



# PEAF: Learnable Power Efficient Analog Acoustic Features for Audio Recognition

Boris Bergsma<sup>†\*</sup> Minhao Yang<sup>◇</sup> Milos Cernak<sup>†</sup>

<sup>\*</sup> École Polytechnique Fédérale de Lausanne (EPFL), Switzerland

<sup>†</sup> Logitech Europe S.A., Lausanne, Switzerland

<sup>◇</sup> Reexen Technology, Zürich, Switzerland

bergsma.boris@gmail.com, yang@reexen.com, milos.cernak@ieee.org

## Abstract

At the end of Moore’s law, new computing paradigms are required to prolong the battery life of wearable and IoT smart audio devices. Theoretical analysis and physical validation have shown that analog signal processing (ASP) can be more power-efficient than its digital counterpart in the realm of low-to-medium signal-to-noise ratio applications. In addition, ASP allows a direct interface with an analog microphone without a power-hungry analog-to-digital converter. Here, we present power-efficient analog acoustic features (PEAF) that are validated by fabricated CMOS chips for running audio recognition. Linear, non-linear, and learnable PEAF variants are evaluated on two speech processing tasks that are demanded in many battery-operated devices: wake word detection (WWD) and keyword spotting (KWS). Compared to digital acoustic features, higher power efficiency with competitive classification accuracy can be obtained. A novel theoretical framework based on information theory is established to analyze the information flow in each individual stage of the feature extraction pipeline. The analysis identifies the information bottleneck and helps improve the KWS accuracy by up to 7%. This work may pave the way to building more power-efficient smart audio devices with best-in-class inference performance.

**Index Terms:** Analog systems, audio classification, power efficiency, information theory

## 1. Introduction

The emerging and expanding wearable and IoT market requires extremely power-efficient electronics for prolonged continuous operation with tiny batteries and/or energy harvesting. Audio devices like wireless stereo earbuds and active noise canceling earbuds, and other wearables like watches and glasses, can benefit from smart acoustic interfaces for gesture-free voice control. Such interface typically comprises always-on functionalities such as voice activity detection (VAD), wake-word detection (WWD), keyword spotting (KWS), etc. As always-on functionalities, the power efficiency of their hardware implementation dramatically affects the overall system battery life, especially when the power-hungry blocks triggered by them are heavily duty-cycled.

With 3nm process looking to be in mass production in 2022/2023, the 50 years of silicon scaling will end in this decade. Not only the active power reduction associated with transistor feature size shrinkage is soon finished, but also the severe leakage problem in advanced technology nodes largely compromises the power efficiency. Other than the dominant digital signal processing (DSP), new computing paradigms are being actively explored, and analog signal processing (ASP)

is one promising alternative. The power advantage of ASP is twofold. Firstly, theoretical analysis [1] and experimental validation [2, 3] show that ASP is more power efficient than its digital counterpart [4, 5] implementing algorithms like Mel-frequency cepstral coefficient (MFCC) when low-to-medium signal-to-noise ratio processing is sufficient. Secondly, ASP can directly interface with an analog MEMS microphone [2, 3], avoiding a power-consuming high-precision analog-to-digital converter (ADC) [6]. Recent chip implementations of ASP-based acoustic feature extraction focusing on the VAD functionality [2, 3] have showed the extreme low power consumption from nanowatts to microwatts with inference accuracy comparable to that of MFCC-based features.

In this work, linear [2] (L-PEAF), non-linear [3] (N-PEAF), and proposed learnable power-efficient analog acoustic features (Learn-PEAF) are evaluated on WWD and KWS, which are highly sought-after functionalities in many wearable and IoT devices. Given that the PEAF variants have spike output in the feature extraction pipeline, spiking neural networks [7, 8] seem to be the natural candidate as the classifier. But we use artificial neural networks because of the state-of-the-art energy efficiency in their hardware implementations. Inspired by a learnable audio frontend (LEAF) [9] intended for improving DSP-based feature extraction, we devise a novel analysis method based on information theory and use it for the learnable PEAF design. Our experimental validation on both WWD and KWS shows improved inference performance with appreciable power reduction. Fig. 1. shows all the considered features in this work.

## 2. Methods

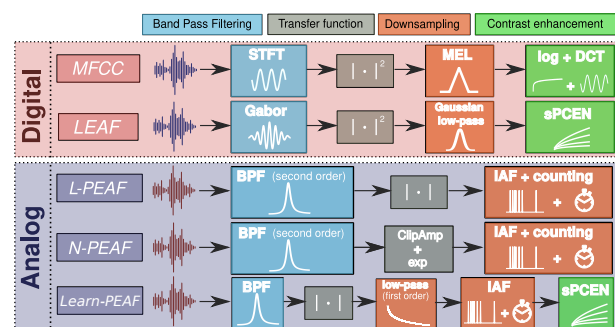


Figure 1: Scheme of all the compared features in this paper. STFT : Short Term Fourier Transform, DCT: Discrete Cosine Transform, PCEN: Per Channel Energy Normalisation, BPF: BandPass Filters, IAF: Integrate And Fire encoder, L-: Linear, N-: Non-linear; PEAF: Power Efficient Analog Features.

Digital handcrafted (MFCCs) and learnable (LEAFs) features are compared to linear (L-PEAF), non-linear (N-PEAF), and learnable (Learn-PEAF) analog features.

### 2.1. (Reference) digital features

Due to our ambition to compare the performance and power efficiency of features on fabricated chips, we have chosen MFCC reference implementation [5] with twenty Mel-frequency bands and ten coefficients. LEAF is a method that outperforms MFCC in multiple tasks, but its neural network based implementation [9] has not provided power measurements and we thus excluded it from comparison. Instead, we used the LEAF representation in information theory-based analysis (Sections 2.5 and 3.1) to identify suboptimal stages in analog feature computation.

### 2.2. Analog features

Fig. 1 shows the schema of analog feature computation that follows the human cochlea processing. The computation comprises three basic operations: 2nd-order bandpass filtering, an activation function, and spike generation based on Integrate-and-Fire (IAF) encoding. The IAF acts as a biological neuron, accumulating a certain amount of electrical charge and releasing a voltage spike after reaching a certain threshold. We call the analog feature extraction with an absolute value activation function “L-PEAF”, given its piece-wise linear nature and the linearity requirement on circuit implementation [2]. On the contrary, a version with a clipped exponential activation function is tagged “N-PEAF”, which can be implemented with a limiting amplifier and a single voltage-to-current conversion transistor [3]. The output of the IAF encoder can be directly used for spiking neural networks [7] or counted in frames and used as a 2D feature like the MFCC.

Learnable PEAF, named “Learn-PEAF”, is our proposed improvement to the actual L-PEAF schema. It has an additional first-order low pass filter (per channel) before the IAF and a PCEN step before entering the digital classifier. All the PEAF features have built-in and fixed parameters, for example, the center frequency of the bandpass filters. The proposed Learn-PEAF feature schema was simulated using PyTorch. This allows to automatically differentiate the PEAF frontend, and the gradient can be back propagated through the classifier and frontend, allowing to optimise PEAF’s parameters.

### 2.3. Digital deep learning classification

We evaluate the different features on multiple classifiers for scaling analysis and robust results. The employed classifiers have different neural network architectures with growing classification capacities: **A)** a small version of EfficientNetV2 [10] (Fig. 2), **B)** LeNet-5 [11], **C)** a depthwise separable CNN (DS-CNN) [12] and **D)** EfficientNetB0 [13]. All the classifiers are trained using Adam optimization [14] with default learning rate, and SpecAugment [15] is used as data augmentation. A scheduler reduces the learning rate geometrically by e-0.01 after 100 epochs for the small neural networks (**A,B**) and 50 for the large ones (**C,D**).

### 2.4. Power consumption estimation

The power consumption of the L-PEAF and the N-PEAF can be estimated from the fabricated chips [2, 3]. The power consumption of the digital MFCC features can be estimated from [5], plus an ADC’s contribution [6]. The power consumption of the

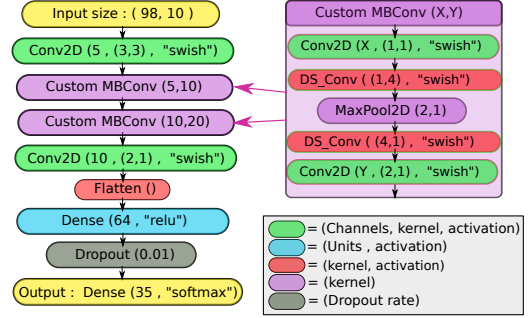


Figure 2: Scheme of the Mini-EfficientNet classifier (A).

Learn-PEAF is obtained by estimating that the 1st-order low-pass filter consumes the same as the bandpass filters (16% of the L-PEAF [2]), and a possible simple DSP-based PCEN adds another 10%. The power consumption estimates are summarized in Table 1.

Table 1: Power consumption of the different features.

Feature	Power
MFCC [5] + ADC [6]	$0.34 + 7.2 = 7.5 \mu\text{W}$
N-PEAF [3]	$0.072 \mu\text{W}$
L-PEAF [2]	$0.38 \mu\text{W}$

To estimate a lower bound of the power consumption of the classifiers, the energy efficiency from one state-of-the-art neural network processor  $E_{eff} = 36.5 \cdot 10^{12}$  OPS/W is used [16]. Table 2 summarizes the power consumption of the considered classifiers. The classifiers’ power is calculated by equation 1, where  $N_{OPS}$  is the number of operations done by the neural network model,  $FR$  is the frame rate. We choose  $FR = 10 \text{ fps}$  for WWD and  $FR = 30 \text{ fps}$  for KWS.

$$P_{tot} = P_{feat} + P_{class} = P_{feat} + \frac{N_{OPS} \times FR}{E_{eff}} \quad (1)$$

Table 2: Power consumption estimates of the classifiers for the MFCC with 10 coefficients.

Classifier	KWS Power	WWD Power
<b>A)</b> Mini-EfficientNet (Fig. 2)	$0.097 \mu\text{W}$	$0.079 \mu\text{W}$
<b>B)</b> LeNet-5 [11]	$0.34 \mu\text{W}$	$0.30 \mu\text{W}$
<b>C)</b> DS-CNN [12]	$8.3 \mu\text{W}$	$5.4 \mu\text{W}$
<b>D)</b> EfficientNet [13]	$48 \mu\text{W}$	-

### 2.5. Information theory

Information theory is used to compare the different features to identify the bottleneck operation in the feature computation. The tool used is the Shannon Entropy [17]:

$$S_{Shannon} = - \sum_{i=0}^N p_i \log(p_i) \quad (2)$$

with  $p_i$  the probability of the data distribution. The estimation of the data distribution in the features can be done in multiple ways. To retain the spatial information of the features this

work encodes spatial information by giving each data point a unique identifier (an integer) in a second dimension. This two-dimensional object is defined as :

$$\text{Encoded feature} = (\text{Flattenedfeature}^T, \text{linspace}(1, N)^T) \quad (3)$$

with “data-point value“ and “spatial label“ dimensions. This two-dimensional object is then fitted with a two-dimensional histogram to obtain an estimation of the  $p_i$ 's. This method should be seen as an approximation, as there is currently no universally acceptable way to extract entropy for matrices without probability priors. The proposed method of histogram computation also requires more research; some initial exploration is presented in [18].

### 3. Experimental Results

In all our experiments, digital MFCCs with various digital classifiers and known overall power consumption was the baseline. We evaluated analog PEAf with the same digital classifiers.

#### 3.1. Feature analysis and PEAf validation

First, we applied information theory to gain an insight into information flow across feature computation. We hypothesized that comparing the information flows from the MFCCs, LEAFs, and PEAfs could identify “information bottleneck“, and inform the learnable PEAf which steps could benefit from data-driven optimization.

A thousand random samples from the Speech Command dataset V2 [19] with equal class repartition are used for this analysis. The features are separated into their computation steps, and the entropy is computed as described in section 2.5. The average of the entropy on all the samples is then plotted with its standard deviation in Fig.3.

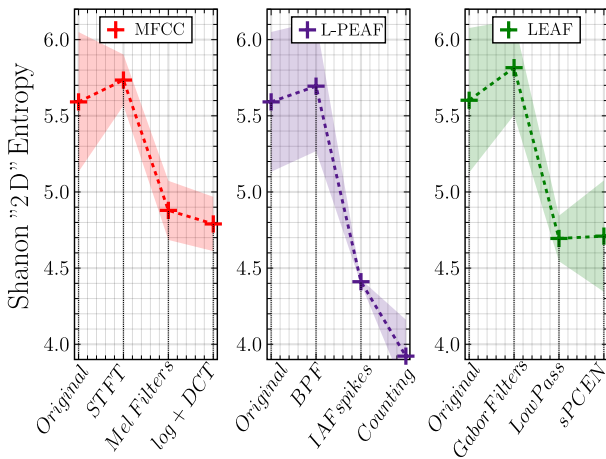


Figure 3: Comparison of the features at each computation step using information theory on KWS.

The obvious conclusion that can be made out of Fig.3 that the last two steps of the computation of analog features (L-PEAF) are the information bottleneck (more details in sec.4). We tested adjustments in the analog feature computation by replacing them with LEAF's counterparts to validate it. The IAF was replaced with the Gaussian low pass filters of LEAF. Then, we added PCEN as the last step of the computation. The results in Tab. 3 confirm that Learn-PEAFs outperform L-PEAF

with the fixed parameters and also MFCCs. We validated the design of Learn-PEAFs and used the novel proposed features in consecutive experiments.

Table 3: Validation using Learn-PEAF with LeNet (B) classifier on KWS. For every row, one computation step of L-PEAF was replaced sequentially by the LEAF equivalent.

Parameters :	Fixed	Optimized
L-PEAF (baseline)	79.1%	-
IAF $\rightarrow$ Gaussian Low Pass	81.7%	82.8%
Adding sPCEN	82.6%	<b>85.1%</b>
Learn-PEAF	-	<b>86.4%</b>
MFCC	84.8%	-

#### 3.2. Feature comparison for KWS task

We compared digital MFCCs, and analog PEAf variants on the Speech Command dataset V2 [19], with all the 35 different classes and no added noise. Fig. 4 summarizes obtained results.

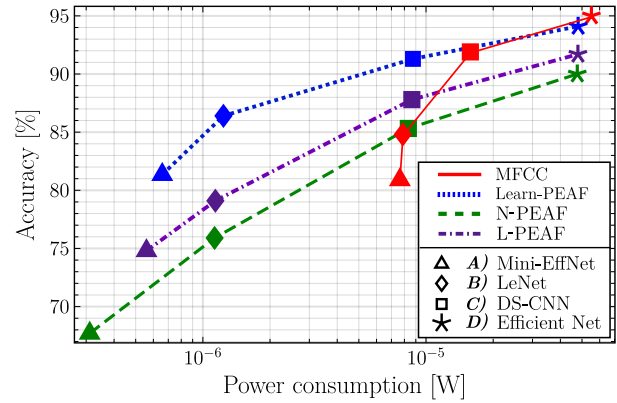


Figure 4: Features comparison on KWS.

#### 3.3. Feature comparison for WWD

Evaluating PEAf performance on the second task is two-fold. First, WWD runs on many battery-operated devices, and its power efficiency is critical. Second, Learn-PEAF was fit on the KWS data, and we were interested in evaluating their cross-dataset performance on the different, wake word data.

For WWD, we adopted the open-source experimental setup Howl from Mozilla [20] including the code provided in [21]. The dataset consisted of multiple positive wake words “Hey Firefox“ and Common Voice dataset [22] prompts as negative background data. Three hundred fifty actual recordings of “Hey Firefox“ were also provided. The noise was added from the MS-Noise dataset [23] at 10dB SNR. The final dataset (training/validation sets) contained 10% positive examples (2700 generated and 300 actual recordings) with a total of 30,000 audio samples of 2.5 seconds. The testing set was composed of 50 actual recordings of “Hey Firefox“, 500 positive generated recordings given by [21], 3500 background noise, and 3500 negative speech samples, both randomly shifted in time.

Fig. 5 shows the ROC curves of different setups. Howl's performance was evaluated using the pre-trained model (*Res-8 small*) provided by the author [21]. Fig. 6 shows the accuracy

vs. power consumption. To get an accuracy value the false alarm rate was set to 4 as done in Howl.

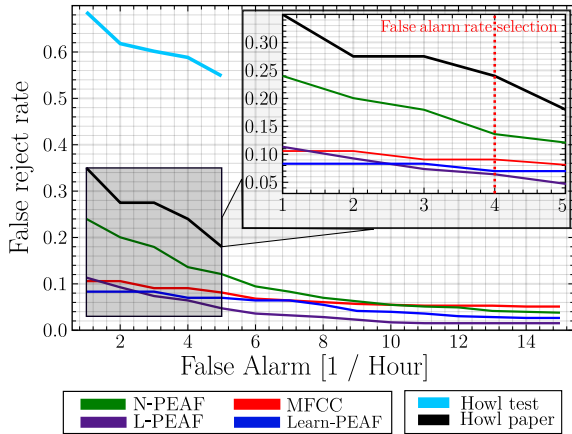


Figure 5: ROC curves for Mini-EfficientNet (A) on WWD.

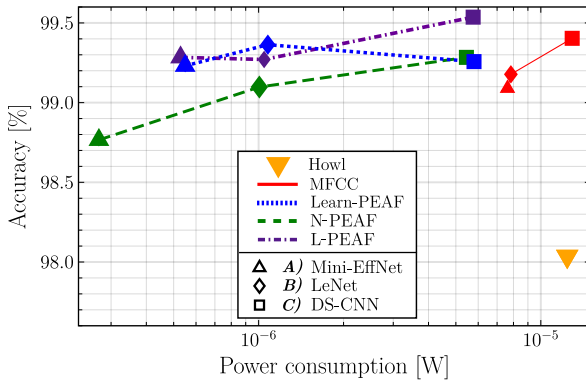


Figure 6: Features comparison for WWD.

## 4. Discussion

Analyzing Fig. 3 allows understanding of information bottlenecks. The pattern is similar for all the features; the first step extracts as much information as possible. The dimensionality of the features is then reduced, inducing a reduction in information. The MFCC loses a bit of information in the last step. Indeed there are twenty Mel-Frequency bands and only ten final coefficients after DCT. LEAF sees its information increase in mean and variance with the use of PCEN. The output entropy of the MFCC is higher than LEAF, with LEAF having a higher variance. Knowing that the MFCC and LEAF have the same performance on KWS [9], it could be a testament to how the variance in entropy is primordial or an issue with our experimental, theoretical framework. The analog features lose the most information with the spikes computation in the IAF and when counting the spikes. Table 3 shows that all the tested replacements increase the accuracy by a significant amount. Optimizing the parameters of the BPF in the original analog extractor allows for significant and easy to implement improvement. The fully improved Learn-PEAF gains 7.3% in accuracy, surpassing the MFCC. The other comparison points added on Fig. 4 and 6 also show a clear improvement. The increase is clear

on KWS and WWD with LeNet (B). Overall this information theory method allowed to get insight into the features and target the computation steps that needed change. A more in depth analysis is required to make it a stable tool.

The results on keyword spotting (Fig.4) indicate that L-PEAF performs worse by about 2 to 5% for all classifiers compared to the digital features. For large classifiers (C,D), the energy consumption of the classifier dominates the total power consumption, and there is no reason to choose L-PEAF or N-PEAF over MFCCs. The accuracy gap increases with the smaller neural network (A,B) for L-PEAF and N-PEAF. Power consumption is around seven times better for all PEAf with LeNet (B) and more than thirteen times greater for the Mini-EfficientNet (A). Learn-PEAF seems to favor the small classifiers (A,B) with the improvement of more than 5 % accuracy compared to L-PEAF. The larger neural networks (C,D) show only a 2 % accuracy improvement. Overall, Learn-PEAF outperforms the MFCC in the low power domain on both axes.

Fig.5 shows that L-PEAF has the overall best ROC curves, followed by Learn-PEAF closely. Evaluated Howl’s pre-trained system demonstrated a non-competitive ROC curve. A difference in pre-processing could explain it. The weights were used as-is and not retrained. The reported Howl ROC curves are in the same magnitude as our tested results. Apart from the results from N-PEAF, all the ROC curves are close. The accuracy results on Fig.6 demonstrate a potential usage of L-PEAF and Learn-PEAF, with results better than the MFCC. The gain is also evident in energy efficiency for all PEAfs, with a clear energy gap between analog and digital features. N-PEAF could be an alternative for the most power efficient WWD application consuming two times less energy compared to L-PEAF with a 1 % accuracy loss. Learn-PEAF outperform all the other features only with LeNet (B). As it was trained with this classifier on keyword spotting, it is interesting to note that it can transfer some of its knowledge to another task. It could be improved further by training its parameters on the WWD task.

One limitation of our comparison is that Learn-PEAF uses a simplified version of the IAF encoder. It was done to increase computing time dramatically and allow the training of the extractor’s parameters. Both L-PEAF and N-PEAF are also simulated numerically, no recording of the real analog chips were taken. This was done in [2, 3]; though only for VAD, the simulation proved to be accurate enough.

## 5. Conclusions

This paper has presented linear, non-linear, and learnable power-efficient acoustic analog features, and evaluated them on two speech processing tasks demanded in many battery-operated devices: WWD and KWS. The WWD task showed that the analog features outperform the digital ones, with an order of magnitude lower power consumption and competitive classification accuracies.

We argue that the novel information theory-based method introduced in this paper could also be applied for other feature analysis. We have used it to identify the information bottleneck in L-PEAF calculation and, inspired by LEAF, devised a novel learn-PEAF. Learn-PEAF improved the KWS accuracy up to 7% compared to L-PEAF, outperforming MFCCs with less power consumption.

Future works include adapting the processing in the modified stages to feasible integrated electronics implementation and exploring the possibility of directly using spiking neural networks as the classifier after IAF for accuracy improvement.

## 6. References

- [1] B. J. Hosticka, "Performance comparison of analog and digital circuits," *Proceedings of the IEEE*, vol. 73, no. 1, pp. 25–29, 1985.
- [2] M. Yang, C.-H. Yeh, Y. Zhou, J. P. Cerqueira, A. A. Lazar, and M. Seok, "A  $1\mu\text{W}$  voice activity detector using analog feature extraction and digital deep neural network," in *2018 IEEE International Solid-State Circuits Conference (ISSCC)*, 2018, pp. 346–347.
- [3] M. Yang, H. Liu, W. Shan, J. Zhang, I. Kiselev, S. J. Kim, C. Enz, and M. Seok, "Nanowatt Acoustic Inference Sensing Exploiting Nonlinear Analog Feature Extraction," *IEEE Journal of Solid-State Circuits*, vol. 56, no. 10, pp. 3123–3133, 2021.
- [4] A. Raychowdhury, C. Tokunaga, W. Beltman, M. Deisher, J. W. Tschanz, and V. De, "A 2.3 nJ/frame voice activity detector-based audio front-end for context-aware system-on-chip applications in 32-nm cmos," *IEEE journal of solid-state circuits*, vol. 48, no. 8, pp. 1963–1969, 2013.
- [5] W. Shan, M. Yang, T. Wang, Y. Lu, H. Cai, L. Zhu, J. Xu, C. Wu, L. Shi, and J. Yang, "A 510-nw wake-up keyword-spotting chip using serial-fft-based mfcc and binarized depthwise separable cnn in 28-nm cmos," *IEEE Journal of Solid-State Circuits*, vol. 56, no. 1, pp. 151–164, 2021.
- [6] H. Chandrakumar and D. Marković, "A 15.2-ENOB 5-kHz BW  $4.5\text{-}\mu\text{W}$  Chopped CT  $\Delta\Sigma$ -ADC for Artifact-Tolerant Neural Recording Front Ends," *IEEE Journal of Solid-State Circuits*, vol. 53, no. 12, pp. 3470–3483, 2018.
- [7] G. Dellaferera, F. Martinelli, and M. Cernak, "A Bin Encoding Training of a Spiking Neural Network Based Voice Activity Detection," in *ICASSP 2020 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, May 2020, pp. 3207–3211.
- [8] F. Martinelli, G. Dellaferera, P. Mainar, and M. Cernak, "Spiking neural networks trained with backpropagation for low power neuromorphic implementation of voice activity detection," in *ICASSP 2020 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, May 2020, pp. 8544–8548.
- [9] N. Zeghidour, O. Teboul, F. de Chaumont Quitry, and M. Tagliasacchi, "Leaf: A learnable frontend for audio classification," 2021.
- [10] M. Tan and Q. V. Le, "Efficientnetv2: Smaller models and faster training," *CoRR*, vol. abs/2104.00298, 2021. [Online]. Available: <https://arxiv.org/abs/2104.00298>
- [11] Y. Lecun, L. Bottou, Y. Bengio, and P. Haffner, "Gradient-based learning applied to document recognition," *Proceedings of the IEEE*, vol. 86, no. 11, pp. 2278–2324, 1998.
- [12] A. Howard, M. Sandler, G. Chu, L.-C. Chen, B. Chen, M. Tan, W. Wang, Y. Zhu, R. Pang, V. Vasudevan, Q. V. Le, and H. Adam, "Searching for MobileNetV3," in *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 2019, pp. 1314–1324.
- [13] M. Tan and Q. Le, "Efficientnet: Rethinking model scaling for convolutional neural networks," in *International conference on machine learning*. PMLR, 2019, pp. 6105–6114.
- [14] D. P. Kingma and J. Ba, "Adam: A method for stochastic optimization," *arXiv preprint arXiv:1412.6980*, 2014.
- [15] 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," in *Proc. Interspeech 2019*, 2019, pp. 2613–2617.
- [16] F. Tu, Y. Wang, Z. Wu, L. Liang, Y. Ding, B. Kim, L. Liu, S. Wei, Y. Xie, and S. Yin, "A 28nm 29.2 TFLOPS/W BF16 and 36.5 TOPS/W INT8 Reconfigurable Digital CIM Processor with Unified FP/INT Pipeline and Bitwise In-Memory Booth Multiplication for Cloud Deep Learning Acceleration," in *2022 IEEE International Solid-State Circuits Conference (ISSCC)*, vol. 65. IEEE, 2022, pp. 1–3.
- [17] C. E. Shannon, "A mathematical theory of communication," *The Bell system technical journal*, vol. 27, no. 3, pp. 379–423, 1948.
- [18] B. Bergsma, "Power efficient acoustic feature representation for IoTs," Master's thesis, École Polytechnique Fédérale de Lausanne, 2021.
- [19] P. Warden, "Speech Commands: A Dataset for Limited-Vocabulary Speech Recognition," *arXiv preprint arXiv:1804.03209*, 2018.
- [20] R. Tang, J. Lee, A. Razi, J. Cambre, I. Bicking, J. Kaye, and J. Lin, "Howl: A deployed, open-source wake word detection system," in *Proceedings of Second Workshop for NLP Open Source Software (NLP-OSS)*. Association for Computational Linguistics, Nov. 2020, pp. 61–65. [Online]. Available: <https://www.aclweb.org/anthology/2020.nlposs-1.9>
- [21] —, "The github page of howl." <https://github.com/castorini/howl>.
- [22] R. Ardila, M. Branson, K. Davis, M. Kohler, J. Meyer, M. Henretty, R. Morais, L. Saunders, F. Tyers, and G. Weber, "Common voice: A massively-multilingual speech corpus," in *Proceedings of the 12th Language Resources and Evaluation Conference*. Marseille, France: European Language Resources Association, May 2020, pp. 4218–4222. [Online]. Available: <https://aclanthology.org/2020.lrec-1.520>
- [23] C. K. Reddy, E. Beyrami, J. Pool, R. Cutler, S. Srinivasan, and J. Gehrke, "A Scalable Noisy Speech Dataset and Online Subjective Test Framework," in *Proc. Interspeech 2019*, 2019, pp. 1816–1820.