



Training Keyword Spotting Models on Non-IID Data with Federated Learning

Andrew Hard, Kurt Partridge, Cameron Nguyen, Niranjan Subrahmanya, Aishanee Shah, Pai Zhu, Ignacio Lopez Moreno, Rajiv Mathews

Google LLC,
Mountain View, CA, U.S.A.

harda@google.com

Abstract

We demonstrate that a production-quality keyword-spotting model can be trained on-device using federated learning and achieve comparable false accept and false reject rates to a centrally-trained model. To overcome the algorithmic constraints associated with fitting on-device data (which are inherently non-independent and identically distributed), we conduct thorough empirical studies of optimization algorithms and hyperparameter configurations using large-scale federated simulations. To overcome resource constraints, we replace memory-intensive MTR data augmentation with SpecAugment, which reduces the false reject rate by 56%. Finally, to label examples (given the zero visibility into on-device data), we explore teacher-student training.

Index Terms: federated learning, on-device learning, keyword spotting, wake word detection, non-iid data, data augmentation

1. Introduction

Keyword spotting has become an essential access point for virtual assistants. Vocalized keywords such as *Alexa*, *Hey Google*, or *Hey Siri* can be used to initiate search queries and issue commands to mobile phones and smart speakers. The underlying algorithms must process streaming audio—the majority of which must be ignored—and trigger quickly and reliably when needed.

Neural networks have achieved state-of-the-art performance in automatic speech recognition tasks [1, 2, 3]. Applications of neural networks to keyword spotting have also been explored, particularly within the contexts of quality improvement and latency reduction for low-resource environments [4, 5, 6, 7, 8, 9] and end-to-end model training [10].

Federated Learning (FL) [11] is a decentralized computation paradigm that can be used to train neural networks directly on-device. In FL, all model updates shared by devices with the server are *ephemeral* (only stored temporarily in server memory), *focused* (only relevant to a specific training task), and *aggregated* (only processed collectively with updates from other devices across the population). In conjunction with techniques such as differential privacy [12, 13] and secure aggregation [14], FL can integrate strong anonymity and privacy guarantees into the neural network training process.

Federated learning provides a path to train keyword models at the edge, on real user data, as opposed to proxy data. In contrast, centrally-trained keyword-spotting models use proxy data, since false accepts (in which the keyword-spotter accidentally triggers) are not logged.

Multiple production models have been trained with federated learning, including next-word prediction [15], emoji prediction [16], n-gram language models [17], and query suggestions [18] for mobile keyboards. Many of these models achieve

better performance as a result of the additional signals and unbiased data available on-device. Recently, the feasibility of training keyword-spotting algorithms with FL has been explored with smaller datasets [19].

On-device training comes with challenges, including the fact that the quantities and characteristics of training examples vary considerably from device to device. Centrally-trained models benefit from the ability to sample data in a controlled, independent and identically distributed (IID) manner, resulting in gradient updates that are unbiased estimates of the total gradient for the dataset. This is not true on-device, where client updates are biased representations of the gradient across the entire population [20].

Non-IID data adversely affect convergence [21], and have been identified as a fundamental challenge to FL [22]. Proposals to fit non-IID data better include optimizers that account for client drift [23], data sharing between client devices [20], and adaptive server optimizers with client learning rate decay [24], among others [25].

The primary contribution of this paper is to demonstrate that keyword-spotting models can be trained on large-scale datasets using FL, and can achieve false accept and false reject rates that rival those of centralized training. Using simulated federated learning experiments on large-scale datasets consisting of thousands of speakers and millions of utterances, we address the algorithmic challenges associated with training on non-IID data, the visibility challenges associated with labeling on-device data, and the physical constraints that limit augmentation capabilities on-device.

2. Model

This paper used the end-to-end architecture described in [8]. End-to-end trainable neural architectures have demonstrated state-of-the-art performance in terms of accuracy as well as lowered resource requirements while providing a highly optimizable system design [10]. The model consisted of an encoder-decoder architecture in which both the encoder and decoder made use of efficiently parameterized SVDF (single value decomposition filter) layers—originally introduced in [26]—to approximate fully-connected layers with low rank approximations. Dense bottleneck layers were used to further reduce computational costs and keep the model size down at only 320,778 parameters (Figure 1).

The encoder took spectral domain features \mathbf{X}_t as input and generated outputs Y^E corresponding to phoneme-like sound units. The decoder model used the encoder output as input and generated binary output Y^D that predicted the existence of a keyword. The model was fed with acoustic input features at each frame (generated every 10ms), and generated prediction labels at each frame in a streaming manner.




Figure 1: *End-to-end topology trained to predict the keyword likelihood score [8]. The encoder consists of 4 SVDF-plus-bottleneck layers, while 3 SVDF layers comprise the decoder.*

Training such an architecture traditionally required frame-level labels generated by LVCSR systems [27] to provide accurate timing information. This approach to label generation is not possible on-device, due to the large computational resources required to store and run a LVCSR system. Therefore, while we retained the same architecture, in our experiments we train the system with just a binary cross entropy loss for keyword presence and did not present any supervised targets to train the encoder. Recent work [8] suggests a better approach to train the encoder without LVCSR targets and will be the focus of future work.

3. Federated Optimization

In federated learning [11], a central server sends neural models to many client devices (such as phones). These clients process local caches of data in parallel and send updated model weights back to the server. The server aggregates the updates, produces a new global model, and repeats this cycle (called a federated training round) until the model converges.

Federated Averaging (FedAvg) [11] was used as a baseline optimization algorithm. During each training round, indexed by t , a subset of $K = 400$ client devices in the experiment population downloaded a global model, w_t , from the server. Each client $k \in K$ had a local data cache consisting of n_k examples. The clients used stochastic gradient descent (SGD) to train over their local examples and derive an average gradient, g_k . For a client learning rate η_c , the local client step, w_{t+1}^k , was defined:

$$w_{t+1}^k = w_t - \eta_c g_k. \quad (1)$$

This equation represents a single step of SGD, but client training typically involved multiple steps of SGD with a batch size of 1. Updated client weights were sent back to the server, which aggregated them to compute a global model update:

$$\Delta_t = \sum_{k=1}^K \frac{n_k}{N} (w_t - w_{t+1}^k). \quad (2)$$

where $N = \sum_k n_k$. For a server learning rate η_s , the updated global model weights, w_{t+1} were computed according to:

$$w_{t+1} = w_t - \eta_s \Delta_t. \quad (3)$$

When phrased in the form of Equation 3, FedAvg clearly consists of an *inner optimizer loop* (SGD over gradients on the clients) and an *outer optimizer loop* (SGD on averaged weight deltas on the server).

Momentum-based variants of FedAvg were explored as in Ref. [15, 28], in which Nesterov accelerated gradients (NAG) [29] were applied to the server updates. The server update in Equation 3 was replaced by:

$$w_{t+1} = w_t - \eta_s (\gamma v_{t+1} + \Delta_t). \quad (4)$$

where γ is the momentum hyperparameter and $v_{t+1} = \gamma v_t + \Delta_t$ is the forward-looking Nesterov accelerated gradient. The advantages of Nesterov momentum over classical momentum [30] have been demonstrated in the central training setting [31], and were expected to translate to FL.

Finally, adaptive variants of FedAvg were investigated, in which the server optimizer function was replaced by Adam [32], Yogi [33], or LAMB [34]. Prior works shown that adaptive per-coordinate updates can improve convergence for FL [19, 24]. Adaptive methods have shown particular strength in environments with heavy-tailed stochastic gradient noise distributions [35]—a common property of non-IID data in FL.

Adaptive methods replace the server optimizer loop, shown in Equation 3, with an adaptive optimization step. As with FedAvg, the classical momentum terms of Adam or Yogi can be replaced with NAG as in NAdam [36].

4. Data and Simulations

Experiments were conducted using simulated federated learning with vendor-collected datasets. A total of 1.5M utterances (1,650 hours in total) were recorded by 8,000 English-speaking participants. Each lasted a few seconds, and most contained one of two spoken keyword phrases. The dataset was divided into a train set (1.3 million utterances) and an eval set (180,000 utterances) with non-overlapping groups of users. IID and non-IID configurations of the training data were prepared, while the eval set was always used in an IID configuration.

Utterances were grouped into non-IID simulated clients in three steps. First, data were clustered according to speaker. The resulting clusters varied significantly in size: though speakers provided a median of 108 utterances each, a few speakers were associated with nearly a thousand unique examples. Next, clusters were further divided on the basis of labels. Individual clusters contained either positive utterances (which contained the keyword) or negative utterances (which lacked the keyword) exclusively. It should be noted that training labels were specified on a per-frame basis, and even positive utterances contained numerous frames with negative targets. Finally, the clusters were randomly subdivided to enforce an exponential distribution of utterances per client (n_k). The final step was motivated by a desire to match the training data distribution to the inference distribution, coupled with observations of power-law feature usage among the general population.

The resulting non-IID training dataset consisted of 160,000 clusters and a median of 6.5 utterances per cluster. Individual clusters were affiliated with individual simulated client devices for federated learning.

For IID simulated clients, the data were randomly divided into clusters consisting of 50 utterances. Each uniformly-sized cluster thus included a mix of speakers and labels. The IID training set contained 26,000 clusters of 50 utterances each. 3,675 clusters comprised the IID eval set.

5. Experiments

This section describes experiments to address the various constraints of on-device training. Specifically, we discuss optimization techniques to overcome the algorithmic challenge of fitting non-IID data, lightweight data augmentation techniques that run with constrained on-device resources, and teacher-student training to provide labels given the inability to peek at federated data.

5.1. Training and evaluation

Model checkpoints generated by the training tasks were periodically saved and sent to a held-out set of client devices for federated eval tasks. The train and eval tasks ran over orthogonal datasets, which were constructed following the description in Section 4. Metrics including the frame-level accuracy and cross-entropy loss were computed on each client and averaged on the server.

Hyperparameters were tuned and checkpoints were selected based on the criterion of minimizing eval loss for non-IID data. Eval loss was measured on the IID dataset described in Section 4. Tasks were also trained on IID datasets in order to compare training under the different data distributions.

Models were evaluated using false accept (FA) and false reject (FR) rates, which were computed offline on large tests sets consisting of negative utterances and positive utterances, respectively. The triggering thresholds were tuned to have FA=0.2%, approximately. FA and FR are more relevant to quality than loss, since they directly correspond to inference performance.

5.2. Optimizers and learning rate schedules

Optimization techniques were explored for non-IID training. First, the algorithms described in Section 3 were tuned via grid searches. For FedAvg, $\eta_s = 1.0$ and a momentum value of 0.99 was found to work best. FedAdam and FedYogi both converged well with $\beta_1 = 0.9$ and $\beta_2 = 0.999$, though Adam worked with the default $\epsilon = 10^{-8}$ and $\eta_s = 10^{-3}$ while Yogi worked best with a larger $\epsilon = 10^{-3}$, $\eta_s = 0.1$, and initial accumulator value of 10^{-6} . Experiments were also performed in which Nesterov accelerated gradients were substituted for classical momenta.

Table 1 compares non-IID training with each server optimization algorithm. Exponentially decayed client learning rates were used with FedAdam and FedYogi, while FedAvg worked better with constant client learning rates. The adaptive optimizers had a decisive advantage over FedAvg on FR. Replacing classical momentum with NAG benefitted FedAvg and FedAdam, but FedYogi with classical momentum had the lowest FR overall.

Table 1: Comparisons of offline false accept and reject rates for various optimizers on non-IID data.

Optimizer	FA [%]	FR [%]
FedAvg	0.21	8.76
FedAvg + NAG	0.21	4.09
FedAdam	0.19	1.95
FedAdam + NAG	0.21	1.68
FedYogi	0.19	1.39
FedYogi + NAG	0.20	2.11

Theoretical and empirical results indicate that client learning rate (LR) decay improves convergence on non-IID data [20, 24, 21]. Fixed client learning rates (with $\eta_c = 0.02$) were compared with exponentially-decayed client learning rate schedules, in which η_c was reduced by a constant factor, $\Gamma_{\eta,c}$, after every N_Γ steps. Hyperparameter scans found that the eval loss was minimized with an initial learning rate $\eta_{c,0} = 0.02$, $\Gamma_{\eta,c} = 0.9$, and $N_\Gamma = 1000$.

Learning rate comparisons are shown in Table 2, for the FedYogi optimizer. LR decay significantly improves both IID

and non-IID training. The difference is most pronounced for non-IID training, where the FR decreases from 2.35% to 1.39% given a fixed FA=0.2%.

Table 2: A comparison of client learning rate schedules.

LR schedule	FR (IID) [%]	FR (Non-IID) [%]
Constant	2.14	2.35
Exponential	1.73	1.39

5.3. Data Augmentation

Two common speech data augmentation methods—MTR [37] and SpecAugment [38]—were tuned for non-IID training.

MTR is an acoustic room simulator that generates noise files which can be applied to spectrogram inputs. Based on *a priori* distributions, MTR generates random room sizes and dimensions, speaker and noise source positions, signal to noise ratios, and reverberation times [39]. The technique is effective for far-field speech recognition and has been used previously for keyword spotting [10, 8].

In the simulation experiments, MTR was used to create up to 100 noised replica of each clean utterance from vendor data. In order to keep a constant number of training examples per simulated client, MTR configurations were randomly sampled every time a given simulated client device was used for training.

Unfortunately, MTR is infeasible for on-device training: users would have to download additional noise data (for additive noise) and room simulation configurations (for reverberations). The extra data processing would also lengthen client training times.

Spectrum augmentation (SpecAugment) is a fast and lightweight alternative for speech data augmentation. It has been used previously for keyword spotting [40], and has been used to achieve state-of-the-art ASR performance [41]. The augmentation policy broadly consists of three components: (1) *Time Masking*, in which consecutive time frames in the spectrogram are masked and replaced with Gaussian-distributed noise, (2) *Frequency Masking*, in which adjacent bins of the spectrogram are zeroed, and (3) *Time Warping*, in which features are linearly displaced along the temporal axis.

SpecAugment is an ideal on-device alternative to MTR, as it requires no config files and minimally increases training time. Tuning in non-IID data simulations found an optimal configuration of 2 time masks of up to 60 frames along with 2 frequency masks of up to 15 bins. TimeWarp was not used.

Table 3: FA and FR comparisons for models trained on IID and non-IID data with different data augmentations.

Data Augmentation	Data type	FA [%]	FR [%]
No augmentation	IID	0.17	4.20
No augmentation	Non-IID	0.20	3.19
MTR	IID	0.13	6.96
MTR	Non-IID	0.18	6.15
SpecAugment	IID	0.20	1.73
SpecAugment	Non-IID	0.19	1.39

Augmentation strategies for IID and non-IID FL are compared in Table 3. SpecAugment reduced the FR with respect to MTR and no augmentation on both data distributions. Thus, we

can reduce communication costs and on-device processing time with SpecAugment while also improving performance.

5.4. Labeling

High-quality labeling can be difficult to obtain on-device, since peeking at data is impossible by design in FL, and user feedback signals are unreliable or infrequent. Given the obstacles to on-device labeling, teacher student training can be used to adapt a model trained on the server (with manually labeled data) to the on-device unlabeled data domain [42, 43, 44]. Models were trained on both IID and non-IID data with supervised and teacher-generated labels. For the semi-supervised setting, the teacher model architecture was identical to the student, but was trained on additional data in a centralized setting.

Table 4: FR comparisons for on-device labeling strategies.

Labeling	FR (IID) [%]	FR (Non-IID) [%]
Supervised	1.73	1.39
Teacher	2.12	2.07

Table 4 compares teacher student training with supervised training. While the FR increases when moving to semi-supervised labels, it is expected that the matched data available in true on-device data, coupled with a limited number of samples labeled with user feedback signals, will close the performance gap.

5.5. Central training comparison and ablation studies

The previously-discussed techniques were applied and then removed individually in an ablation study, and the results were compared with a model trained in the centralized setting on the exact same vendor dataset using 4×10^8 steps of asynchronous SGD. This provided a direct comparison of centralized training and FL on IID and non-IID data.

Two additional techniques were explored to fit non-IID data. Client update clipping, based on $\|w_{t+1}^k - w_t\|^2$, was tuned for non-IID fitting. Multiple client training epochs were also studied, and have been shown to improve convergence [24].

The following settings were used for the ablation study FL baseline: data were augmented with SpecAugment, 10 client epochs were used, the client LR was decayed, client L2 weight norms were clipped to 20, and the FedYogi optimizer was used to train with supervised labels.

Ablation results are shown in Figure 2. The metrics favor non-IID data because the hyperparameters were tuned to minimize non-IID eval loss. Had the hyperparameters been tuned for IID data, the IID FR would be lower than the FR tuned for non-IID data.

FL achieves comparable FR performance with the centrally-trained model. In absolute terms, SpecAugment and multiple client epochs provided the largest contributions to both IID and non-IID performance. Interestingly, decayed client learning rates were more important to non-IID training than IID training. And contrary to prior studies [20], we found that additional client training epochs benefitted non-IID training.

6. Conclusions

Empirical studies were conducted to train a keyword-spotting model using FL on non-IID data. Adaptive server optimizers like FedYogi helped train a model with a lower false reject




Figure 2: FA and FR for the ablation study, with models trained on IID and non-IID data compared with centralized training (dashed line).

rate in fewer training rounds. We also demonstrated the necessity and utility of replacing MTR with SpecAugment for on-device training. Ablation studies revealed the importance of multiple client epochs and reduced client clipping. And we provided strong empirical evidence in favor of client learning rate decay for training with non-IID data. Finally, we overcome the visibility limitations of on-device training by demonstrating that, in the absence of high-quality on-device labels, teacher-student training can achieve comparable performance.

7. Acknowledgements

The authors would like to thank Google Research colleagues for providing the FL framework, Manzil Zaheer for his optimizer expertise, and Daniel Park for SpecAugment discussions.

8. References

- [1] G. Hinton, L. Deng, D. Yu, G. Dahl, A. rahman Mohamed, N. Jaitly, A. Senior, V. Vanhoucke, P. Nguyen, T. Sainath, and B. Kingsbury, “Deep neural networks for acoustic modeling in speech recognition,” *Signal Processing Magazine*, 2012.
- [2] N. Jaitly, P. Nguyen, A. Senior, and V. Vanhoucke, “Application of pretrained deep neural networks to large vocabulary speech recognition,” in *Proceedings of Interspeech 2012*.
- [3] Y. He, T. N. Sainath, R. Prabhavalkar, I. McGraw, R. Alvarez, D. Zhao, D. Rybach, A. Kannan, Y. Wu, R. Pang, and et al., “Streaming end-to-end speech recognition for mobile devices,” *ICASSP 2019 - 2019 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*.
- [4] G. Chen, C. Parada, and G. Heigold, “Small-footprint keyword spotting using deep neural networks,” in *ICASSP*, 2014.
- [5] S. Panchapagesan, M. Sun, A. Khare, S. Matsoukas, A. Mandal, B. Hoffmeister, and S. Vitaladevuni, “Multi-task learning and weighted cross-entropy for DNN-based keyword spotting,” in *INTERSPEECH*, 2016.
- [6] S. Team, “Hey siri: An on-device dnn-powered voice trigger for apple’s personal assistant,”

- <https://machinelearning.apple.com/2017/10/01/hey-siri.html>, accessed: 2020-04-30.
- [7] M. Sun, D. Snyder, Y. Gao, V. K. Nagaraja, M. Rodehorst, S. Panchapagesan, N. Strom, S. Matsoukas, and S. Vitaladevuni, "Compressed time delay neural network for small-footprint keyword spotting," in *INTERSPEECH*, 2017.
 - [8] H.-J. Park, P. Violette, and N. Subrahmanya, "Learning to detect keyword parts and whole by smoothed max pooling," *ICASSP 2020 - 2020 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*.
 - [9] A. Gruenstein, R. Alvarez, C. Thornton, and M. Ghodrat, "A cascade architecture for keyword spotting on mobile devices," 2017.
 - [10] R. Alvarez and H.-J. Park, "End-to-end streaming keyword spotting," *ICASSP 2019 - 2019 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*.
 - [11] B. McMahan, E. Moore, D. Ramage, S. Hampson, and B. A. y Arcas, "Communication-efficient learning of deep networks from decentralized data," in *Proceedings of the 20th International Conference on Artificial Intelligence and Statistics, AISTATS 2017, 20-22 April 2017, Fort Lauderdale, FL, USA*.
 - [12] C. Dwork, "Differential privacy," in *33rd International Colloquium on Automata, Languages and Programming, part II (ICALP 2006)*, vol. 4052. Venice, Italy: Springer Verlag.
 - [13] M. Abadi, A. Chu, I. Goodfellow, H. B. McMahan, I. Mironov, K. Talwar, and L. Zhang, "Deep learning with differential privacy," *Proceedings of the 2016 ACM SIGSAC Conference on Computer and Communications Security*.
 - [14] K. Bonawitz, V. Ivanov, B. Kreuter, A. Marcedone, H. B. McMahan, S. Patel, D. Ramage, A. Segal, and K. Seth, "Practical secure aggregation for federated learning on user-held data," in *NIPS Workshop on Private Multi-Party Machine Learning*, 2016.
 - [15] A. Hard, K. Rao, R. Mathews, S. Ramaswamy, F. Beaufays, S. Augenstein, H. Eichner, C. Kiddon, and D. Ramage, "Federated learning for mobile keyboard prediction," 2018.
 - [16] S. Ramaswamy, R. Mathews, K. Rao, and F. Beaufays, "Federated learning for emoji prediction in a mobile keyboard," 2019.
 - [17] M. Chen, A. T. Suresh, R. Mathews, A. Wong, C. Allauzen, F. Beaufays, and M. Riley, "Federated learning of n-gram language models," in *Proceedings of the 23rd Conference on Computational Natural Language Learning, CoNLL 2019, Hong Kong, China, November 3-4, 2019*.
 - [18] T. Yang, G. Andrew, H. Eichner, H. Sun, W. Li, N. Kong, D. Ramage, and F. Beaufays, "Applied federated learning: Improving google keyboard query suggestions," 2018.
 - [19] D. Leroy, A. Coucke, T. Lavril, T. Gisselbrecht, and J. Dureau, "Federated learning for keyword spotting," *ICASSP 2019 - 2019 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*.
 - [20] Y. Zhao, M. Li, L. Lai, N. Suda, D. Civin, and V. Chandra, "Federated learning with non-IID data," 2018.
 - [21] X. Li, K. Huang, W. Yang, S. Wang, and Z. Zhang, "On the convergence of fedavg on non-IID data," 2019.
 - [22] P. K. et al., "Advances and open problems in federated learning," 2019.
 - [23] S. P. Karimireddy, S. Kale, M. Mohri, S. J. Reddi, S. U. Stich, and A. T. Suresh, "Scaffold: Stochastic controlled averaging for federated learning," 2019.
 - [24] S. Reddi, Z. Charles, M. Zaheer, Z. Garrett, K. Rush, J. Konečný, S. Kumar, and H. B. McMahan, "Adaptive Federated Optimization," *arXiv e-prints*, Feb. 2020.
 - [25] M. Dundar, B. Krishnapuram, J. Bi, and R. B. Rao, "Learning classifiers when the training data is not iid," in *Proceedings of the 20th International Joint Conference on Artificial Intelligence*, ser. IJCAI'07. San Francisco, CA, USA: Morgan Kaufmann Publishers Inc., 2007.
 - [26] P. Nakkiran, R. Alvarez, R. Prabhavalkar, and C. Parada, "Compressing deep neural networks using a rank-constrained topology," in *Proceedings of Annual Conference of the International Speech Communication Association (Interspeech)*, 2015.
 - [27] T. Sainath and C. Parada, "Convolutional neural networks for small-footprint keyword spotting," in *Proceedings of Annual Conference of the International Speech Communication Association (Interspeech)*, 2015.
 - [28] T.-M. H. Hsu, H. Qi, and M. Brown, "Measuring the effects of non-identical data distribution for federated visual classification," 2019.
 - [29] Y. Nesterov, "A method for solving the convex programming problem with convergence rate $o(1/k^2)$," *Dokl. Akad. Nauk SSSR*, vol. 269, 1983.
 - [30] B. Polyak, "Some methods of speeding up the convergence of iteration methods," *USSR Computational Mathematics and Mathematical Physics*, vol. 4, no. 5, 1964.
 - [31] I. Sutskever, J. Martens, G. Dahl, and G. Hinton, "On the importance of initialization and momentum in deep learning," in *Proceedings of the 30th International Conference on Machine Learning*. PMLR, Jun 2013.
 - [32] D. P. Kingma and J. Ba, "Adam: A Method for Stochastic Optimization," *ArXiv e-prints*, Dec. 2014.
 - [33] M. Zaheer, S. J. Reddi, D. Sachan, S. Kale, and S. Kumar, "Adaptive methods for nonconvex optimization," in *Proceedings of the 32nd International Conference on Neural Information Processing Systems*, ser. NIPS'18. Red Hook, NY, USA: Curran Associates Inc., 2018.
 - [34] Y. You, J. Li, S. Reddi, J. Hseu, S. Kumar, S. Bhojanapalli, X. Song, J. Demmel, K. Keutzer, and C.-J. Hsieh, "Large batch optimization for deep learning: Training bert in 76 minutes," 2019.
 - [35] J. Zhang, S. P. Karimireddy, A. Veit, S. Kim, S. J. Reddi, S. Kumar, and S. Sra, "Why adam beats sgd for attention models," 2019.
 - [36] T. Dozat, "Incorporating nesterov momentum into adam," 2016.
 - [37] C. Kim, A. Misra, K. K. Chin, T. Hughes, A. Narayanan, T. N. Sainath, and M. Bacchiani, "Generation of large-scale simulated utterances in virtual rooms to train deep-neural networks for far-field speech recognition in google home," in *INTERSPEECH*, 2017.
 - [38] D. S. Park, W. Chan, Y. Zhang, C.-C. Chiu, B. Zoph, E. D. Cubuk, and Q. V. Le, "SpecAugment: A simple data augmentation method for automatic speech recognition," *Interspeech 2019*.
 - [39] R. Prabhavalkar, R. Alvarez, C. Parada, P. Nakkiran, and T. Sainath, "Automatic gain control and multi-style training for robust small-footprint keyword spotting with deep neural networks," in *Proceedings of International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2015.
 - [40] J. Hou, Y. Shi, M. Ostendorf, M. Hwang, and L. Xie, "Mining effective negative training samples for keyword spotting," in *ICASSP 2020 - 2020 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*.
 - [41] D. S. Park, Y. Zhang, C.-C. Chiu, Y. Chen, B. Li, W. Chan, Q. V. Le, and Y. Wu, "SpecAugment on large scale datasets," *ICASSP 2020 - 2020 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*.
 - [42] G. Hinton, O. Vinyals, and J. Dean, "Distilling the knowledge in a neural network," 2015.
 - [43] J. Li, M. L. Seltzer, X. Wang, R. Zhao, and Y. Gong, "Large-scale domain adaptation via teacher-student learning," *Interspeech 2017*.
 - [44] S. H. Krishnan Parthasarathi and N. Strom, "Lessons from building acoustic models with a million hours of speech," *ICASSP 2019 - 2019 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*.