



# Spectrum Correction: Acoustic Scene Classification with Mismatched Recording Devices

Michał Kośmider

Samsung R&D Institute Poland

m.kosmider@samsung.com

## Abstract

Machine learning algorithms, when trained on audio recordings from a limited set of devices, may not generalize well to samples recorded using other devices with different frequency responses. In this work, a relatively straightforward method is introduced to address this problem. Two variants of the approach are presented. First requires aligned examples from multiple devices, the second approach alleviates this requirement. This method works for both time and frequency domain representations of audio recordings. Further, a relation to standardization and Cepstral Mean Subtraction is analysed. The proposed approach becomes effective even when very few examples are provided. This method was developed during the *Detection and Classification of Acoustic Scenes and Events* (DCASE) 2019 challenge and won the 1st place in the scenario with mismatched recording devices with the accuracy of 75%. Source code for the experiments can be found online.

**Index Terms:** channel normalization, acoustic scene classification, mismatched recording devices, convolutional neural network

## 1. Introduction

Audio machine learning has many sub-fields. Music and speech processing are well established, while other such as event detection, acoustic scene classification (ASC) or behaviour classification are less common. Moreover, the more specific or unusual the task, the harder it is to find publicly available datasets.

When a dataset of audio recordings for a specific purpose is collected, it is usually done with some selected set of recording devices [1, 2]. This is sufficient to test the performance of various algorithms, but might not be enough for practical applications. Models trained using such dataset, when evaluated on samples from the same devices, appear to work well. However, their performance deteriorates when applied on samples from devices that were not in the training set or were under-represented. This suggests that machine learning algorithms do not generalize across samples recorded using different devices [3, 4, 5, 6, 7]. In extreme cases, this forces collection of completely new datasets for the targeted devices. Severity of the problem varies between tasks. This effect is clearly visible in acoustic scene classification.

The problem of device mismatch was popularized in the *Detection and Classification of Acoustic Scenes and Events* (DCASE) challenge. In 2018, a task targeting this problem became part of the challenge [2] and has been continued in 2019, which inspired research in this area resulting in various publications [5, 6, 8]. Some of the solutions were based on domain adaptation [9], aggressive regularization [9, 10] and most were model ensembles [11].

Independently from DCASE challenge, an approach was proposed to use transfer learning [4]. The idea was to firstly

train the model using a relatively large dataset and then fine tune it using a smaller dataset containing recordings from the targeted devices.

There currently does not appear to be any general approach to address this problem for ASC. Most solutions focus on tailoring specific models to become less sensitive to characteristics of the recording device. On the other hand, in the speech processing community there are multiple well known methods for channel normalization [12, 3, 13]. However a lot of those methods, such as RASTA [14], are leveraging specific characteristics of speech. Others such as Cepstral Mean Subtraction (CMS) [15], Feature Warping [16] or more recent MVA [17] are relying on quefrequency domain, which is not usually applied for ASC and similar tasks. For frequency domain, there exists a recently introduced Per-Channel Energy Normalization (PCEN) [18, 19], which additionally reduces noise. However, noise reduction may not be desirable for ASC, where we would like to retain as much information about the environment as possible. Noise and reverberation are likely to contain information that differentiates various environments. Because of these considerations, a more universal method could be beneficial.

In this work, a method addressing this problem is described. The approach is straightforward and relies on real-world recordings. It does not influence the design of the architecture of the neural network. This method was designed for ASC, but can be applied to other tasks as well. An initial test of this method was performed during the DCASE 2019 challenge. A submission utilizing this approach won the 1<sup>st</sup> place in the scenario with mismatched recording devices with the accuracy of 75% [20].

From the acoustic point of view, the presented method appears to bear a similarity to standard microphone calibration techniques [21], for example to the secondary calibration of a microphone by the comparison. This method relies on a direct comparison of a microphone with a physically identical microphone having a known calibration and is repeated at each frequency of interest. These methods assume laboratory settings and require expert knowledge, which makes them hard to apply for practitioners without a background in signal processing. Databases of recordings are sometimes provided without explicitly defining the types of devices used in their creation. Also, sometimes the devices are not easily available. The proposed method is simpler, more approachable and could help with the problem of mismatched recording devices.

## 2. Proposed Method

### 2.1. Linear Distortion Model

Throughout this paper a linear distortion model will be assumed. Also all amplitudes and responses are positive and real. It can be described in the following way.

Let  $w_s$  be the undistorted waveform representation of some audio signal  $s$ . Let  $\tilde{w}_{s,d}$  denote the distorted representation of

the same signal produced by some device  $d$ . The signal  $w_s$  can be analysed in terms of its amplitude spectrum. Let  $A_s(t, f)$  denote the amplitude of the frequency  $f$  at time  $t$  in that signal. Similarly, we define  $\tilde{A}_{s,d}(t, f)$  for signal  $\tilde{w}_{s,d}$ .

The assumption is that the distortion introduced by the device can be approximated using the frequency response of that device  $H_d(f)$  as:

$$\tilde{A}_{s,d}(t, f) \approx H_d(f) A_s(t, f) \quad (1)$$

$$= H_d(f) H_s(f) A'_s(t, f), \quad (2)$$

where  $H_s(f)$  is the frequency response specific to the signal  $s$  and independent from the recording device, for example a room impulse response. The amplitude  $A'_s(t, f)$  is a remainder after decomposing  $A_s(t, f)$  into  $H_s(f)$  and everything else. For most tasks, we would want to eliminate both  $H_d(f)$  and  $H_s(f)$ , but for ASC preserving  $H_s(f)$  could be desirable, because it is likely to be specific to the environment.

This model is naïve, because it does not take into account for example non-linear nor phase distortions. Nonetheless, it works well in practice.

## 2.2. Aligned Spectrum Correction

In this work, the *spectrum correction* (SC) method is introduced. The goal is to transform the distortion produced by some device  $d$  into the distortion produced by some chosen reference device  $r$  by matching their spectra. The device  $d$  can be any known device. When all recording devices are known, this method can be used to transform all samples, so that they appear as if they were recorded using the reference device.

As before, let  $\tilde{A}_{s,d}(t, f)$  and  $\tilde{A}_{s,r}(t, f)$  denote amplitudes of frequency  $f$  in signals  $\tilde{w}_d$  and  $\tilde{w}_r$ , recorded using devices  $d$  and  $r$ , respectively. If device  $r$  is the reference device, the SC is derived by applying approximation from Equation (1) twice:

$$\tilde{A}_{s,r}(t, f) \approx H_r(f) A_s(t, f) \quad (3)$$

$$= \frac{H_d(f)}{H_r(f)} H_r(f) A_s(t, f) \quad (4)$$

$$\approx \frac{H_r(f)}{H_d(f)} \tilde{A}_{s,d}(t, f). \quad (5)$$

Then *correction coefficients* can be defined as:

$$C_{d \rightarrow r}(f) := \frac{H_r(f)}{H_d(f)}. \quad (6)$$

Note that the correction coefficients usually cannot be directly computed and have to be estimated from the data, which is explained in Section 2.3. In order to transform a recording from device  $d$  to look like a recording from device  $r$ , one can simply perform element-wise multiplication

$$\tilde{A}_{s,r}(t, f) \approx C_{d \rightarrow r}(f) \tilde{A}_{s,d}(t, f). \quad (7)$$

The correction coefficients can be used to directly convert spectrograms. They can also be used in order to re-synthesise audio using a converted Short Time Fourier Transform (STFT). Alternatively, as shown in Section 2.8, using a finite impulse response filter, recordings can be converted directly in the time domain.

## 2.3. Computing Correction Coefficients

In order to apply spectrum correction, one needs to compute the correction coefficients. To estimate the coefficients, one can

divide the amplitude spectrum from the reference device by the amplitude spectrum from the source device, element-wise. This can be inferred by transforming Equations (5) and (6)

$$C_{d \rightarrow r}(f) = \frac{H_r(f)}{H_d(f)} \approx \text{mean}_{s,t} \frac{\tilde{A}_{s,r}(t, f)}{\tilde{A}_{s,d}(t, f)}. \quad (8)$$

Here mean refers to the geometric mean, as it appears to be an appropriate choice for averaging scaling factors. Note that the mean is calculated over both time and examples. The significance will be explained in Section 2.7.

In Equation (5), it is implicitly assumed that the two devices recorded exactly the same signal. This could be possible only if the devices were in the same point in space at the exactly same moment. In a controlled laboratory environment it is possible to approximately reproduce the same synthetic signal twice. Then by placing the two devices in the exact same location and repeating the experiment, one could obtain such recordings.

## 2.4. Unaligned Spectrum Correction

Procedure described above requires aligned audio recordings of the same signal. For most applications such examples might not be available. However, this requirement can be easily relaxed. Here the only information that is required is which device produced the recording. The trick is to use the geometric mean and its property that

$$C_{d \rightarrow r}(f) \approx \text{mean}_{s,t} \frac{\tilde{A}_{s,r}(t, f)}{\tilde{A}_{s,d}(t, f)} = \frac{\text{mean}_{s,t} \tilde{A}_{s,r}(t, f)}{\text{mean}_{s,t} \tilde{A}_{s,d}(t, f)}. \quad (9)$$

The averages can be computed independently. Also the recordings do not need to be aligned and can be of different signals. The recorded signals only need to be from similar distributions (requirements are described in the following section).

## 2.5. Relation to Standardization

It is possible to omit the reference device, when multiplicative constants are of no concern. A recording converted in the following way has a characteristic that does not match any device, but is the same across all of the devices. It does not possess any practical advantages over the unaligned spectrum correction, but it is useful for theoretical analysis and simplified implementation. Based on the Equation (9) the following proportion holds

$$\frac{\text{mean}_{s,t} \tilde{A}_{s,r}(t, f)}{\text{mean}_{s,t} \tilde{A}_{s,d}(t, f)} \propto \frac{1}{\text{mean}_{s,t} \tilde{A}_{s,d}(t, f)} \approx C_d(f), \quad (10)$$

where  $C_d(f)$  are the correction coefficients for device  $d$ . Then the conversion is performed in the following way

$$C_d(f) \tilde{A}_{s,d}(t, f) \approx \frac{\tilde{A}_{s,d}(t, f)}{\text{mean}_{s,t} \tilde{A}_{s,d}(t, f)}, \quad \forall s \in D_d, \quad (11)$$

where  $D_d$  is a collection of signals that were recorded using the device  $d$ . Using Equation (1), this can be simplified to

$$\frac{\tilde{A}_{s,d}(t, f)}{\text{mean}_{s,t} \tilde{A}_{s,d}(t, f)} = \frac{H_d(f)}{H_d(f)} \frac{A_s(t, f)}{\text{mean}_{s,t} A_s(t, f)}. \quad (12)$$

The converted recordings are the same for all devices if  $\text{mean}_{s,t} A_s(t, f)$  is independent from the device. This is true when devices recorded the same signals. However, it is also true

if the signals are not the same, but are from the same distribution.

It turns out that a simple  $z$ -score standardization can be implemented in such a way as to perform this simplified version. Usual way to perform standardization is to subtract the mean and then divide by the standard deviation. If all averages are assumed to be geometric, then

$$\log \left( C_d(f) \tilde{A}_{s,d}(t, f) \right) \quad (13)$$

$$\approx \log \left( \frac{\tilde{A}_{s,d}(t, f)}{\left( \prod_{s \in D_d} \tilde{A}_{s,d}(f) \right)^{\frac{1}{|D_d|}}} \right) \quad (14)$$

$$= \log \tilde{A}_{s,d}(t, f) - \frac{1}{|D_d|} \sum_{s \in D_d} \log \tilde{A}_{s,d}(f) \quad (15)$$

$$= \log \tilde{A}_{s,d}(t, f) - \frac{1}{|D_d|T} \sum_{s \in D_d} \sum_t \log \tilde{A}_{s,d}(t, f), \quad (16)$$

where  $\tilde{A}_{s,d}(f)$  denotes  $\tilde{A}_{s,d}(t, f)$  averaged over time. This result shows that mean subtraction, when working with log scaled features, such as log-spectrograms or log-melspectrograms is really a special case of the presented method. Note that for this to work the standardization has to be performed for each device and frequency separately. For example average over all examples does not create this effect. This should not be confused with spectral subtraction [22]. A somewhat similar result was used for deconvolution [23].

## 2.6. Relation to Cepstral Mean Subtraction

An interesting side note is that something resembling Cepstral Mean Subtraction can be derived from Equation (16) using linearity of the inverse Fourier transform.

$$\mathcal{F}^{-1} \left[ \log \left( C_d(f) \tilde{A}_s(t, f) \right) \right] = \quad (17)$$

$$\mathcal{F}^{-1} \left[ \log \tilde{A}_s(t, f) \right] - \frac{1}{|D_d|T} \sum_{s \in D_d} \sum_t \mathcal{F}^{-1} \left[ \log \tilde{A}_s(t, f) \right].$$

The usual formulation of CMS computes average only over time for each example separately [24] (see Section 2.7).

## 2.7. Environment Response

A detail worth pointing out is that the mean in all of the four formulations is always computed for a number of recordings, never for the current recording. For example Equation (16) has a different effect than

$$\log \tilde{A}_{s,d}(t, f) - \frac{1}{T} \sum_t \log \tilde{A}_{s,d}(t, f). \quad (18)$$

This can be seen after substituting Equation (2) into both equations. For Equation (16) only the  $H_d(f)$  will cancel out. However, for Equation (18)  $H_d(f)$  will cancel out, but so will  $H_s(f)$ . This will result in loss of information about the impulse response of the environment. Note that the equations have different effect even if  $D_d$  contains only a single example.

## 2.8. Implementation in time domain

The procedure described in Sections 2.2 and 2.4 requires transformation of the signal into the frequency domain. It is completely unnecessary for models taking raw wave-forms as an input. Computing STFT can also be expensive and cumbersome in some situations.

Table 1: Architecture of the network.

layer	outputs	kernel	stride
Conv2D+ReLU+BN	16	3	1
Conv2D+ReLU+BN	32	3	2
Conv2D+ReLU+BN	32	3	1
Conv2D+ReLU+BN	64	3	2
Conv2D+ReLU+BN	64	3	1
AveragePooling	64	-	-
Dense+Softmax	10	-	-

Spectrum correction can be applied directly in the time domain. Correction coefficients can be viewed as gains for frequency bands, then standard techniques from filter design can be used to construct a finite input response (FIR) filter.

## 3. Experiments

The experiments for SC were performed based on ASC. The goal of this task is to determine the environment the device is situated in, e.g. that the surroundings of a mobile phone sound like a metro station. The source code for the experiments can be found online (<https://github.com/SRPOL-AUI/spectrum-correction>).

### 3.1. Dataset

The *TAU Urban Acoustic Scenes 2019 Mobile* [2] dataset was used for all of the reported experiments. This was the dataset used in the DCASE 2019 challenge. It consists of 46 hours of audio recordings divided into equal 10 s samples with a sampling rate of 44.1 kHz. There are three different recording devices in the dataset. Out of all the recordings, 40 hours are from a high quality device (device A) and 3 hours from each of the two mobile devices, Samsung Galaxy S7 (device B) and iPhone SE (device C). The dataset was collected in twelve cities in Europe.

### 3.2. Model and Training

The classifier was a neural network consisting of five convolutional layers with ReLU activations and batch normalization [25] followed by a global average pooling and a single dense layer with a softmax activation, resulting in 71k parameters. The details of this architecture [20] can be seen in Table 1.

The parameters of the network were optimized using Adadelta [26] with a cross-entropy loss. The learning rate was reduced by half every time accuracy did not increase for 16 epochs by at least 0.1 of a percentage point. Batch size was set to 64.

Audio samples were transformed to log-melspectrograms in the following way. First, the STFT was computed using a window of 2048 samples and a hop length of 512. To avoid unnecessary computations, spectrum correction was applied at this point (except for the implementation from Equation (16)). The resulting output was then converted to mel scale with 256 mel-bins. Next, a logarithm was applied to the amplitudes. Frequency bins were separately standardized to zero mean and unit variance each, using statistics for the entire training dataset. For the implementation from Equation (16) statistics are computed separately for each device. During training, mixup augmentation [27] was used with alpha parameter set to 0.4. Mixing was performed without the log scaling.




Figure 1: Plot of correction coefficients for devices A, B and C calculated based on 256 examples each. Device A was chosen as the reference device.

For all experiments device A was chosen as the reference, unless otherwise indicated. Figure 1 illustrates the computed coefficients.

### 3.3. Time Domain and Reference Device

To validate the claim that SC can be applied directly in the time domain, a FIR filter was constructed and tested for various reference devices. The filter was created using the least squares method using 1025 taps. Results for this experiment are shown in Table 2.

Table 2: Accuracy computed as macro average for devices A, B and C. Columns represent different choices for the reference device. Spectrum correction is applied both in the frequency domain (STFT) and in the time domain (FIR).

		reference device		
method	variant	A	B	C
aligned	STFT	70.1%	69.4%	69.6%
aligned	FIR	70.6%	70.2%	70.4%
unaligned	STFT	70.2%	69.7%	69.5%
unaligned	FIR	70.7%	70.1%	70.5%

The difference in accuracy between models trained with SC applied in the time domain using a FIR filter versus in frequency domain using STFT is negligible. It appears that a FIR filter can be used without a loss in performance. Also it is clear that the choice of the reference device does not significantly affect the results and that aligned and unaligned variants are indeed equivalent. On a side note, it is not advisable to use the simplified formulation from Section 2.5 to create the filters (due to the range of the gains).

### 3.4. Required Size of the Dataset

Various numbers of recordings can be used to compute the correction coefficients. It is interesting to see how many are actually needed. Performance of models trained using the correction coefficients computed using a limited number of examples is shown in Table 3. It is surprising to see that the accuracy does not change with the increasing number of recordings as much as one would expect.

### 3.5. Comparison with Other Methods

An important question is how models trained with and without SC compare to models trained using other methods, in terms of their accuracy. Results for those experiments are shown in Table 4. For all methods, best found configurations were reported.

Table 3: Accuracy for devices A, B and C with respect to the number of recordings used to compute the correction coefficients. Models were trained for 100 epochs.

		number of recordings			
device		1	4	16	128
A		72.2%	73.1%	72.4%	72.9%
B		65.6%	63.7%	64.1%	64.6%
C		69.0%	70.0%	69.6%	71.1%

RASTA used alternative with pole at 0.98. For PCEN (a)  $T = 0.01$  s,  $\varepsilon = 10^{-6}$ ,  $\alpha = 0.95$ ,  $\delta = 2$  and  $r = 0.5$ . For PCEN (b)  $T = 0.2$  s,  $\alpha = 1.0$ . The two variants for PCEN trade off accuracy between device A and the mobile devices. Both implementations use Librosa [28]. RASTA implementation is based on MATLAB implementation [29]. CMS is performed in frequency domain (see Equation (18)).

Table 4: Accuracy for devices A, B and C for models trained with and without SC and using other methods. Models were trained for 100 epochs.

		source device			
method		A	B	C	mean
RASTA		60.2%	59.7%	60.9%	60.1%
PCEN (a)		69.5%	55.4%	57.0%	60.6%
PCEN (b)		64.0%	58.6%	64.4%	62.3%
-		71.6%	59.2%	61.2%	64.0%
CMS		63.9%	61.4%	67.1%	64.2%
SC aligned		<b>73.4%</b>	<b>65.4%</b>	<b>71.8%</b>	<b>70.2%</b>
SC unaligned		<b>73.1%</b>	<b>65.9%</b>	<b>71.4%</b>	<b>70.4%</b>
SC mean subtraction		<b>73.2%</b>	<b>65.8%</b>	<b>72.2%</b>	<b>70.3%</b>

Methods from speech processing do not appear to perform well on ASC. However, CMS provided improvement for the mobile devices, nonetheless it deteriorated the accuracy on the main high-quality device A. This was consistent with the expectations [24]. In contrast all implementations of SC resulted in significantly better accuracy for mobile devices. Interestingly, accuracy for device A also increased. It is likely caused by more efficient use of recordings from mobile devices.

## 4. Summary

In this paper, a new approach to address learning with mismatched recording devices was proposed. It was experimentally shown that using just a few examples, it is possible to significantly improve the accuracy of a model. Particularly, accuracy for two mobile devices increased from 59% to 66% and from 61% to 72%, when SC was applied. It also outperforms classical methods such as RASTA or CMS, when applied on ASC. This method can be applied both in the frequency domain using the STFT, and in the time domain using FIR filters.

## 5. Acknowledgements

I thank Zuzanna Kwiatkowska, Sławomir Kapka, Kornel Jankowski, Krzysztof Rykaczewski, Mateusz Matuszewski, Michał Łopuszyński and Piotr Masztalski for review of this paper and Jakub Tkaczuk for help with administrative duties.

## 6. References

- [1] T. Hain, L. Burget, J. Dines, P. N. Garner, F. Grezl, A. E. Hannani, M. Huijbregts, M. Karafiat, M. Lincoln, and V. Wan, "Transcribing meetings with the amida systems," *IEEE Transactions on Audio, Speech, and Language Processing*, vol. 20, no. 2, pp. 486–498, Feb 2012.
- [2] A. Mesaros, T. Heittola, and T. Virtanen, "A multi-device dataset for urban acoustic scene classification," in *Proceedings of the Detection and Classification of Acoustic Scenes and Events 2018 Workshop (DCASE2018)*, November 2018, pp. 9–13. [Online]. Available: <https://arxiv.org/abs/1807.09840>
- [3] R. M. Stern, A. Acero, F.-H. Liu, and Y. Ohshima, "Signal Processing for Robust Speech Recognition," in *Automatic Speech and Speaker Recognition: Advanced Topics*, ser. The Kluwer International Series in Engineering and Computer Science, C.-H. Lee, F. K. Soong, and K. K. Paliwal, Eds. Boston, MA: Springer US, 1996, pp. 357–384.
- [4] G. Gosztolya and T. Grósz, "Domain adaptation of deep neural networks for automatic speech recognition via wireless sensors," *Journal of Electrical Engineering*, vol. 67, 04 2016.
- [5] T. Nguyen and F. Pernkopf, "Acoustic scene classification with mismatched recording devices using mixture of experts layer," in *2019 IEEE International Conference on Multimedia and Expo (ICME)*, July 2019, pp. 1666–1671.
- [6] S. Gharib, K. Drossos, E. Cakir, D. Serdyuk, and T. Virtanen, "Unsupervised Adversarial Domain Adaptation for Acoustic Scene Classification," in *Detection and Classification of Acoustic Scenes and Events*. Tampere University of Technology, 2018. [Online]. Available: <https://arxiv.org/abs/1808.05777>
- [7] B. Ma, H. M. Meng, and M. Mak, "Effects of device mismatch, language mismatch and environmental mismatch on speaker verification," in *2007 IEEE International Conference on Acoustics, Speech and Signal Processing - ICASSP '07*, vol. 4, April 2007, pp. IV–301–IV–304.
- [8] T. Nguyen, F. Pernkopf, and M. Kosmider, "Acoustic Scene Classification for Mismatched Recording Devices Using Heated-Up Softmax and Spectrum Correction," in *ICASSP 2020 - 2020 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, May 2020, pp. 126–130, iSSN: 2379-190X.
- [9] H. Eghbal-zadeh, K. Koutini, and G. Widmer, "Acoustic scene classification and audio tagging with receptive-field-regularized CNNs," DCASE2019 Challenge, Tech. Rep., June 2019.
- [10] W. Gao and M. McDonnell, "Acoustic scene classification using deep residual networks with late fusion of separated high and low frequency paths," DCASE2019 Challenge, Tech. Rep., June 2019.
- [11] T. Nguyen and F. Pernkopf, "Acoustic scene classification using a convolutional neural network ensemble and nearest neighbor filters," DCASE2018 Challenge, Tech. Rep., September 2018.
- [12] J.-C. Junqua, G. van Noord, N. Ide, and J. Véronis, Eds., *Robustness in Language and Speech Technology*, ser. Text, Speech and Language Technology. Dordrecht: Springer Netherlands, 2001, vol. 17. [Online]. Available: <http://link.springer.com/10.1007/978-94-015-9719-7>
- [13] L. Burget, P. Matejka, P. Schwarz, O. Glembek, and J. H. Cernocky, "Analysis of Feature Extraction and Channel Compensation in a GMM Speaker Recognition System," *IEEE Transactions on Audio, Speech, and Language Processing*, vol. 15, no. 7, pp. 1979–1986, Sep. 2007.
- [14] H. Hermansky and N. Morgan, "RASTA processing of speech," *IEEE Transactions on Speech and Audio Processing*, vol. 2, no. 4, pp. 578–589, Oct. 1994.
- [15] B. S. Atal, "Effectiveness of linear prediction characteristics of the speech wave for automatic speaker identification and verification," *The Journal of the Acoustical Society of America*, vol. 55, no. 6, pp. 1304–1312, Jun. 1974, publisher: Acoustical Society of America. [Online]. Available: <http://asa.scitation.org/doi/10.1121/1.1914702>
- [16] J. Pelecanos and S. Sridharan, "Feature Warping for Robust Speaker Verification," in *Proceedings of 2001 A Speaker Odyssey: The Speaker Recognition Workshop*. Crete, Greece: European Speech Communication Association, 2001, pp. 213–218, conference Name: 2001 A Speaker Odyssey: The Speaker Recognition Workshop Meeting Name: 2001 A Speaker Odyssey: The Speaker Recognition Workshop. [Online]. Available: <https://eprints.qut.edu.au/10408/>
- [17] C.-P. Chen and J. A. Bilmes, "MVA Processing of Speech Features," *IEEE Transactions on Audio, Speech, and Language Processing*, vol. 15, no. 1, pp. 257–270, Jan. 2007, conference Name: IEEE Transactions on Audio, Speech, and Language Processing.
- [18] Y. Wang, P. Getreuer, T. Hughes, R. F. Lyon, and R. A. Saurous, "Trainable Frontend For Robust and Far-Field Keyword Spotting," in *Proc. IEEE ICASSP 2017*, New Orleans, LA, 2017.
- [19] V. Lostanlen, J. Salamon, M. Cartwright, B. McFee, A. Farnsworth, S. Kelling, and J. P. Bello, "Per-Channel Energy Normalization: Why and How," *IEEE Signal Processing Letters*, vol. 26, no. 1, pp. 39–43, Jan. 2019. [Online]. Available: <https://ieeexplore.ieee.org/document/8514023/>
- [20] M. Kosmider, "Calibrating neural networks for secondary recording devices," DCASE2019 Challenge, Tech. Rep., Jun. 2019. [Online]. Available: [http://dcase.community/documents/challenge2019/technical\\_reports/DCASE2019\\_Kosmider.61.pdf](http://dcase.community/documents/challenge2019/technical_reports/DCASE2019_Kosmider.61.pdf)
- [21] "Electroacoustics - measurement microphones - part 5: Methods for pressure calibration of working standard microphones by comparison," International Electrotechnical Commission, Geneva, CH, International Standard IEC 61094-5:2016, 2016. [Online]. Available: <https://webstore.iec.ch/publication/24988>
- [22] S. Boll, "Suppression of acoustic noise in speech using spectral subtraction," *IEEE Transactions on Acoustics, Speech, and Signal Processing*, vol. 27, no. 2, pp. 113–120, Apr. 1979, conference Name: IEEE Transactions on Acoustics, Speech, and Signal Processing.
- [23] T. Stockham, T. Cannon, and R. Ingebretsen, "Blind deconvolution through digital signal processing," *Proceedings of the IEEE*, vol. 63, no. 4, pp. 678–692, Apr. 1975.
- [24] R. Mammone, X. Zhang, and R. Ramchandran, "Robust speaker recognition: A feature-based approach," *Signal Processing Magazine, IEEE*, vol. 13, p. 58, Oct. 1996.
- [25] S. Ioffe and C. Szegedy, "Batch normalization: Accelerating deep network training by reducing internal covariate shift," in *Proceedings of the 32nd International Conference on Machine Learning*, ser. Proceedings of Machine Learning Research, F. Bach and D. Blei, Eds., vol. 37. Lille, France: PMLR, 07–09 Jul 2015, pp. 448–456. [Online]. Available: <http://proceedings.mlr.press/v37/loff15.html>
- [26] M. D. Zeiler, "Adadelta: an adaptive learning rate method," *arXiv preprint arXiv:1212.5701*, 2012.
- [27] H. Zhang, M. Cisse, Y. N. Dauphin, and D. Lopez-Paz, "mixup: Beyond empirical risk minimization," in *International Conference on Learning Representations*, 2018. [Online]. Available: <https://openreview.net/forum?id=r1Ddp1-Rb>
- [28] B. McFee, C. Raffel, D. Liang, D. Ellis, M. McVicar, E. Battenberg, and O. Nieto, "librosa: Audio and music signal analysis in python," in *Proceedings of the 14th python in science conference*, Jan. 2015, pp. 18–24.
- [29] D. P. W. Ellis, "PLP and RASTA (and MFCC, and inversion) in Matlab," 2005. [Online]. Available: <http://www.ee.columbia.edu/~dpwe/resources/matlab/rastamat/>