$ approaches the true set of $\{S_k\}_{k=1}^{K_n^*}$ with a high probability when the embedding dimension L is large enough. Here, D represents the maximum of node degrees in graph G . Specifically for some constant $c > 0$

$$\mathbb{P}(\{\widehat{S}_k\}_{k=1}^{K_n} = \{S_k\}_{k=1}^{K_n^*}) \geq 1 - \exp(-c \log n) \rightarrow 1. \quad (11)$$

We further let

$$W_{\widehat{S}_k, n} = \frac{1}{|\widehat{S}_k|} \sum_{i \in \widehat{S}_k} Z_i \quad (12)$$

be the statistic of interest from the k th inferred block, where $|\widehat{S}_k|$ refers to the cardinality of the set of \widehat{S}_k . Consider the corresponding blockwise variance estimator

$$\hat{\sigma}_n^2 = \frac{d_n}{K_n} \sum_{k=1}^{K_n} (W_{\widehat{S}_k, n} - W_n)^2, \quad (13)$$

where $d_n = \lfloor n/K_n \rfloor$ denotes the average number of utterances in the inferred blocks. The following establishes its L_2 -consistency.

Theorem 1 Assume the asymptotic variance of W_n is

$$\lim_{n \rightarrow \infty} n \mathbf{E}(W_n - \mathbf{E}(W_n))^2 = \sigma^2 \in (0, \infty) \quad (14)$$

and $\mu = \mathbf{E}(Z_i)$ for any $i = 1, \dots, n$. Let d_n^* be s.t. $d_n^* \rightarrow \infty$ and $K_n^* \rightarrow \infty$ as $n \rightarrow \infty$. If $n^2 \mathbf{E}(W_n - \mu)^4$ is uniformly bounded and the conditions of Theorem 1 in [21] are satisfied, then

$$\hat{\sigma}_n^2 \rightarrow_{L_2} \sigma^2 \text{ as } n \rightarrow \infty. \quad (15)$$

The proof is straightforward by connecting the graph structure recovery property (11) with Theorem 1 in [8]. Specifically, the integrals of expectation and variance of $\hat{\sigma}_n^2$ can be decomposed into the correct graph structure recovery part (with a high probability) and the incorrect graph structure recovery part (with a low probability). The former follows the proof of Theorem 1 in [8]. Notice that each $(W_{\widehat{S}_k, n} - W_n)^2$ is bounded at the worse case, then the latter part can be dominated by $\exp(-c \log n)$ and become negligible as n goes to infinity.

This theorem shows statistical guarantees in the dependent structure inference for utterances and provides the convergence property of the corresponding blockwise estimator.

4. Experiments

4.1. Dataset and Setup

We experiment with the LibriSpeech data [22], consisting of 960 hours transcribed training utterances from 2,338 speakers. The evaluation dataset has the splits of `test-clean` and `test-other`. Table 1 shows detail of LibriSpeech evaluation set on number of utterances and number of speakers.

Table 1: Summary of LibriSpeech evaluation dataset.

	test-clean	test-other
Number of utterances	2,620	2,939
Number of speakers	40	33

We consider two ASR systems A and B in this investigation. The ASR system A is a RNN-T model with the Emformer encoder [23], LSTM predictor, and a joiner. The ASR system B utilizes the fast-slow cascaded encoders [24]. Both models have around 80 million parameters in total and are trained from scratch using the training utterances of LibriSpeech.

Table 2 displays the empirical WERs of both ASR systems as well as the absolute and relative WER difference of ASR system B compared with system A . The reported WERs are multiplied by 100.0 as per convention. We observe that ASR system B achieves large improvement over system A .

Table 2: WER results of ASR systems A and B .

	test-clean	test-other
WER of ASR System A	3.77	10.10
WER of ASR System B	3.25	8.13
Absolute WER difference	-0.52%	-1.97%
Relative WER difference	-13.8%	-19.5%

With this setup, the goal is to calculate the 95% confidence intervals for WER of ASR system A , relative and absolute WER difference between ASR system B and system A . We compare various approaches in the experiments including vanilla bootstrap (Bootstrap) [5], blockwise bootstrap (BlockBootstrap) [8], and the proposed method (InferredBlockBootstrap).

We set B^{boot} to 10,000 for each bootstrap based method. For the **BlockBootstrap** method, we treat the utterances from the same speaker as a block. To apply the **InferredBlockBootstrap** approach, the pre-trained BERT (base model, uncased) [25, 26] is leveraged to obtain the sentence embedding of each reference in evaluation utterances. The representation has 768 dimensions. Then among the utterances of each speaker, we adopt the graphical lasso on their sentence embedding vectors to estimate the underlying graph structure as well as corresponding independent blocks. Cross-validation is utilized to select the regularization parameter λ . Then the blockwise bootstrap method is applied on the inferred blocks. Here, the assumption we made is that ASR errors on any utterances from different speakers are independent with each other, while any utterances from the same speaker but having very different semantic or linguistic characteristics are also considered as independent.

4.2. Results

For each method in the comparison, Table 3 shows the calculated 95% confidence intervals for WER of ASR system A. For both of `test-clean` and `test-other` sets, we can see that the 95% confidence intervals from **InferredBlockBootstrap** are wider than the ones generated from **Bootstrap**, but are narrower than the ones computed by **BlockBootstrap**.

Table 3: Confidence intervals for WER of ASR system A.

	test-clean	test-other
Bootstrap	[3.54, 4.00]	[9.66, 10.53]
BlockBootstrap	[3.33, 4.23]	[8.40, 12.20]
InferredBlockBootstrap	[3.41, 4.14]	[8.79, 11.65]

Table 4 and Table 5 also display the computed 95% confidence intervals for the absolute and relative WER difference between ASR systems B and A , respectively. The findings are similar to the ones that we observe above. On the `test-other` dataset in Table 5, the width of the 95% confidence interval from **Bootstrap** is 4.8%, while the width of the one from **BlockBootstrap** is 7.9%. Also the 95% confidence interval from **InferredBlockBootstrap** has the width of 6.7%, which is around 40% wider than the one from **Bootstrap** but 15% narrower than the one from **BlockBootstrap**.

We observe that all of these 95% confidence intervals computed in Table 4 and Table 5 do not cover the point of 0, which means that the ASR system B is statistically significantly better than the system A in term of ASR quality.

Seen from these results, the vanilla **Bootstrap** method ignores the dependency among utterances and thus leads to narrower confidence intervals, which might result in over-optimistic conclusions due to false-positive discoveries. At another extreme, due to the subjective definitions of block structures, the **BlockBootstrap** method could be conservative on claiming statistical significance, where the corresponding confidence intervals are wider in the experiments. Instead, the presented **InferredBlockBootstrap** approach utilizes data-driven methods to explicitly model the dependent structure among utterances and draw sound conclusions.

Figure 1 shows the number of utterances and number of inferred blocks from the **InferredBlockBootstrap** method for each speaker. In total, there exist 822 blocks for the `test-clean` set and 912 blocks for the `test-other` set. The median of the ratios between the block count and utterance count per speaker is around 0.30.

Table 4: Confidence intervals for the absolute WER difference of ASR system B compared with ASR system A .

	test-clean	test-other
Bootstrap	[-0.67%, -0.35%]	[-2.24%, -1.69%]
BlockBootstrap	[-0.71%, -0.32%]	[-2.63%, -1.44%]
InferredBlockBootstrap	[-0.70%, -0.34%]	[-2.50%, -1.53%]

Table 5: Confidence intervals for the relative WER difference of ASR system B compared with ASR system A .

	test-clean	test-other
Bootstrap	[-17.4%, -9.6%]	[-21.9%, -17.1%]
BlockBootstrap	[-18.0%, -8.9%]	[-23.5%, -15.6%]
InferredBlockBootstrap	[-18.1%, -9.2%]	[-22.8%, -16.1%]

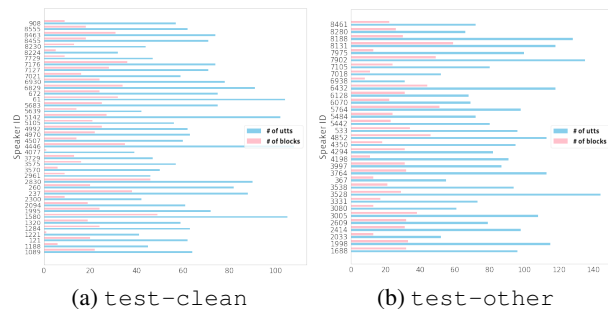


Figure 1: Number of utterances and number of inferred blocks for each speaker.

For the introduced **InferredBlockBootstrap** approach, we also experiment with the use of nonparanormal, which weakens the Gaussian assumption imposed on the embedding vectors for evaluation utterances. Table 6 shows the computed 95% confidence intervals on WER of ASR system A for graphical lasso and nonparanormal based **InferredBlockBootstrap**. From the results, we observe that nonparanormal based method results in wider confidence intervals compared with the ones obtained from graphical lasso, but they are still narrower than those computed by **BlockBootstrap**.

Table 6: Comparison of graphical lasso and nonparanormal based **InferredBlockBootstrap** on the computation of confidence intervals for WER of ASR system A .

	test-clean	test-other
InferredBlockBootstrap		
· graphical lasso	[3.41, 4.14]	[8.79, 11.65]
· nonparanormal	[3.37, 4.18]	[8.47, 12.07]

5. Conclusions

In this paper we propose to use the graphical lasso and its variant of nonparanormal to estimate the dependent structure among speech utterances in ASR evaluation. We demonstrate its applications on WER significance analysis through blockwise bootstrap resampling with inferred independent blocks. Results are also presented on LibriSpeech for the comparison of 95% confidence intervals computed from various approaches.

6. References

- [1] G. Navarro, "A guided tour to approximate string matching," *ACM Computing Surveys (CSUR)*, vol. 33, no. 1, pp. 31–88, 2001.
- [2] L. Gillick and S. J. Cox, "Some statistical issues in the comparison of speech recognition algorithms," in *IEEE International Conference on Acoustics, Speech, and Signal Processing (ICASSP)*, 1989, pp. 532–535.
- [3] D. S. Pallett, J. G. Fiscus, W. M. Fisher, and J. S. Garofolo, "Benchmark tests for the DARPA spoken language program," in *Proceedings of the Workshop on Human Language Technology*, 1993, pp. 7–18.
- [4] H. Strik, C. Cucchiari, and J. M. Kessens, "Comparing the recognition performance of CSRs: in search of an adequate metric and statistical significance test," in *Proceeding of International Conference on Spoken Language Processing*, 2000, pp. 740–743.
- [5] M. Bisani and H. Ney, "Bootstrap estimates for confidence intervals in ASR performance evaluation," in *IEEE International Conference on Acoustics, Speech, and Signal Processing (ICASSP)*, vol. 1, 2004, pp. 1–409.
- [6] D. Vilar, G. Leusch, H. Ney, and R. E. Banchs, "Human evaluation of machine translation through binary system comparisons," in *Proceedings of the Second Workshop on Statistical Machine Translation*. Association for Computational Linguistics, 2007, pp. 96–103.
- [7] J. M. Vilar, "Efficient computation of confidence intervals for word error rates," in *IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2008, pp. 5101–5104.
- [8] Z. Liu and F. Peng, "Statistical testing on ASR performance via blockwise bootstrap," in *Proceedings of Interspeech*, 2020, pp. 596–600.
- [9] B. Efron and R. J. Tibshirani, *An introduction to the bootstrap*. CRC Press, 1994.
- [10] B. Efron, "Second thoughts on the bootstrap," *Statistical Science*, vol. 18, no. 2, pp. 135–140, 2003.
- [11] S. L. Lauritzen, *Graphical models*. Clarendon Press, 1996, vol. 17.
- [12] M. Yuan and Y. Lin, "Model selection and estimation in the Gaussian graphical model," *Biometrika*, vol. 94, no. 1, pp. 19–35, 2007.
- [13] J. Friedman, T. Hastie, and R. Tibshirani, "Sparse inverse covariance estimation with the graphical lasso," *Biostatistics*, vol. 9, no. 3, pp. 432–441, 2008.
- [14] H. Liu, J. Lafferty, and L. Wasserman, "The nonparanormal: semiparametric estimation of high dimensional undirected graphs," *Journal of Machine Learning Research*, vol. 10, no. 10, 2009.
- [15] J. Friedman, T. Hastie, H. Höfling, and R. Tibshirani, "Pathwise coordinate optimization," *The Annals of Applied Statistics*, vol. 1, no. 2, pp. 302–332, 2007.
- [16] T. T. Wu and K. Lange, "Coordinate descent algorithms for lasso penalized regression," *The Annals of Applied Statistics*, vol. 2, no. 1, pp. 224–244, 2008.
- [17] J. 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.
- [18] A. Stuart and M. G. Kendall, *The advanced theory of statistics*. Griffin, 1963.
- [19] D. R. Cox and D. V. Hinkley, *Theoretical statistics*. CRC Press, 1979.
- [20] T. Amemiya, *Advanced econometrics*. Harvard university press, 1985.
- [21] G. Raskutti, B. Yu, M. J. Wainwright, and P. Ravikumar, "Model selection in gaussian graphical models: High-dimensional consistency of ℓ_1 -regularized MLE," *Advances in Neural Information Processing Systems*, vol. 21, 2008.
- [22] V. Panayotov, G. Chen, D. Povey, and S. Khudanpur, "Librispeech: an ASR corpus based on public domain audio books," in *IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2015, pp. 5206–5210.
- [23] Y. Shi, Y. Wang, C. Wu, C.-F. Yeh, J. Chan, F. Zhang, D. Le, and M. Seltzer, "Emformer: Efficient memory transformer based acoustic model for low latency streaming speech recognition," in *IEEE International Conference on Acoustics, Speech, and Signal Processing (ICASSP)*, 2021.
- [24] J. Mahadeokar, Y. Shi, K. Li, D. Le, J. Zhu, V. Chandra, O. Kalinli, and M. L. Seltzer, "Streaming parallel transducer beam search with fast-slow cascaded encoders," in *Interspeech*, 2022.
- [25] J. Devlin, M.-W. Chang, K. Lee, and K. Toutanova, "BERT: pre-training of deep bidirectional transformers for language understanding," *arXiv preprint arXiv:1810.04805*, 2018.
- [26] T. Wolf, L. Debut, V. Sanh, J. Chaumond, C. Delangue, A. Moi, P. Cistac, T. Rault, R. Louf, M. Funtowicz *et al.*, "Transformers: State-of-the-art natural language processing," in *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing: System Demonstrations*, 2020, pp. 38–45.