

# WEKWS: A PRODUCTION FIRST SMALL-FOOTPRINT END-TO-END KEYWORD SPOTTING TOOLKIT

Jie Wang<sup>1,4</sup>, Menglong Xu<sup>3,4</sup>, Jingyong Hou<sup>2,4</sup>, Binbin Zhang<sup>3,4</sup>, Xiao-Lei Zhang<sup>1</sup>, Lei Xie<sup>2</sup>, Fuping Pan<sup>3</sup>

<sup>1</sup>School of Marine Science and Technology, Northwestern Polytechnical University, Xi'an, China

<sup>2</sup>Audio, Speech and Language Processing Group (ASLP@NPU),  
School of Computer Science, Northwestern Polytechnical University, Xi'an, China

<sup>3</sup>Horizon Robotics, Beijing, China

<sup>4</sup>WeNet Open Source Community

## ABSTRACT

Keyword spotting (KWS) enables speech-based user interaction and gradually becomes an indispensable component of smart devices. Recently, end-to-end (E2E) methods have become the most popular approach for on-device KWS tasks. However, there is still a gap between the research and deployment of E2E KWS methods. In this paper, we introduce WeKws, a production-quality, easy-to-build, and convenient-to-be-applied E2E KWS toolkit. WeKws contains the implementations of several state-of-the-art backbone networks, making it achieve highly competitive results on three publicly available datasets. To make WeKws a pure E2E toolkit, we utilize a refined max-pooling loss to make the model learn the ending position of the keyword by itself, which significantly simplifies the training pipeline and makes WeKws very efficient to be applied in real-world scenarios. The toolkit is publicly available at <https://github.com/wenet-e2e/wekws>.

**Index Terms**— Keyword spotting, production first, end-to-end

## 1. INTRODUCTION

Keyword spotting (KWS) is the task of detecting predefined keywords from continuous audio streams. Wake-up word (WuW) detection, as a special task of KWS, has become a typical and indispensable component in internet of things (IoT) devices, e.g. smart speakers and mobile phones, enabling users to have a fully hands-free voice interactive user experience. A WuW detection system needs to process streaming audio and runs locally and persistently on IoT devices. Therefore, the system should have a small memory footprint and computational cost, while maintaining a low latency (real-time response) and high detection accuracy.

Thanks to the emergence of some neural networks with long-term modeling capabilities, end-to-end (E2E) based KWS methods recently have drawn much attention for their simplicity of training/decoding [1–5]. Under the E2E setting,

KWS becomes a keyword/non-keyword binary classification task. Posteriors corresponding to different keywords are predicted directly by a model. Then a system can easily detect keywords, by comparing the posteriors with manually defined thresholds of the keywords. E2E based methods yield significant better performances over conventional methods [6–10].

However, the present E2E KWS methods still have defects. Firstly, to train the model, they usually need a force-alignment procedure to get the start-end position of keywords in utterances [1, 3–5]. Secondly, some of them rely on the entire input audio sequence to make decisions [11–13], so are not applicable to streaming tasks. In addition to the above shortcomings of the present E2E KWS methods, there is a lack of an open-source toolkit like Wenet [14] in the field of small-footprint KWS to bridge the gap between research and production. There are several good speech-processing toolkits, e.g. Kaldi [15], Fairseq [16] and Honk [17], that have implemented KWS as part of their features. However, they are either too complex [15, 16] or far away from real production in design [17].

In this paper, to tackle the above-mentioned problems, we introduce WeKws, a production-oriented lightweight E2E KWS toolkit. The key advantages of WeKws are as follows.

- **Alignment-free:** WeKws is an alignment-free E2E toolkit. There is no need to use an automatic speech recognition (ASR) or speech activity detection (SAD) system to get the alignments or ending timestamps of keywords, which significantly simplifies the KWS training pipeline.

- **Production ready:** We try our best to bridge the gap between research and production when we design WeKws. Wekws uses causal convolutions to achieve streaming KWS. All modules of WeKws conform to the requirements of TorchScript<sup>1</sup>. Therefore, a model trained with WeKws can be exported by Torch Just In Time (JIT), transformed to Open Neural Network Exchange<sup>2</sup> (ONNX) format, and easily adopted in many deployment environments.

<sup>1</sup><https://pytorch.org/docs/stable/jit.html>

<sup>2</sup><https://github.com/onnx/onnx>

- **Light weight:** WeKws is designed specifically for E2E KWS with clean and simple codes, and it only depends on PyTorch. The trained model is lightweight, and is able to run on embedded devices.

- **Competitive results:** WeKws achieves competitive results on several public KWS benchmarks compared with other recently proposed KWS systems. Moreover, WeKws does not require alignments and decoding graphs, while the comparison systems are unable to do so, making their training pipelines complicated and rely on other heavy toolkits.

## 2. WEKWS

### 2.1. System design

As shown in Figure 1, our proposed WeKws is designed with three layers.

#### 2.1.1. Layer 1: Data preparation module and an on-the-fly feature extraction and argumentation.

In the data preparation module, all we need to do is to prepare an audio list and utterance-level keyword category labels required for model training. WeKws uses an on-the-fly feature extraction and argumentation procedure. Each utterance is first re-sampled to a specific sampling rate, followed by speed perturbation and Mel-filter bank feature extraction. Feature-level SpecAugment [18] are applied to generate the input of the KWS model. It is worth noticing that the above procedure is performed dynamically and on-the-fly in each training mini-batch. Compared with traditional offline feature preparation, this on-the-fly procedure not only helps save disk usage but also enriches the diversity of training examples across different epochs, which improves the robustness of the model.

#### 2.1.2. Layer 2: Model training and testing.

We have provided an easy-to-use platform for researchers to develop and evaluate their own algorithms. At present, we have implemented a unified model architecture with several popular KWS backbones and a refined max-pooling KWS training objective which greatly simplified the KWS training pipeline. The detailed model architecture and the refined max-pooling training objective will be introduced in the next two sub-sections. We also provided standard KWS testing criteria. Competitive results on several public KWS benchmarks were achieved with the proposed model and training method.

#### 2.1.3. Layer 3: Model exportation and development.

The trained model supports TorchScript and ONNX exportation, and therefore can be easily applied to different platforms. Currently, we support deploying the WeKws production model on three mainstream platforms, i.e., x86 server,




Fig. 1. System design of the proposed WeKws toolkit.

android, and Raspberry Pi. Furthermore, both *float32* model and quantized *int8* model are supported in WeKws, where the quantized *int8* model can increase the inference speed when hosted on embedded devices such as the ARM-based Android and Raspberry Pi platforms.

### 2.2. Model architecture

The overall model architecture of WeKws is shown in Figure 2. It consists of four parts, starting with a global cepstral mean and variance normalization (CMVN) layer to normalize the input acoustic features to a normal distribution. This is followed by a linear layer which maps the dimensions of the input features to the required dimensions. Then there is a backbone network, which can be recurrent neural network (RNN), temporal convolutional network (TCN) [3] or multi-scale depthwise temporal convolution (MDTC) [19]. At the end of the model, there are several binary classifiers, each of which has a single output node with a sigmoid activation to predict the posterior probability of the corresponding keyword. As [4] did, for each keyword, we add an independent binary classifier after the backbone network to deal with the scenarios with multiple keywords. Note that the backbone network is shared by these binary classifiers.

WeKws current supports the following three backbones: 1) RNN or its improved version LSTM [1], which has been widely used in speech recognition and other speech tasks; 2) TCN [3] or its lightweight version depthwise separable TCN (DS-TCN) with dilated convolutions to increase the receptive field; 3) MDTC [19], an advanced backbone network that extracts multi-scale feature representation efficiently. Note that as we aim to achieve a streaming KWS system, we use causal convolutions in all convolution-based neural networks.

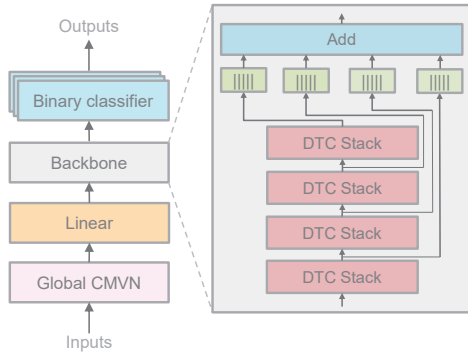
### 2.3. Refined max-pooling based training objective

A refined max-pooling loss is used as the training objective of WeKws, which is formulated as:

$$\mathcal{L} = -[y_i \ln p_i^* + (1 - y_i) \ln (1 - p_i^*)] \quad (1)$$

with

$$p_i^* = \max_j (\mathbf{p}_{ij}), \quad \forall j = m + 1, m + 2, \dots, N. \quad (2)$$



**Fig. 2.** A WeKws implementation with MDTC as the backbone, where the depthwise temporal convolution (DTC) as a basic block. The details of the MDTC network can be found in [19].

**Table 1.** FRR(%) comparison of the proposed WeKws and other KWS systems with FAH fixed at 0.5, on keywords “Hi xiaowen” and “Nihao wenwen”.

System	#Params	Hi xiaowen	Nihao wenwen
Wang at al. [10]	57k	0.7	0.5
WeKws	153k	<b>0.50</b>	<b>0.43</b>

where  $\mathbf{p}_{ij}$  is the predicted posterior probabilities,  $m$  is the minimum duration frames of the keyword, and  $N$  is the number of frames of the  $i$ -th utterance in each training mini-batch. Note that the minimum duration  $m$  is calculated statistically on all positive samples in the training set.

By using the max-pooling loss, the model learns the ending timestamps of the keyword automatically, and thus gets rid of the dependence on alignments or the ending timestamps of the keyword. Specifically, for the positive samples, the max-pooling loss will only optimize the frames with the highest posterior probabilities, and ignore the other frames. For the negative samples, the max-pooling loss will minimize the frames with the highest posterior probabilities, therefore, the posterior probabilities of all frames in the negative samples will be minimized.

### 3. EXPERIMENTS

We conducted experiments on several public KWS benchmarks. We will introduce our experimental setup and report the results in this section.

#### 3.1. Experimental setup

We used Mobvoi (SLR87)<sup>3</sup>, Snips<sup>4</sup> and Google Speech Command (GSC) [23] corpora to evaluate our proposed

<sup>3</sup><https://www.openslr.org/87>

<sup>4</sup><https://research.snips.ai/datasets/keyword-spotting>

**Table 2.** FRR (%) comparison of the proposed WeKws and other KWS systems on keyword “Hey snips”.

System	#Params	FAH=0.5	FAH=1
Zhang at al. [5]	244k	3.53	2.82
Coucke at al. [20]	222k	<b>0.12</b>	–
WeKws	153k	<b>0.12</b>	<b>0.08</b>

**Table 3.** Comparison of the proposed WeKws and other KWS systems on Google Speech Command.

System	#Params	Accuracy (%)
Zhang at al. [21]	109K	97.20
Ding at al. [22]	404K	97.56
WeKws	158K	<b>97.97</b>

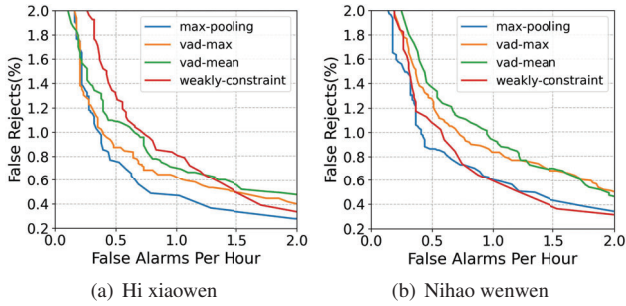
WeKws. The Mobvoi is a Mandarin corpus designed specifically for WuW tasks. There are two keywords in the corpus, each having about 36K utterances. The number of non-keyword utterances is about 183K. We followed the same setup as [10] to use the corpus. The Snips is a crowd-sourced WuW corpus. The keyword of the corpus is “Hey snips”, and there are around 11K keyword utterances and 86.5K non-keyword utterances in the corpus. The detailed corpus information can be found in [3] and we followed the recommended setup in [3] to carry out our experiments. The GSC corpus consists of 64,721 one-second-long recordings of 30 words by 1881 different speakers. Following the typical setup like [19], we used utterances in *validation.list.txt* and *testing.list.txt* as validation and testing data, respectively, and used the other utterances as training data.

We used 40-dimensional Mel-filter bank (Fbank) features, with 25 ms window size and 10 ms window shift, as the model input. We used Adam, with an initial learning rate of  $1e-3$  and an L2 weight decay of  $1e-4$ , as an optimizer for the model training. The batch size was set to 128 utterances. The proposed WeKws was trained with 80 epochs. We selected 30 best models on the development set, from the total 80 models saved after each epoch, and then averaged them to obtain the final model.

#### 3.2. Experimental results

##### 3.2.1. WeKws vs. other KWS systems

In this section, we have evaluated our proposed WeKws on three public benchmarks and compared WeKws with other competitive KWS systems proposed recently [5, 10, 20–22]. We used MDTC with the best setup in [19] as the backbone of our model, and the refined max-pooling loss introduced in Section 2.3 to train the model. In Table 1, compared with an LF-MMI-based KWS system which relies on a graph-based decoding algorithm [10], our WeKws improves the false rejection rate (FRR) relatively by around 28% and 14% on “Hi



**Fig. 3.** DET curve comparison of different models trained with max-pooling, vad-max, vad-mean, and weakly-constraint training objectives, on keywords “Hi xiaowen” and “Nihao wenwen”.

xiaowen” and “Nihao wenwen”, respectively, on the Mobvoi corpus. In Table 2, we compare our WeKws with two E2E methods on the “Snips” corpus. Our WeKws achieves competitive results with the two methods, and a smaller memory footprint than the latter. Moreover, the comparison methods need a SAD system to get the ending timestamps of keywords in training, which is complicated, while our method avoids to do so. In Table 3, WeKws achieves the best result, compared with two recent KWS methods on the GSC dataset. The above results support the superiority of our WeKws.

### 3.2.2. Comparison of different training objectives

To evaluate the effect of max-pooling loss, we compared it with several representative loss functions in literature. The **vad-mean** loss mainly follows [3, 5], in which the authors utilize the SAD technique to get the ending timestamps of the keyword in the positive samples, and maximize the posteriors of all frames within a given time interval around the ending timestamp. In our implementation, we maximized the average of the posteriors within the given time interval, and the time interval was set to 5 frames. The **vad-max** loss follows [1], where a max-pooling function is used to pick the most informative frame within the alignment of the keyword. Different from [1], we restrict the picked frame to be around the ending timestamp of the keyword, and the range was set to 40 frames. The **weakly-constraint** loss was proposed by [4]. It only uses the constraint of the ending timestamp at the early stages of training and relaxes the constraint after a preset training epoch. In our experiments, we used the constraint in the first 5 epochs.

All experiments in this section were based on the DS-TCN backbone, and the comparison results are shown in Figure 3 and Table 4. From Figure 3, we see that the **max-pooling** loss achieves the best performance in most cases on the two keywords. More specifically, from Table 4 we see that the **max-pooling** loss overcomes all comparison losses at 0.5/1 False Alarm per hour. In addition, it is worth men-

**Table 4.** FRR(%) comparison of different training objectives with FAH fixed at 0.5/1.0, on keywords “Hi xiaowen” and “Nihao wenwen”.

Training objective	Hi xiaowen	Nihao wenwen
vad-max	0.87/0.64	1.28/0.84
weakly-constraint	1.32/0.81	1.08/0.65
vad-mean	1.09/0.70	1.38/0.95
max-pooling	<b>0.75/0.47</b>	<b>0.87/0.61</b>

**Table 5.** FRR(%) comparison of different backbones with FAH fixed at 0.5/1.0.

Backbone	#Params	Hi xiaowen	Nihao wenwen
TCN	2M	0.63/0.42	0.69/0.44
DSTCN	287K	0.75/0.47	0.87/0.61
GDSTCN	124K	0.84/0.36	0.64/0.35
MDTC	153K	<b>0.50/0.22</b>	<b>0.43/0.29</b>

tioning that the **max-pooling** loss does not use any additional supervised information, which makes WeKws a simple and effective KWS system.

### 3.2.3. Comparison of different backbone networks

To further verify the effectiveness of the **max-pooling** loss, we applied it to multiple different backbone networks such as dilated TCN [24], grouped DS-TCN [25], and MDTC [19]. The comparison results are listed in Table 5. It shows that all backbone networks obtain good results with the **max-pooling** loss. From Table 5, we also see that the depthwise separable convolution significantly reduces the number of parameters with a small or no performance degradation. In addition, we see that MDTC achieves the best trade-off between performance and model footprint. Compared with DSTCN, MDTC uses almost half of the parameters to achieve 33.3%, 53.2%, 50.6%, and 52.5% relative reduction on FRR at 0.5/1.0 false alarm per hour on “Hi xiaowen” and “Nihao wenwen” respectively.

## 4. CONCLUSIONS

In this paper, we present a production-quality, easy-to-build, and convenient-to-be-applied open-source E2E KWS toolkit called WeKws. With a pure E2E training objective and elaborate designed modules, WeKws can be trained, exported and applied on different platforms easily. WeKws helps to bridge the gap between research and production of KWS methods and provide an easy-to-use platform for researcher and engineer. Evaluations on three KWS benchmarks demonstrate that WeKws achieves competitive or better results compared with many state-of-the-art KWS methods.

## 5. REFERENCES

- [1] Ming Sun, Anirudh Raju, George Tucker, Sankaran Panchapagesan, Gengshen Fu, Arindam Mandal, Spyros Matsoukas, Nikko Strom, and Shiv Vitaladevuni, “Max-pooling loss training of long short-term memory networks for small-footprint keyword spotting,” in *Proc. SLT*, 2016, pp. 474–480.
- [2] Changhao Shan, Junbo Zhang, Yujun Wang, and Lei Xie, “Attention-based end-to-end models for small-footprint keyword spotting,” in *Proc. Interspeech*, 2018, pp. 2037–2041.
- [3] Alice Coucke, Mohammed Chlieh, Thibault Gisselbrecht, David Leroy, Mathieu Poumeyrol, and Thibaut Lavril, “Efficient keyword spotting using dilated convolutions and gating,” in *Proc. ICASSP*, 2019, pp. 6351–6355.
- [4] Jingyong Hou, Yangyang Shi, Mari Ostendorf, Mei-Yuh Hwang, and Lei Xie, “Mining effective negative training samples for keyword spotting,” in *Proc. ICASSP*, 2020, pp. 7444–7448.
- [5] Kun Zhang, Zhiyong Wu, Daode Yuan, Jian Luan, Jia Jia, Helen Meng, and Binheng Song, “Re-weighted interval loss for handling data imbalance problem of end-to-end keyword spotting,” in *Proc. Interspeech*, 2020, pp. 2567–2571.
- [6] Marius-Calin Silaghi, “Spotting subsequences matching an HMM using the average observation probability criteria with application to keyword spotting,” in *Proc. AAI*, 2005, pp. 1118–1123.
- [7] Guoguo Chen, Carolina Parada, and Georg Heigold, “Small-footprint keyword spotting using deep neural networks,” in *Proc. ICASSP*, 2014, pp. 4087–4091.
- [8] Raphael Tang and Jimmy Lin, “Deep residual learning for small-footprint keyword spotting,” in *Proc. ICASSP*, 2018, pp. 5484–5488.
- [9] Yiming Wang, Hang Lv, Daniel Povey, Lei Xie, and Sanjeev Khudanpur, “Wake word detection with alignment-free lattice-free MMI,” in *Proc. Interspeech*, 2020, pp. 4258–4262.
- [10] Yiming Wang, Hang Lv, Daniel Povey, Lei Xie, and Sanjeev Khudanpur, “Wake word detection with streaming transformers,” in *Proc. ICASSP*, 2021, pp. 5864–5868.
- [11] Ye Bai, Jiangyan Yi, Jianhua Tao, Zhengqi Wen, Zhengkun Tian, Chenghao Zhao, and Cunhang Fan, “A time delay neural network with shared weight self-attention for small-footprint keyword spotting,” in *Proc. Interspeech*, 2019, pp. 2190–2194.
- [12] Mengjun Zeng and Nanfeng Xiao, “Effective combination of densenet and bilstm for keyword spotting,” *IEEE Access*, vol. 7, pp. 10767–10775, 2019.
- [13] Ephrem Tibebe Mekonnen, Alessio Brutti, and Daniele Falavigna, “End-to-end low resource keyword spotting through character recognition and beam-search re-scoring,” in *Proc. ICASSP*, 2022, pp. 8182–8186.
- [14] Zhuoyuan Yao, Di Wu, Xiong Wang, Binbin Zhang, Fan Yu, Chao Yang, Zhendong Peng, Xiaoyu Chen, Lei Xie, and Xin Lei, “Wenet: Production oriented streaming and non-streaming end-to-end speech recognition toolkit,” *arXiv preprint arXiv:2102.01547*, 2021.
- [15] Daniel Povey, Arnab Ghoshal, Gilles Boulianne, Lukas Burget, Ondrej Glembek, Nagendra Goel, Mirko Hannemann, Petr Motlicek, Yanmin Qian, Petr Schwarz, et al., “The kaldı speech recognition toolkit,” in *Proc. ASRU*, 2011.
- [16] Myle Ott, Sergey Edunov, Alexei Baevski, Angela Fan, Sam Gross, Nathan Ng, David Grangier, and Michael Auli, “Fairseq: A fast, extensible toolkit for sequence modeling,” in *Proc. NAACL-HLT (Demonstrations)*, 2019.
- [17] Raphael Tang and Jimmy Lin, “Honk: A pytorch reimplementation of convolutional neural networks for keyword spotting,” *arXiv preprint arXiv:1710.06554*, 2017.
- [18] Daniel S Park, William Chan, Yu Zhang, Chung-Cheng Chiu, Barret Zoph, Ekin D Cubuk, and Quoc V Le, “SpecAugment: A simple data augmentation method for automatic speech recognition,” *arXiv preprint arXiv:1904.08779*, 2019.
- [19] Jingyong Hou, Lei Xie, and Shilei Zhang, “Two-stage streaming keyword detection and localization with multi-scale depthwise temporal convolution,” *Neural Networks*, vol. 150, pp. 28–42, 2022.
- [20] Alice Coucke, Mohammed Chlieh, Thibault Gisselbrecht, David Leroy, Mathieu Poumeyrol, and Thibaut Lavril, “Efficient keyword spotting using dilated convolutions and gating,” in *Proc. ICASSP*, 2019, pp. 6351–6355.
- [21] Bo Zhang, Wenfeng Li, Qingyuan Li, Weiji Zhuang, Xiangxiang Chu, and Yujun Wang, “AutoKWS: Keyword spotting with differentiable architecture search,” in *Proc. ICASSP*, 2021, pp. 2830–2834.
- [22] Kevin Ding, Martin Zong, Jiakui Li, and Baoxiang Li, “LETR: A lightweight and efficient transformer for keyword spotting,” in *Proc. ICASSP*, 2022, pp. 7987–7991.
- [23] Pete Warden, “Speech commands: A dataset for limited-vocabulary speech recognition,” *arXiv preprint arXiv:1804.03209*, 2018.
- [24] Shaojie Bai, J Zico Kolter, and Vladlen Koltun, “An empirical evaluation of generic convolutional and recurrent networks for sequence modeling,” *arXiv preprint arXiv:1803.01271*, 2018.
- [25] Simon Mittermaier, Ludwig Kürzinger, Bernd Waschneck, and Gerhard Rigoll, “Small-footprint keyword spotting on raw audio data with sinc-convolutions,” in *Proc. ICASSP*, 2020, pp. 7454–7458.