



An Anchor-Free Detector for Continuous Speech Keyword Spotting

Zhiyuan Zhao¹, Chuanxin Tang¹, Chengdong Yao², Chong Luo¹

¹Microsoft Research Asia

²University of Technology Sydney

zhiyuzh@microsoft.com, chutan@microsoft.com, chengdong.yao-1@student.uts.edu.au,
cluoluo@microsoft.com

Abstract

Continuous Speech Keyword Spotting (CSKWS) is a task to detect predefined keywords in a continuous speech. In this paper, we regard CSKWS as a one-dimensional object detection task and propose a novel anchor-free detector, named AF-KWS, to solve the problem. AF-KWS directly regresses the center locations and lengths of the keywords through a single-stage deep neural network. In particular, AF-KWS is tailored for this speech task as we introduce an auxiliary unknown class to exclude other words from non-speech or silent background. We have built two benchmark datasets named LibriTop-20 and continuous meeting analysis keywords (CMAK) dataset for CSKWS. Evaluations on these two datasets show that our proposed AF-KWS outperforms reference schemes by a large margin, and therefore provides a decent baseline for future research. **Index Terms:** keyword spotting, continuous speech keyword spotting, speech recognition, anchor-free detector, open dataset

1. Introduction

Keyword spotting is a classical problem in speech recognition. The most well-known application is trigger word detection for smart speakers and smart devices, such as “Okay Google” [1, 2] and “Hey Siri” [3, 4, 5]. Another widely discussed application is speech command recognition for remote control, such as “Yes”, “No”, “Up”, and “Down” [6]. In this paper, however, we are interested in continuous speech keywords spotting (CSKWS), which finds many applications in prohibited words filtering and recorded meeting understanding. The latter is attracting increasing interests due to the explosive increase of on-line meetings and meeting recordings since the COVID-19 pandemic. By defining a small set of meeting-related keywords and key phrases, such as “Start”, “Agenda”, “To-do”, and “Summarize” and detecting them with a dedicated model, one can divide a meeting into chapters in a cost-effective way, compared to running automatic speech recognition (ASR) or large-vocabulary continuous speech recognition (LVCSR) throughout the meeting recording.

Existing technologies and datasets developed for trigger word detection [7] and speech command recognition [8, 9] are not readily applicable to CSKWS. In trigger word detection, the objective is to detect a single pre-defined trigger word or trigger phrase on an edge device. In speech command recognition, it is often assumed that the keywords are solitary, with clear and easily detectable boundaries with other words. In Google Speech Commands dataset [6], each one-second-long speech sample contains one separated keyword. The main challenge of CSKWS, which makes it a unique research problem, lies in the presence of multiple target keywords and the interference of context in continuous speech.

In this paper, we attempt to standardize the definition,

benchmark dataset, and evaluation criteria of the CSKWS problem. Previously, Rostami et al. [10] collected the football keyword dataset (FKD) for CSKWS, but they only collected the lone word audio and then synthesized the continuous speech. As only keywords are authentic speech, one could expect a simple authentic/synthetic speech discriminator to perform well on FKD. We propose two benchmark datasets for the CSKWS problem. The first one, named LibriTop-20, is derived from the LibriSpeech dataset [11]. It contains 20 most frequently appeared words with at least two syllables. There are a total of 256k utterances and 530k keywords. The second one is a brand new dataset named continuous meeting analysis keywords (CMAK). We define 24 meeting-structure-related keywords, such as “Agenda”, “Start”, and “Question”. We collect over 15 sentences for each keyword and employ 32 human voices and 136 synthetic voices to produce 350k utterances.

We present a novel algorithm for the CSKWS problem. Specifically, we treat CSKWS as an object detection problem rather than a classification problem. In computer vision (CV), anchor-based object detectors utilize predefined anchors to enumerate possible locations, scales, and aspect ratios which introduce redundant computing, while anchor-free object detectors directly learn the object existence possibility and bounding box coordinates. Inspired by the work in CV [12, 13, 14], we propose an anchor-free method, named AF-KWS, for speech keyword spotting. We predict the center location of a keyword in continuous speech and regress the length of the keyword by a simple convolutional head. An additional offset head is employed to refine the predicted center position. While these designs are directly extended from visual object detection, we introduce an auxiliary “unknown” class for the interfering words in continuous speech. Such a design differentiates interfering words from the non-word background containing silence and environment noise, and thus significantly improves the accuracy of the keyword spotting model.

We evaluate the proposed AF-KWS method on the two benchmark datasets we have built. In addition to the commonly adopted evaluation metrics, such as false rejection rate (FRR), false alarms (FAs) per hour, and accuracy, we also adopt the object detection metric mAP. We compare AF-KWS with two reference methods adapted from the SOTA trigger word detection algorithms by an inserted sliding-window region proposal module. Results show that, AF-KWS outperforms these adapted methods by a large margin. This justifies our claim that CSKWS is a unique research problem and should not be viewed as a simple extension of trigger words detection or speech command recognition. We will release our benchmark datasets¹ as well as the codes of AF-KWS. We hope our work can pave the way for future research on this interesting CSKWS problem.

¹Utterances of human voice will be excluded from CMAK to avoid privacy concerns.

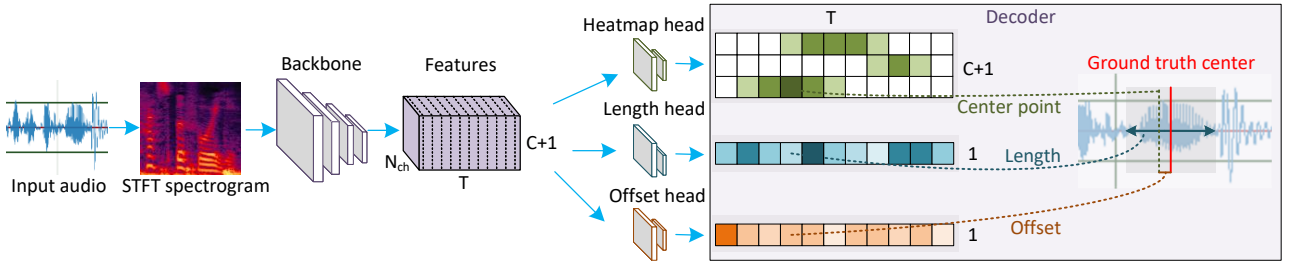


Figure 1: Overview of the proposed anchor-free keywords spotting (AF-KWS) method.

2. Anchor-Free Keywords Spotting

2.1. Overview of AF-KWS

CSKWS is treated as a 1D object detection problem. Given a predefined keywords set $K = \{k_1, k_2, \dots, k_C\}$ of size C and an input audio of length r , the task of CSKWS is to find the locations and lengths of all the keywords in the input audio.

Figure 1 provides an overview of the proposed method AF-KWS. This end-to-end framework starts with computing Short-Time Fourier Transform (STFT) spectrogram for the input audio. It then employs a trainable backbone to extract the feature map $F \in \mathbb{R}^{T \times (C+1) \times N_{ch}}$, where T is the temporal resolution and N_{ch} is the number of channels. Inspired by the anchor-free detection method in computer vision [12], we employ three trainable prediction heads to regress a keyword heatmap $\hat{Y} \in [0, 1]^{T \times (C+1)}$, a keyword length vector $\hat{L} \in \mathbb{R}^T$, and an offset vector $\hat{O} \in \mathbb{R}^T$, respectively. Offsets are used to correct word positions to compensate for the loss of precision in limited temporal resolution caused by down-sampling and framing in frequency domain.

A large $\hat{Y}_{t,c}$ suggests a high probability that a keyword k_c is centered at time index t . When the value of c is from 1 to C , the detected keyword belongs to the predefined keywords set K . However, when $c = C + 1$, the detected keyword belongs to an auxiliary class “unknown” (word). Introducing an auxiliary class presents a different design from what is commonly adopted in visual object detection, where undefined objects and non-objects are all treated as background. We believe that such an auxiliary class which differentiates other words from non-speech background is necessary in CSKWS. This is validated in the ablation studies in Section 4.

2.2. Training

The proposed AF-KWS network is trained in a supervised manner. Thanks to the off-the-shelf forced alignment tool[15], we are able to label a speech with the locations of each spoken word in an automatic way. Let $W = \{w_1, w_2, \dots, w_N\}$ be the set of aligned words in an input speech utterance, and let $cls(w_i)$, $loc_{pc}(w_i)$, and $len(w_i)$ be the class label, center location, and length of word w_i ($i = 1, \dots, N$), respectively. The value of $cls(\cdot)$ is between 1 and $C + 1$, and the values of $loc_{pc}(\cdot)$ a real number not exceeding T . We further split the precise location of w_i into an integer part $loc(w_i) = \lfloor loc_{pc}(w_i) \rfloor$ and an offset $ofs(w_i) = loc_{pc}(w_i) - loc(w_i)$. For training, we shall convert these automatic alignment results into the ground truth of network output, denoted by Y , L , and O . We shall also define the loss functions for the three prediction tasks.

The ground truth keyword heatmap Y is a smoothed heatmap. For each word $w_i \in W$, we first set $Y_{loc(w_i), cls(w_i)}$ to 1 and then use a Gaussian kernel to spread the heat to its

neighborhood along the temporal axis. The standard deviation σ_{w_i} of the Gaussian kernel is proportional to the length of the word, or $\sigma_{w_i} = len(w_i) \times \gamma$, where γ is a size adaptive factor. We set $\gamma = 0.125$ based on ablation studies.

For keypoint heatmap prediction, we use a penalty-reduced pixel-wise logistic regression with focal loss [16]

$$L_h = -\frac{1}{N} \begin{cases} \sum_{t,c} (1 - \hat{Y}_{t,c})^\alpha \log(\hat{Y}_{t,c}) & \text{if } Y_{t,c} = 1 \\ \sum_{t,c} (1 - Y_{t,c})^\beta (\hat{Y}_{t,c})^\alpha & \\ \sum_{t,c} \log(1 - \hat{Y}_{t,c}) & \text{otherwise} \end{cases} \quad (1)$$

where α and β are hyper-parameters of the focal loss, and N is the number of keywords in the input audio.

The groundtruth for keyword length is a T -dimensional vector L , where L_t is the length of the keyword centered at time index t . Note that the temporal resolution T is selected such that the center locations of two adjacent words never appear at the same time index. For each of the N aligned words, we set $L_{loc(w_i)} = len(w_i)$. All the unassigned elements are initialized to zero. We use L1 loss for length prediction and it is computed by:

$$L_{len} = \frac{1}{N} \sum_{i=1}^N |\hat{L}_{t_i} - L_{t_i}| \quad (2)$$

The offset prediction results provides an adjustment to the center locations of detected keywords. Same as the keyword length prediction, the offset ground truth is a T -dimensional vector, denoted by $O \in \mathbb{R}^T$. We set $O_{loc(w_i)} = ofs(w_i)$ for each word w_i and the rest elements are initialized to zero. During training, we use L1 loss for offset prediction:

$$L_{offset} = \frac{1}{N} \sum_{i=1}^N |\hat{O}_{loc(w_i)} - O_{loc(w_i)}| \quad (3)$$

The overall training loss is:

$$L_{cskws} = L_h + \lambda_{len} L_{len} + \lambda_{offset} L_{offset} \quad (4)$$

where λ_{len} and λ_{offset} is a loss weight factor for length loss and offset loss. We follow previous practice [12] to set $\lambda_{len} = 0.1$ and $\lambda_{offset} = 1$.

2.3. Inference process

During inference, we obtain \hat{Y} , \hat{L} , and \hat{O} for each audio sample. We need to convert them into a set of detected keywords $\hat{W} = \{\hat{w}_1, \dots, \hat{w}_M\}$ and associated properties including $cls(\hat{w}_i)$, $loc_{pc}(\hat{w}_i)$, and $len(\hat{w}_i)$.

Table 1: Statistics of the LibriTop-20 dataset

Split	Samples	Words	Keywords
Train	251,229	9,521,509	520,701
Dev	2,953	80,210	5,337
Test	2,684	74,952	4,602
Total	256,866	9,676,671	530,640

With the predicted keywords heatmap \hat{Y} , we examine the vector \hat{Y}_c for each keyword class c . We identify the peak locations whose response score is greater than its neighbors along the temporal axis. After examine all the keyword classes, we select the top M locations by their response scores. If $\hat{Y}_{t,c}$ is selected, a new keyword w with $cls(w) = c$, $len(w) = \hat{L}_t$, and $loc_{pc}(w) = t + \hat{O}_t$ is added to set \hat{W} . Finally, we can obtain a bounding box for this keyword as $[loc_{pc}(w) - len(w)/2, loc_{pc}(w) + len(w)/2]$. $\hat{Y}_{t,c}$ is directly used as the confidence score of this detection result. We do not perform non-maximum suppression (NMS) on the output.

2.4. Implementation details

The input audio is 16,000 Hz mono audio, and the input length for the model is $r=5.11s$. We first extract STFT features with $hop_length = 160$, $win_length = 400$ and $filter_length = 510$, and then down-sampling them to temporal resolution $T = 128$. We use up-convolutional residual networks (ResNet) [17] Res34 as our backbone model. The maximum number of detected keywords is set to $M = 30$, translating to a speech rate of about 350 words per minute (wpm), which is much faster than a normal speed of 160 to 180 wpm as specified by BBC [18]. In the ablation study, we compare our method with a setting without the auxiliary unknown class. In that case, M is set to 3 or 2, depending on the datasets, to avoid too many false alarms. We apply several types of data augmentation, with the same probability of 0.2, to the training data, including additive noise, reverberation, and pitch shifting.

3. Dataset

The availability and quality of benchmark datasets are of great importance to the development of a research topic. There are few dedicated datasets for the CSKWS task. One related dataset is the hybrid real-synthetic football keyword dataset FWD [10]. An earlier work [19] tried to build a dataset by selecting the six most frequently occurring words in TIMIT [20]. However, TIMIT is a small dataset with only 5.4 hours speech. In our work, we follow a similar practice as [19] to build a much larger benchmark dataset from LibriSpeech. We also define and collect a new meeting keywords dataset named CMAK to facilitate meeting-related downstream tasks.

3.1. LibriTop dataset

The LibriTop dataset is derived from LibriSpeech [11], which is a continuous speech corpus of approximately 1000 hours of 16kHz read English speech. As LibriSpeech is built for the ASR task, all samples in it have been carefully segmented and aligned. We first calculate the word frequencies in LibriSpeech and then try to pick the 20 most frequently used words. However, due to the nature of KWS, it usually cannot be applied to single-syllable words. Therefore, the resulting LibriTop-20

Table 2: Statistics of the CMAK dataset

Source	Speakers	Utterances
Recorded	32	15,092
Synthetic	136	334,847
Total	168	349,939

dataset contains 20 top-frequency words with at least two syllables. These words are: Very, Into, Little, About, Only, Upon, Any, Before, Other, Over, After, Never, Our, Mister, Again, Himself, Away, Even, Without, Every.

There are a total of 256k utterances and 530k occurrences of keywords in LibriTop-20. We keep the same data split as LibriSpeech. The detailed statistics are listed in Table 1.

3.2. CMAK dataset

CMAK dataset is collected for speech-based meeting structure analysis. We defined a total of 24 keywords, relating to the start, the end, or some specific sections of a meeting. They are:

- Start: Begin, Start, Agenda, Outline, Today, Introduce, Talk about
- End: Summarize, Conclude, Action item, To do, Follow up, Next step, Share, Plan
- Specific sections: Timeline, Milestone, Question, Comments, Additions, Highlight, Important, Crucial, Attention

CMAK is a hybrid real-synthetic dataset. Different from FWD which contains hybrid utterances, each utterance in CMAK is either real or synthetic. To collect the real data, we select more than 15 sentences from the Oxford English dictionary [21] for each keyword, and ask 32 speakers to read the scripts and record the speech. To collect the synthetic data, we randomly generate a script with 10 ~ 15 words containing one keyword at a time, and employ ESPNet [22, 23, 24, 25, 26] to synthesize the speech. We use 136 different speaker embeddings to generate 100 sentences for each keyword. The detailed statistics for the dataset are listed in Table 2.

We carry out most of the experiments in this paper using a subset of CMAK, called CMAK-7. It contains the seven keywords relating to the start of a meeting. A model trained on this subset can be applied to the first few minutes of a recorded meeting to enable the *Skip greetings* feature. Further, CMAK-7-s is a subset of CMAK-7 that only contains the synthetic speech.

4. Experiments

We evaluate the proposed AF-KWS method on LibriTop-20 and CMAK-7 datasets. LibriTop-20 is divided in the same way as LibriSpeech, and CMAK-7 is divided into 80% training set and 20% testing set. For different versions of a method, the input length is fixed. We apply repeated padding for shorter input and random crop for longer input during training. We train the model with a batch size of 64 and a learning rate of 0.00125.

4.1. Evaluation metrics

In trigger words detection, false rejection rate (FRR) and false accepts per hour (FAs) are commonly used as the evaluation metrics [27, 28, 29, 30]. In speech command recognition, accuracy is commonly used for evaluation [31, 32, 33]. CSKWS can borrow the evaluation metrics from these related tasks. However, there are more keywords in CSKWS and they appear more

Table 3: Performance comparison between the proposed AF-KWS and the reference schemes on LibriTop-20. AP@K means the AP at an IoU threshold of K%. FRR@K means the FRR when the FAs per hour is K.

Model	AP@5 \uparrow	AP@75 \uparrow	mAP \uparrow	FRR@5 \downarrow	FRR@15 \downarrow	FRR@25 \downarrow	Classification Accuracy \uparrow	RTF \downarrow
DSTC-ResNet	0.748	0.058	0.398	0.647	0.519	0.402	0.961	0.018
MHAtt-RNN	0.795	0.076	0.426	0.530	0.418	0.374	0.978	0.057
AF-KWS (ours)	0.952	0.886	0.860	0.140	0.074	0.049	N/A	0.031

Table 4: Ablation study on LibriTop-20

Model	AP@5	AP@75	FRR@5	FRR@25
AF-KWS-cls	0.876	0.062	0.570	0.265
w/o unknown	0.867	0.800	0.234	0.117
AF-KWS (ours)	0.952	0.886	0.140	0.049

frequently than in trigger words detection. Besides, CSKWS is essentially a detection task instead of classification task, so we propose to evaluate its solutions with object detection metrics known as 1D AP(average precision) and mAP (mean average precision). 1D IoU (Intersection over Union) is computed for each detection result with respect to the groundtruth, and mAP is computed in an IoU range of (0.05, 0.95) by step 0.05. This is a slightly different setting from the mAP used in visual object detection [34], because there is a higher tolerance of positional accuracy in CSKWS.

4.2. Performance comparison

CSKWS is a less studied task than speech command recognition, or wake-up words detection. There are few reference schemes and none of them have publicly available codes. Therefore, we choose to adapt the SOTA speech commands recognition methods for CSKWS task and use them as reference schemes. We pick two Google implemented models [8], MHAtt-RNN [8] and DSTC-ResNet based on MatchboxNet [9], which achieve top results on Google speech command dataset.

We insert a sliding window region proposal module before these speech commands recognition methods, which are essentially classifiers. The sliding window strategy has two parameters, namely the window size L_{in} and the sliding step L_{step} . To make sure that each keyword is included in at least one window, we apply the constraints that $L_{in} > MAX_x$ and $L_{step} < L_{in} - MAX_x$, where MAX_x is the maximum length of the keywords. We set $L_{in} = 0.5$ for LibriTop-20. There is a trade-off between detection performance and inference speed, so we grid search the best step size from 0.1s to 0.4s with 0.1s interval. We re-train the adapted KWS methods on our dataset.

Table 3 shows the comparison between AF-KWS and the two reference schemes. We list the results on both detection-oriented metrics and conventional KWS metrics. Both DSTC-ResNet and MHAtt-RNN can achieve a high classification accuracy on trimmed input, showing that our re-training is successful. However, they give poor AP numbers, because sliding window region proposal cannot provide high temporal resolution, or the inference cost will be too high. The results show that our AF-KWS model outperforms the two methods on AP and FRR with almost same inference speed².

²The real time factor (RTF) is tested on our local CPU machine with Intel Xeon CPU E5-2690 v4 @2.6GHz.

Table 5: Performance on CMAK-7 with different datasets

Dataset	Extra Data	AP@5	AP@75	FRR@5	FRR@25
Synthetic	N	0.859	0.697	0.360	0.137
Hybrid	N	0.874	0.633	0.356	0.159
Hybrid	Y	0.916	0.731	0.238	0.126

4.3. Ablation experiments

Our AF-KWS detector outperforms two reference schemes adapted from classifiers. To verify that the gain is not due to a stronger backbone, we carry out an ablation study which replaces the three detection heads in AF-KWS with a classification head. Accordingly, the input length is adjusted to 0.5s and a sliding window inference strategy is adopted. Table 4 shows that this AF-KWS-cls model does not perform well, confirming our claim that CSKWS should be treated as a detection task.

AF-KWS is inspired by anchor-free visual object detectors, but we make one key design choice that differentiates it from normal detectors. AF-KWS defines an auxiliary “unknown” class in addition to the predefined keyword classes. The reason is that if only the defined keywords are treated as foreground, the background will include other words, non-speech sound, and silence, which is too diverse to handle. To verify this, we implement a “w/o unknown” version of AF-KWS. As Table 4 shows, its performance is significantly inferior to our method.

Finally, we study the impact of dataset using different versions of CMAK-7. CMAK is a hybrid dataset with both real recordings and synthetic speech. Due to privacy concerns, we may only able to release the synthetic part of CMAK. In Table 5, the results of the first two rows show that our proposed model achieves similar results on hybrid and synthetic datasets, given that the test set has the same data distribution as the training set. The last row shows the results when we use additional LibriSpeech data for training. The performance is significantly boosted on the same test set as used in the second row.

5. Conclusion

In this paper, we have discussed the continuous speech keyword spotting (CSKWS) task which finds many important applications in the post-pandemic era. We have introduced the LibriTop-20 dataset derived from LibriSpeech and have collected a new continuous meeting analysis keywords dataset named CMAK. We have proposed a novel anchor-free keywords detection method which is tailored for the CSKWS task. The proposed AF-KWS method achieves better performance than the adapted speech command recognition methods, which verifies our claim that CSKWS should be treated as a detection problem instead of a classification problem. AF-KWS provides a reasonable baseline for future research on this interesting CSKWS problem.

6. References

- [1] G. Chen, C. Parada, and G. Heigold, "Small-footprint keyword spotting using deep neural networks," in *2014 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2014, pp. 4087–4091.
- [2] R. Alvarez and H.-J. Park, "End-to-end streaming keyword spotting," in *ICASSP 2019-2019 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2019, pp. 6336–6340.
- [3] T. Higuchi, M. Ghasemzadeh, K. You, and C. Dhir, "Stacked 1d convolutional networks for end-to-end small footprint voice trigger detection," *Proc. Interspeech 2020*, pp. 2592–2596, 2020.
- [4] S. Sigtia, J. Bridle, H. Richards, P. Clark, E. Marchi, and V. Garg, "Progressive voice trigger detection: Accuracy vs latency," in *ICASSP 2021-2021 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2021, pp. 6843–6847.
- [5] S. Adya, V. Garg, S. Sigtia, P. Simha, and C. Dhir, "Hybrid transformer/ctc networks for hardware efficient voice triggering," *Proc. Interspeech 2020*, pp. 3351–3355, 2020.
- [6] P. Warden, "Speech commands: A dataset for limited-vocabulary speech recognition," *arXiv preprint arXiv:1804.03209*, 2018.
- [7] J. Hou, Y. Shi, M. Ostendorf, M.-Y. Hwang, and L. Xie, "Region proposal network based small-footprint keyword spotting," *IEEE Signal Processing Letters*, vol. 26, no. 10, pp. 1471–1475, 2019.
- [8] O. Rybakov, N. Kononenko, N. Subrahmanya, M. Visontai, and S. Laurenzo, "Streaming keyword spotting on mobile devices," *Proc. Interspeech 2020*, pp. 2277–2281, 2020.
- [9] S. Majumdar and B. Ginsburg, "Matchboxnet: 1d time-channel separable convolutional neural network architecture for speech commands recognition," *Proc. Interspeech 2020*, pp. 3356–3360, 2020.
- [10] A. M. Rostami, A. Karimi, and M. A. Akhaee, "Efficientnet-absolute zero for continuous speech keyword spotting," *arXiv preprint arXiv:2012.15695*, 2020.
- [11] V. Panayotov, G. Chen, D. Povey, and S. Khudanpur, "Librispeech: an asr corpus based on public domain audio books," in *2015 IEEE international conference on acoustics, speech and signal processing (ICASSP)*. IEEE, 2015, pp. 5206–5210.
- [12] X. Zhou, D. Wang, and P. Krähenbühl, "Objects as points," *arXiv preprint arXiv:1904.07850*, 2019.
- [13] C. Zhu, Y. He, and M. Savvides, "Feature selective anchor-free module for single-shot object detection," in *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 2019, pp. 840–849.
- [14] T. Kong, F. Sun, H. Liu, Y. Jiang, L. Li, and J. Shi, "Foveabox: Beyond anchor-based object detection," *IEEE Transactions on Image Processing*, vol. 29, pp. 7389–7398, 2020.
- [15] M. McAuliffe, M. Socolof, S. Mihuc, M. Wagner, and M. Sonderegger, "Montreal forced aligner: Trainable text-speech alignment using kaldii," in *Proc. Interspeech*, vol. 2017, 2017, pp. 498–502.
- [16] T.-Y. Lin, P. Goyal, R. Girshick, K. He, and P. Dollár, "Focal loss for dense object detection," in *Proceedings of the IEEE international conference on computer vision*, 2017, pp. 2980–2988.
- [17] B. Xiao, H. Wu, and Y. Wei, "Simple baselines for human pose estimation and tracking," in *Proceedings of the European conference on computer vision (ECCV)*, 2018, pp. 466–481.
- [18] E. Rodero, "A comparative analysis of speech rate and perception in radio bulletins," *Text & Talk*, vol. 32, no. 3, pp. 391–411, 2012.
- [19] P. Baljekar, J. F. Lehman, and R. Singh, "Online word-spotting in continuous speech with recurrent neural networks," in *2014 IEEE Spoken Language Technology Workshop (SLT)*. IEEE, 2014, pp. 536–541.
- [20] J. S. Garofolo, L. F. Lamel, W. M. Fisher, J. G. Fiscus, and D. S. Pallett, "Darpa timit acoustic-phonetic continuous speech corpus cd-rom. nist speech disc 1-1.1," *NASA STI/Recon technical report n*, vol. 93, p. 27403, 1993.
- [21] O. E. Dictionary, "Oxford english dictionary," *Simpson, Ja & Weiner, Esc*, p. 3, 1989.
- [22] S. Watanabe, T. Hori, S. Karita, T. Hayashi, J. Nishitoba, Y. Unno, N. Enrique Yalta Soplin, J. Heymann, M. Wiesner, N. Chen, A. Renduchintala, and T. Ochiai, "ESPnet: End-to-end speech processing toolkit," in *Proc. Interspeech 2018*, 2018, pp. 2207–2211. [Online]. Available: <http://dx.doi.org/10.21437/Interspeech.2018-1456>
- [23] T. Hayashi, R. Yamamoto, K. Inoue, T. Yoshimura, S. Watanabe, T. Toda, K. Takeda, Y. Zhang, and X. Tan, "Espnet-TTS: Unified, reproducible, and integratable open source end-to-end text-to-speech toolkit," in *Proceedings of IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2020, pp. 7654–7658.
- [24] H. Inaguma, S. Kiyono, K. Duh, S. Karita, N. Yalta, T. Hayashi, and S. Watanabe, "ESPnet-ST: All-in-one speech translation toolkit," in *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics: System Demonstrations*. Online: Association for Computational Linguistics, Jul. 2020, pp. 302–311. [Online]. Available: <https://www.aclweb.org/anthology/2020.acl-demos.34>
- [25] C. Li, J. Shi, W. Zhang, A. S. Subramanian, X. Chang, N. Kamo, M. Hira, T. Hayashi, C. Boeddeker, Z. Chen, and S. Watanabe, "ESPnet-SE: End-to-end speech enhancement and separation toolkit designed for ASR integration," in *Proceedings of IEEE Spoken Language Technology Workshop (SLT)*. IEEE, 2021, pp. 785–792.
- [26] S. Arora, S. Dalmia, P. Denisov, X. Chang, Y. Ueda, Y. Peng, Y. Zhang, S. Kumar, K. Ganesan, B. Yan *et al.*, "Espnet-slu: Advancing spoken language understanding through espnet," *arXiv preprint arXiv:2111.14706*, 2021.
- [27] A. H. Michaely, X. Zhang, G. Simko, C. Parada, and P. Aleksic, "Keyword spotting for google assistant using contextual speech recognition," in *2017 IEEE Automatic Speech Recognition and Understanding Workshop (ASRU)*. IEEE, 2017, pp. 272–278.
- [28] A. Coucke, M. Chlieh, T. Gisselbrecht, D. Leroy, M. Poumeyrol, and T. Lavril, "Efficient keyword spotting using dilated convolutions and gating," in *ICASSP 2019-2019 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2019, pp. 6351–6355.
- [29] C. Shan, J. Zhang, Y. Wang, and L. Xie, "Attention-based end-to-end models for small-footprint keyword spotting," *Proc. Interspeech 2018*, pp. 2037–2041, 2018.
- [30] 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 *Proc. Interspeech 2017*, 2017, pp. 3607–3611.
- [31] B. Kim, S. Chang, J. Lee, and D. Sung, "Broadcasted residual learning for efficient keyword spotting," *arXiv preprint arXiv:2106.04140*, 2021.
- [32] C.-C. Kao, M. Sun, Y. Gao, S. Vitaladevuni, and C. Wang, "Sub-band convolutional neural networks for small-footprint spoken term classification," *Proc. Interspeech 2019*, pp. 2195–2199, 2019.
- [33] S. Choi, S. Seo, B. Shin, H. Byun, M. Kersner, B. Kim, D. Kim, and S. Ha, "Temporal convolution for real-time keyword spotting on mobile devices," *Proc. Interspeech 2019*, pp. 3372–3376, 2019.
- [34] Tsung-Yi, G. Lin, M. Patterson, R., Y. Ronchi, M. Cui, S. Maire, L. Belongie, R. Bourdev, J. Girshick, P. Hays, D. Perona, L. Ramanan, P. Zitnick, and Dollar. cocodataset-detection-eval. [Online]. Available: <https://cocodataset.org/#detection-eval>