

FILTERBANK LEARNING FOR NOISE-ROBUST SMALL-FOOTPRINT KEYWORD SPOTTING

Iván López-Espejo^{1,2}, Ram C. M. C. Shekar², Zheng-Hua Tan¹, Jesper Jensen^{1,3}, John H. L. Hansen²

¹Department of Electronic Systems, Aalborg University, Denmark

²Center for Robust Speech Systems (CRSS), The University of Texas at Dallas, USA

³Oticon A/S, Denmark

{ivl, zt, jje}@es.aau.dk, {ramcharan.chandrashekhar, john.hansen}@utdallas.edu, jesj@demant.com

ABSTRACT

In the context of keyword spotting (KWS), the replacement of handcrafted speech features by learnable features has not yielded superior KWS performance. In this study, we demonstrate that filterbank learning outperforms handcrafted speech features for KWS whenever the number of filterbank channels is severely decreased. Reducing the number of channels might yield certain KWS performance drop, but also a substantial energy consumption reduction, which is key when deploying common always-on KWS on low-resource devices. Experimental results on a noisy version of the Google Speech Commands Dataset show that filterbank learning adapts to noise characteristics to provide a higher degree of robustness to noise, especially when dropout is integrated. Thus, switching from typically used 40-channel log-Mel features to 8-channel learned features leads to a relative KWS accuracy loss of only 3.5% while simultaneously achieving a $6.3 \times$ energy consumption reduction.

Index Terms— Keyword spotting, filterbank learning, small footprint, noise robustness, end-to-end

1. INTRODUCTION

Over recent years, there has been an increasing interest in developing end-to-end deep learning systems in which the feature extraction process is also optimized towards the task goal [1]. The keyword spotting (KWS) task, which deals with the recognition of a very limited vocabulary in speech signals, is, indeed, not alien to this trend [2–4]. In KWS, researchers seek to replace solid *handcrafted* speech features —e.g., log-Mel features and Mel-frequency cepstral coefficients (MFCCs)— by learnable features that are able to yield better KWS performance.

In spite of recent attempts [2–4], the above goal has not been achieved. In [2], Mittermaier *et al.* studied the

integration of a trainable filterbank —the so-called SincNet [5]— into a convolutional neural network (CNN)-based KWS pipeline. In that system, the cut-off frequencies of a sinc-convolution-based filterbank are trained jointly along with the CNN acoustic model. While a formal comparison between SincNet and handcrafted features is missing in [2], this analysis has recently been done in [4], where Peter *et al.* demonstrate that MFCCs remain superior for KWS.

In previous work [3], we found no statistically significant differences between using learned filterbanks and log-Mel features for KWS. From this observation, we conjectured that much of the spectral information is redundant when it comes to the recognition of a set of few keywords¹. Inspired by this, in this paper, we prove —for the first time, to the best of our knowledge— that filterbank learning outperforms handcrafted speech features for KWS as long as the number of filterbank channels is drastically reduced. While this might lead to certain KWS performance degradation, it can also yield a significant energy consumption and inference time reduction², which is of utmost importance when deploying typical always-on KWS on low-resource devices.

In this work, we conduct experiments on a noisy version of the Google Speech Commands Dataset [7, 8] by using a CNN-based KWS system integrating residual connections [9]. Experimental results prompt the following observations:

1. Filterbank learning adapts to noise spectral characteristics to offer a higher degree of robustness to noise;
2. The use of dropout [10] enhances robustness to noise and generalization capabilities of learned filterbanks;
3. Switching from 40-channel log-Mel features to 8-channel learned features yields a relative KWS accuracy loss of only 3.5% while involving a $6.3 \times$ energy consumption reduction;

¹Although in the context of reduced-precision speech features for KWS, this same hypothesis was independently made in [6].

²This is as a result of the number of multiplications in the acoustic model directly depending on the size of the input feature matrix.

This work has been funded by the European Union's Horizon 2021 research and innovation program under the Marie Skłodowska-Curie grant agreement No. 101062614.

- Switching from 8-channel log-Mel features to 5-channel learned features allows us for essentially maintaining KWS accuracy while leading to a $2\times$ energy consumption reduction.

The rest of the paper is organized as follows. Section 2 presents our filterbank learning methodology. The experimental setup is described in Section 3. Results are shown and discussed in Section 4, and Section 5 concludes this work.

2. FILTERBANK LEARNING METHODOLOGY

A block diagram of the KWS system employed in this work integrating filterbank learning can be seen in Fig. 1. First, the short-time Fourier transform (STFT) $X(t, f)$ of a discrete-time input speech signal $x(m)$, potentially comprising a keyword, is calculated. Note that $t = 1, 2, \dots, T$ and $f = 1, 2, \dots, F$ indicate, respectively, the time frame and linear frequency bin. Furthermore, T and F denote the total number of time frames and linear frequency bins, respectively, in the input speech signal. We can define a $T \times F$ matrix \mathbf{X} representing $x(m)$ in the linear power spectral domain as follows:

$$\mathbf{X} = \begin{bmatrix} |X(1, 1)|^2 & \cdots & |X(1, F)|^2 \\ \vdots & \ddots & \vdots \\ |X(T, 1)|^2 & \cdots & |X(T, F)|^2 \end{bmatrix}. \quad (1)$$

Next, the filterbank layer —optionally considering dropout [10] at training time— implements the matrix multiplication

$$\mathbf{Y} = \mathbf{X} \cdot g(\mathbf{W}), \quad (2)$$

where $\mathbf{W} \in \mathbb{R}^{F \times K}$ is the learnable K -channel filterbank matrix, which is optimized *jointly* along with the acoustic model (i.e., back-end) by backpropagation. Moreover, $g(\cdot) = \max(\cdot, 0)$ is the rectified linear unit function, which is element-wise applied to guarantee the positivity of the filterbank weights. The result of log-compressing the $T \times K$ filterbank representation \mathbf{Y} as in $\log(\max(\mathbf{Y}, \eta = e^{-50}))$, where the operators $\log(\cdot)$ and $\max(\cdot)$ are element-wise applied, is input to a batch normalization layer producing the speech features for the acoustic model.

Notice that the only difference between this filterbank learning scheme and the one that we already proposed in [3] is the optional consideration of dropout [10] at training time, the goal of which is to improve robustness and generalization of individual filterbank channels.

For acoustic modeling, we use a deep residual CNN integrating dilated convolutions [9], since in a recent overview article [1] we concluded that state-of-the-art KWS acoustic modeling relies on CNNs incorporating both residual connections and a scheme to seize long time-frequency patterns. This model produces word-level posteriors from one-second long input speech segments $x(m)$. A keyword is spotted every time that its associated posterior probability is the largest among all posteriors.

3. EXPERIMENTAL SETUP

3.1. Noisy Google Speech Commands Dataset

For experimental purposes, we employ the noisy version of the Google Speech Commands Dataset (GSCD) [7] that was introduced in [8]. This noisy GSCD defines training, validation and test sets composed of one-second long speech utterances —comprising one word each— that are contaminated by additive noises from the datasets NOISEX-92 [11] and CHiME-3 [12]. On the one hand, the types of noise present in the training and validation sets are vehicle interior, factory 1, bus, pedestrian street, white noise, babble, machine gun and F-16 cockpit. On the other hand, test set noise types are vehicle interior, factory 1, bus and pedestrian street (*seen noises*), as well as factory 2, Buccaneer jet cockpit, café and street junction (*unseen noises*). In addition, apart from the clean case also being considered, speech signals are contaminated at the following signal-to-noise ratios (SNRs): $\{0, 5, 10, 15, 20\}$ dB in the training and validation sets, and $\{-10, -5, 0, 5, 10, 15, 20\}$ dB in the test set. The training, validation and test sets contain 3,699, 427 and 497 utterances, respectively, from each combination of noise type and SNR level. The reader is referred to [8] for further details about this noisy version of the GSCD.

3.2. Configuration and System Training

3.2.1. Front-End

For STFT computation, an analysis window of 30-ms duration —corresponding to $N = 480$ samples at a 16 kHz sampling rate— with a 10-ms skip is employed. Hence, $F = (N/2) + 1 = 241$ is the total number of linear frequency bins. In addition, when dropout is considered, the dropout rate is 0.4. Departing from the standard number of filterbank channels in the KWS literature $K = 40$ [1, 3, 8, 13], different KWS systems with a reduced number of filterbank channels are trained to assess prospective benefits of filterbank learning. Note that \mathbf{W} is initialized by a Mel filterbank.

3.2.2. Back-End

The CNN back-end is trained to model 11 different classes: the 10 standard keywords of the GSCD “yes”, “no”, “up”, “down”, “left”, “right”, “on”, “off”, “stop” and “go” [7], plus the filler (i.e., non-keyword) class. As is typical in the context of GSCD, all classes are close to be balanced across the training, validation and test sets [1, 7].

The CNN acoustic model is trained by means of Adam [14] towards minimizing cross-entropy loss. The size of the mini-batch is 64 samples and early-stopping (monitoring the validation loss) with a patience of 5 epochs is used for regularization purposes. Finally, note that every KWS experiment is repeated 5 times by training 5 different acoustic models with

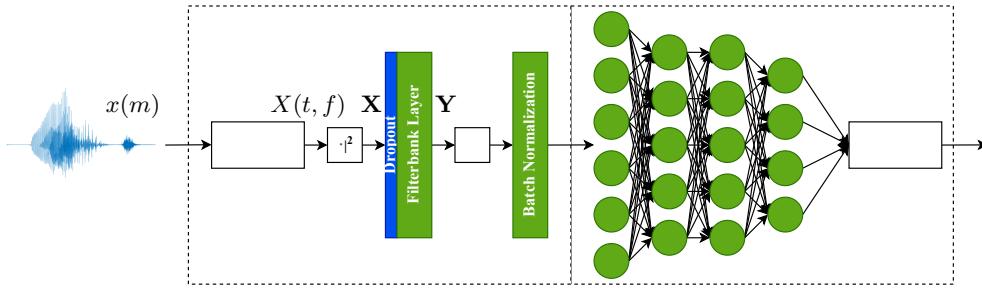


Fig. 1: Block diagram of the deep keyword spotting system employed in this work. See the text for further details.

different random initialization of their parameters in order to make sound conclusions from statistical tests.

4. RESULTS AND DISCUSSION

In this section, we compare the use of log-Mel features with the utilization of the learnable features described in Section 2 with —Learned+D— and without —Learned— dropout. Table 1 shows KWS accuracy results (averaged across 5 experiment repetitions), in percentages, as a function of the number of filterbank channels K . *Due to space constraints*, we only report results for a more relevant selection of K values (i.e., $\{40, 10, 8, 7, 5\}$). Statistically significant accuracy improvements are identified by means of two-sample t -tests [15] with a significance level of 0.05. In Table 1, statistically significant improvements with respect to log-Mel (Learned) are indicated in boldface (underlined).

Table 1 also shows the number of multiplications of the acoustic model per second of input speech, which highly depends on K due to the use of residual connections. Because *the number of multiplications exhibits a strong positive linear relationship with the energy consumption* of the KWS system, ($R^2 = 0.9641, p = 0.0001$) [16], we use the former as a proxy for the latter in this work.

Consistent with what we determined in [3], from Table 1, we see no statistically significant performance differences between log-Mel and learned features when employing a standard number of filterbank channels (i.e., $K = 40$). However, filterbank learning starts to become beneficial when $K < 10$, especially when dealing with unseen noise types. In particular, for $K = 5$, Learned+D outperforms log-Mel features in all the evaluated noisy conditions. We conjecture that the combination of dropout and a reduced K creates a bottleneck in the filterbank layer (quite similar to the ones observed in autoencoders), thereby providing a better knowledge representation for filterbank learning.

While it is true that KWS performance tends to drop when reducing K , filterbank learning cushions this drop while achieving noticeable energy savings. For example, switching from standard 40-channel log-Mel (avg. accuracy of 81.95%) to Learned+D when $K = 8$ (avg. accuracy of 79.11%) leads to a relative KWS accuracy loss of only 3.5%

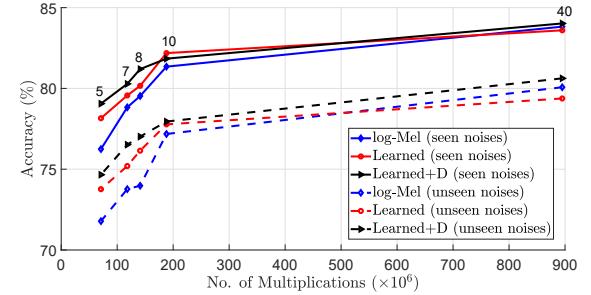


Fig. 2: Keyword spotting accuracy (%) averaged across SNRs, from using log-Mel and learned features, as a function of the number of multiplications of the acoustic model. The corresponding number of filterbank channels, K , is indicated above the curves.

while involving a $895M/141M \approx 6.3\times$ energy consumption reduction³. Furthermore, switching from 8-channel log-Mel (avg. accuracy of 76.75%) to Learned+D when $K = 5$ (avg. accuracy of 76.86%) allows us to essentially maintain KWS performance while yielding a reduction in energy consumption of a factor $141M/71M \approx 2$. To assess at a glance the more robust behavior of filterbank learning to the decreasing of K (for both seen and unseen noises), Fig. 2 plots KWS accuracy (averaged across SNRs) as a function of the number of multiplications of the acoustic model.

4.1. Filterbank Examination

Figure 3 depicts the average (across the 5 experiment repetitions) 5-channel learned filterbanks with (bottom) and without (top) dropout. As can be seen from this figure, the inclusion of dropout produces a smoother version of the filterbank learned without dropout, especially at lower frequencies.

It is interesting to note how learned filterbanks suppress frequency components near 2.7 kHz and 4.3 kHz. We pinpoint that this is due to the long-term spectrum of the F-16 cockpit training noise exhibiting strong peaks at those fre-

³For comparison, the relative KWS accuracy loss between 40-channel and 8-channel log-Mel (avg. accuracy of 76.75%) is of 6.3%.

Table 1: Average keyword spotting accuracy results (%), from using log-Mel and learned features, as a function of the number of filterbank channels K . Accuracy values are broken down by SNR as well as by seen and unseen noises during the training phase. The number of multiplications of the acoustic model, which depends on the number of filterbank channels, is also shown. Statistically significant improvements with respect to log-Mel (Learned) are indicated in boldface (underlined).

#Ch.	#Mult.	Method	SNR (dB) - Seen Noises							SNR (dB) - Unseen Noises								
			-10	-5	0	5	10	15	20	Clean	-10	-5	0	5	10	15	20	
40	895M	log-Mel	51.26	67.61	82.27	90.48	93.43	94.08	95.36	96.09	37.17	61.38	78.86	86.87	91.92	93.66	94.58	96.13
		Learned	49.61	66.93	82.46	90.58	93.96	94.23	95.41	95.58	35.41	58.84	77.94	86.82	92.24	93.47	94.70	95.60
		Learned+D	51.21	68.16	83.50	90.29	93.89	93.84	95.43	95.85	38.14	<u>63.29</u>	78.96	87.35	92.29	93.86	94.90	96.20
10	188M	log-Mel	45.41	63.09	79.30	87.85	91.93	92.90	94.59	95.70	32.00	55.67	74.34	83.94	90.23	92.29	93.91	95.07
		Learned	47.37	64.15	80.10	88.86	92.71	93.55	94.88	95.87	33.54	55.91	74.85	84.67	91.22	92.48	94.22	95.31
		Learned+D	46.72	63.57	80.39	88.36	92.13	93.29	94.71	95.58	32.50	57.39	75.14	84.57	90.67	93.23	94.39	95.72
8	141M	log-Mel	44.64	60.51	76.01	85.00	90.29	91.88	93.31	94.59	28.25	49.94	68.92	80.82	87.45	90.40	92.12	93.93
		Learned	45.70	62.22	77.97	86.06	90.75	91.45	92.97	94.15	29.65	53.11	72.24	84.04	90.33	92.16	93.40	94.24
		Learned+D	46.23	63.67	78.99	87.58	90.94	92.32	94.13	95.70	30.52	55.02	73.76	84.62	90.45	92.94	94.00	94.82
7	118M	log-Mel	42.90	60.72	76.28	84.49	89.18	90.31	92.51	94.23	27.88	48.20	68.51	80.53	87.06	90.40	92.72	<u>94.82</u>
		Learned	44.08	62.27	76.91	85.53	89.54	91.01	92.68	94.49	30.47	51.80	72.12	82.08	88.05	90.79	92.53	93.74
		Learned+D	45.53	61.35	78.00	87.00	90.80	91.43	93.26	94.88	31.32	54.68	73.30	83.94	89.50	92.09	92.99	94.32
5	71M	log-Mel	40.12	56.52	73.31	82.56	87.13	88.41	89.78	92.13	24.52	46.84	66.26	79.93	85.03	88.59	90.59	92.50
		Learned	42.92	59.40	75.34	84.69	88.26	89.57	91.52	93.55	28.34	50.47	70.11	80.82	86.99	89.41	91.22	92.74
		Learned+D	44.66	61.76	76.93	84.57	88.77	90.07	91.79	93.91	28.39	51.29	70.04	82.54	87.67	91.27	92.26	93.81

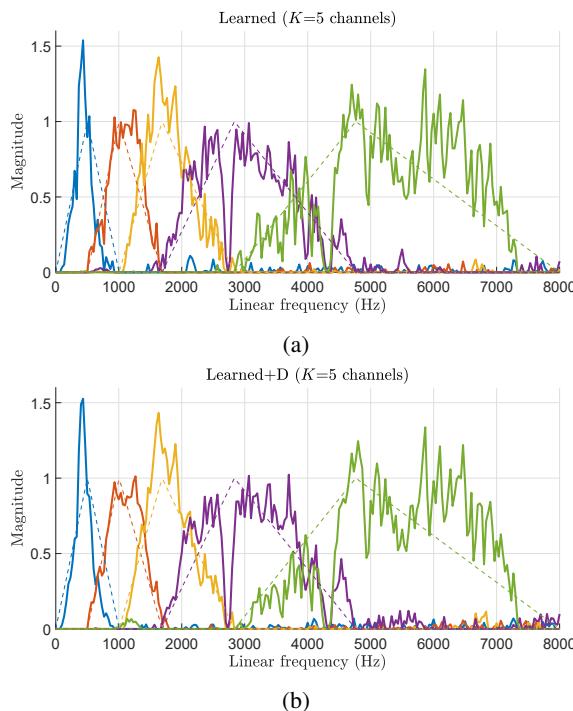


Fig. 3: Average (across the 5 experiment repetitions) learned filterbanks (solid lines) when considering $K = 5$ channels: (a) without dropout and (b) with dropout. For the sake of comparison, 5-channel Mel filterbanks (dashed lines) are overlaid.

quencies. In addition, the learned filterbanks seem to give more emphasis to the frequency range 5.8-7 kHz than the Mel filterbank (dashed lines in Fig. 3). This may be due to the fact that most of the training noise long-term spectra have fading tails at higher frequencies, so speech information potentially useful for KWS is less distorted in the 5.8-7 kHz range. In conclusion, filterbank learning seems to adapt to noise spectral characteristics to offer a higher degree of robustness to both seen and unseen noises.

It is remarkable that this analysis, illustrated by considering $K = 5$, holds also valid for $K > 5$, with similar plots to those in Fig. 3. Therefore, we believe that the potential advantages of filterbank learning adapting to noise characteristics are masked by the aforementioned information redundancy issue [3, 6] when $K \geq 10$ (see Table 1).

5. CONCLUDING REMARKS

For the first time, to the best of our knowledge, we have shown in this study that, when the number of filterbank channels is substantially decreased, filterbank learning is able to outperform handcrafted speech features for KWS. The entailed KWS performance-energy consumption trade-off is further compensated by using dropout for filterbank learning, which might be of particular interest to practitioners who deploy common always-on KWS on low-resource devices.

Future work envisages a comprehensive study on this matter aiming at obtaining deeper insight helping to design superior speech features for robust, small-footprint KWS.

6. REFERENCES

- [1] Iván López-Espejo, Zheng-Hua Tan, John H. L. Hansen, and Jesper Jensen, “Deep spoken keyword spotting: An overview,” *IEEE Access*, vol. 10, pp. 4169–4199, 2021.
- [2] Simon Mittermaier, Ludwig Kürzinger, Bernd Waschneck, and Gerhard Rigoll, “Small-footprint keyword spotting on raw audio data with sinc-convolutions,” in *Proceedings of ICASSP 2020 – 45th IEEE International Conference on Acoustics, Speech and Signal Processing, May 4-8, Barcelona, Spain*, 2020, pp. 7454–7458.
- [3] Iván López-Espejo, Zheng-Hua Tan, and Jesper Jensen, “Exploring filterbank learning for keyword spotting,” in *Proceedings of EUSIPCO 2020 – 28th European Signal Processing Conference, January 18-21, Amsterdam, Netherlands*, 2021, pp. 331–335.
- [4] David Peter, Wolfgang Roth, and Franz Pernkopf, “End-to-end keyword spotting using neural architecture search and quantization,” in *Proceedings of ICASSP 2022 – 47th IEEE International Conference on Acoustics, Speech and Signal Processing, May 23-27, Singapore, Singapore*, 2022, pp. 3423–3427.
- [5] Mirco Ravanelli and Yoshua Bengio, “Speaker recognition from raw waveform with SincNet,” in *Proceedings of SLT 2018 – IEEE Spoken Language Technology Workshop, December 18-21, Athens, Greece*, 2018, pp. 1021–1028.
- [6] Alexandre Riviello and Jean-Pierre David, “Binary speech features for keyword spotting tasks,” in *Proceedings of INTERSPEECH 2019 – 20th Annual Conference of the International Speech Communication Association, September 15-19, Graz, Austria*, 2019, pp. 3460–3464.
- [7] Pete Warden, “Speech Commands: A dataset for limited-vocabulary speech recognition,” *arXiv:1804.03209v1*, 2018.
- [8] Iván López-Espejo, Zheng-Hua Tan, and Jesper Jensen, “A novel loss function and training strategy for noise-robust keyword spotting,” *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, vol. 29, pp. 2254–2266, 2021.
- [9] Raphael Tang and Jimmy Lin, “Deep residual learning for small-footprint keyword spotting,” in *Proceedings of ICASSP 2018 – 43rd IEEE International Conference on Acoustics, Speech and Signal Processing, April 15-20, Calgary, Canada*, 2018, pp. 5484–5488.
- [10] Nitish Srivastava, Geoffrey Hinton, Alex Krizhevsky, Ilya Sutskever, and Ruslan Salakhutdinov, “Dropout: A simple way to prevent neural networks from overfitting,” *Journal of Machine Learning Research*, vol. 15, pp. 1929–1958, 2014.
- [11] Andrew Varga and Herman J.M. Steeneken, “Assessment for automatic speech recognition: II. NOISEX-92: A database and an experiment to study the effect of additive noise on speech recognition systems,” *Speech Communication*, vol. 12, pp. 247–251, 1993.
- [12] Jon Barker, Ricard Marxer, Emmanuel Vincent, and Shinji Watanabe, “The third ‘CHiME’ speech separation and recognition challenge: Analysis and outcomes,” *Computer Speech & Language*, vol. 46, pp. 605–626, 2017.
- [13] Byeonggeun Kim, Simyung Chang, Jinkyu Lee, and Dooyong Sung, “Broadcasted residual learning for efficient keyword spotting,” in *Proceedings of INTERSPEECH 2021 – 22nd Annual Conference of the International Speech Communication Association, August 30-September 3, Brno, Czechia*, 2021, pp. 4538–4542.
- [14] Diederik P. Kingma and Jimmy Ba, “Adam: A method for stochastic optimization,” in *Proceedings of ICLR 2015 – 3rd International Conference on Learning Representations, May 7-9, San Diego, USA*, 2015.
- [15] George W. Snedecor and William G. Cochran, *Statistical Methods (8th edition)*, Iowa State University Press, 1989.
- [16] Raphael Tang, Weijie Wang, Zhucheng Tu, and Jimmy Lin, “An experimental analysis of the power consumption of convolutional neural networks for keyword spotting,” in *Proceedings of ICASSP 2018 – 43rd IEEE International Conference on Acoustics, Speech and Signal Processing, April 15-20, Calgary, Canada*, 2018, pp. 5479–5483.