

Unsupervised Feature Learning for Speech using Correspondence and Siamese Networks

Petri-Johan Last, Herman A. Engelbrecht, and Herman Kamper

Abstract—In zero-resource settings where transcribed speech audio is unavailable, unsupervised feature learning is essential for downstream speech processing tasks. Here we compare two recent methods for frame-level acoustic feature learning. For both methods, unsupervised term discovery is used to find pairs of word examples of the same unknown type. Dynamic programming is then used to align the feature frames between each word pair, serving as weak top-down supervision for the two models. For the *correspondence autoencoder* (CAE), matching frames are presented as input-output pairs. The *Triamese network* uses a contrastive loss to reduce the distance between frames of the same predicted word type while increasing the distance between negative examples. For the first time, these feature extractors are compared on the same discrimination tasks using the same weak supervision pairs. We find that, on the two datasets considered here, the CAE outperforms the Triamese network. However, we show that a new hybrid *correspondence-Triamese* approach (CTriamese), consistently outperforms both the CAE and Triamese models in terms of average precision and ABX error rates on both English and Xitsonga evaluation data.

Index Terms—Unsupervised learning, correspondence autoencoders, Siamese networks, zero-resource speech processing.

I. INTRODUCTION

AUTOMATIC speech recognition systems are usually developed using large transcribed speech datasets. However, for many languages it is difficult or impossible to collect the annotated resources required for training supervised speech recognition models [1]. In so-called *zero-resource* environments, where unlabelled speech is the only available resource, unsupervised acoustic feature learning is essential for downstream tasks such as speech retrieval or indexing [2]–[5]. Features should ideally disregard irrelevant information (such as speaker and gender), while capturing linguistically meaningful contrasts (such as phone or word categories). Several different unsupervised frame-level acoustic feature learning methods have been developed over the last few years [6]–[12], with neural networks being used in a number of studies [13]–[17].

One type of neural approach that has received particular attention is *Siamese networks* [18]–[22]. A Siamese network consists of two identical sub-networks with tied weights taking in a pair of inputs [23]. The standard Siamese network is trained so that the distance between the outputs of the sub-networks are minimised when the inputs are of the same type, and maximised when the inputs are of different types. To train a Siamese network in an unsupervised fashion on unannotated speech, paired input is required. Building on an idea from [24], Thiollière et al. [19] used an unsupervised term discovery (UTD) system [25], [26] to automatically discover pairs of word-like segments predicted to be of the same type;

these segments were aligned using dynamic time warping (DTW), giving matching frame pairs which were presented to the Siamese network as positive pairs. Features from the resulting unsupervised Siamese network performed well in a minimal-pair triphone discrimination task [7], [27].

Siamese networks have subsequently been applied and further developed in a number of studies [20]–[22]. Notably, [21] introduced the *Triamese network*, which takes triplets instead of paired input, with a triplet consisting of two positive items and one negative item. A contrastive loss is used to minimise the distance between the two positive items such that it is smaller (by some margin) than the distance between a positive and negative item. In discrimination tasks, it is often not the absolute distance between items that is of interest, but rather the relative distance—in contrast to the standard Siamese network, the Triamese network optimises relative distances. Apart from showing that the Triamese approach is superior, [21] also attempted to explicitly disentangle speaker and phonetic information within a single network (but this had limited effect).

Another type of neural approach that has seen success in learning acoustic features is the *correspondence autoencoder* (CAE) [28]. Instead of trying to reconstruct its input, as is done in a standard autoencoder [29], the CAE tries to reconstruct another instance of the same type as its input. To obtain input-output pairs from unlabelled speech, [28] followed the same approach as used in the Siamese studies [19]: a UTD system with subsequent DTW is used to find matching frame pairs to serve as weak supervision. By using the terms produced by the UTD system to train the CAE, the approach as a whole is unsupervised. The CAE was also applied and extended in a number of subsequent studies [15], [30]–[34].

Although the one uses a reconstruction-like loss while the other explicitly minimises or maximises a distance-based loss, both the CAE and Triamese networks rely on the same weak supervision from another unsupervised system. Despite their commonalities, the performance of these networks has never been compared when trained on exactly the same paired data using the same evaluation metric. The results in [15] and [19] are somewhat comparable (here the Siamese model seems to outperform an early version of the CAE), but the supervision pairs in the two studies are not the same and the feature dimensionalities are very different, so it is not possible to draw definitive conclusions. In this paper we compare the performance of the two approaches in a controlled experimental setup: we train each network on the same discovered terms and evaluate the networks on the same tasks. In this setting, we find that the CAE outperforms the Triamese network. More importantly, we show that the two approaches are

complementary and that they can be combined to form a *correspondence-Triamese* (CTriamese) network. This method consistently outperforms both the CAE and Triamese models in an evaluation on English and Xitsonga data.

II. UNSUPERVISED SPEECH FEATURE LEARNING

We consider three approaches to unsupervised feature learning of frame-level acoustic features. The first two (the CAE and Triamese models) were developed in previous work, while the last is a new hybrid approach. All three approaches use weak top-down constraints in the form of discovered word pairs. Concretely, we apply an unsupervised term discovery (UTD) system [26] to an unlabelled speech collection, resulting in pairs of word segments predicted to be of the same unknown type. For each discovered word pair, dynamic time warping (DTW) is used to align the acoustic frames between the two words, yielding a set of frame pairs $\{(x_a, x_b)\}$, where each $x \in \mathbb{R}^D$ is an acoustic feature vector (e.g. MFCCs). Each model below uses these frame pairs differently as weak supervision.

A. Correspondence autoencoder (CAE)

Instead of using the same frames as input and output as in a standard autoencoder, the *correspondence autoencoder* (CAE) uses the aligned frames as input-output pairs [28]. Formally, given an input frame x_a , the CAE is trained using the loss $\ell_{\text{cae}}(x_a, x_b) = \|\hat{x} - x_b\|^2$, where \hat{x} denotes the output of the network and x_b the frame aligned to x_a . The overall approach is illustrated in Figure 1. The CAE network consists of a number of encoder layers, a 39-dimensional bottleneck layer, and decoder layers. After training, we use the network up to the bottleneck layer e as our final feature extractor.

The idea behind the CAE is that it will result in features that are insensitive to factors not common to the discovered word pairs, such as speaker, gender and channel, while remaining dependent on factors that are, such as the word identity.

B. Triamese network

An alternative to the reconstruction-like loss of the CAE is to train a model which explicitly learns to minimise a distance between paired inputs. A *Siamese network* [23] refers to a pair of networks with tied parameters which is trained to minimise the distance between two instances of the same type while




Fig. 1. The correspondence autoencoder (CAE) is trained to reconstruct one acoustic frame x_b in a discovered word from another x_a . The word pairs are discovered using unsupervised term discovery and aligned using DTW.

maximising the distance between instances of different types. This idea was first used for speech representation learning in [18]. In [21], Zeghidour et al. extended this to *Triamese networks*. Rather than minimising or maximising an absolute distance, Triamese networks optimise the relative distance between instances of the same type and a negative example.

Concretely, a Triamese network consists of three branches. Parameters are shared between the branches, and each branch produces a feature embedding of its corresponding input. Inputs consists of triples (x_a, x_b, x') , where x_a and x_b are of the same type and x' is a negative example; the network produces corresponding embeddings (e_a, e_b, e') . The network is then trained using the following contrastive loss function [35]:

$$\ell_{\text{triplet}}(x_a, x_b, x') = \max\{0, m + d_{\text{cos}}(e_a, e_b) - d_{\text{cos}}(e_a, e')\} \quad (1)$$

where $d_{\text{cos}}(\cdot, \cdot)$ denotes the cosine distance and m denotes a margin parameter. This loss is at a minimum when all embeddings e_a and e_b of the same type are more similar by a margin m than embeddings e_a and e' of different types. After training, one of the tied branches is used as a feature extractor. The approach is illustrated in Figure 2.

The motivation for this approach is similar to that of the CAE in that we hope to obtain features that retain meaningful information while disregarding nuisance factors. However, while the CAE employs a soft reconstruction loss, the Triamese model optimises distances between representations discriminatively. The CAE is also only trained on positive pairs, while the Triamese model requires negative examples. Here we sample negative examples randomly [21], but [22] showed that more involved sampling schemes can give improvements.

C. Correspondence-Triamese network (CTriamese)

Although the CAE and Triamese networks are both trained on discovered word pairs, they use this information very differently—the two approaches could therefore be complementary. The *correspondence-Triamese network* (CTriamese) applies the triplet loss function of the Triamese network to intermediate representations obtained from the CAE. Concretely, each branch of the Triamese network above is replaced with a full CAE. We then apply the triplet loss to the bottleneck layers from each CAE. The full CTriamese network takes inputs (x_a, x_b, x'_a, x'_b) , where (x_a, x_b) is a pair of the same predicted type, which is different from (x'_a, x'_b) . The three CAE networks have the following respective input-output




Fig. 2. The Triamese network is trained so that the embeddings e_a and e_b of the same type are more similar by a margin m than embeddings e_a and e' of different types.




Fig. 3. The correspondence-Triamese (CTriamese) network is trained using the combined loss of the three CAE branches and the contrastive triplet loss.

pairs: (x_a, x_b) , (x_b, x_a) and (x'_a, x'_b) , as shown in Figure 3. Given the inputs, embeddings (e_a, e_b, e'_a) are produced and the network is trained with the following loss:

$$\begin{aligned} \ell(x_a, x_b, x'_a, x'_b) = & \ell_{\text{cae}}(x_a, x_b) + \ell_{\text{cae}}(x_b, x_a) + \ell_{\text{cae}}(x'_a, x'_b) \\ & + \ell_{\text{triplet}}(x_a, x_b, x'_a) \end{aligned}$$

Starting from a random initialisation, the parameters of the entire network is trained jointly to optimise this loss. For the CTriamese model, we also additionally include speaker information: after the bottleneck layer of each CAE, a speaker embedding is incorporated based on the output speaker identity for that CAE branch. These embeddings are learned jointly with the rest of the model and are concatenated to the first hidden layer after the bottleneck. This means that speaker information is not required at test time (when features are taken from the bottleneck layers). In the experiments below, we incorporate speaker information in a similar way into the basic CAE.

III. EXPERIMENTAL SETUP

A. Data

We perform experiments on English and Xitsonga data. For the English experiments, we use a subset of the conversational Buckeye corpus [36]. As training data, around 12k terms from 12 speakers are discovered by applying the UTD system of [26] to the complete set of unlabelled data. For validation and testing, around 3k and 4k true word segments are respectively used. There is no speaker overlap between the training, validation or test data. On English, early stopping on the validation set is performed to determine when to terminate training.

For the Xitsonga experiments, we use data from the NCHLT corpus [37]. Around 12k UTD terms from 24 speakers are discovered from the complete unlabelled training set, with 6.5k true words used for testing. More details of the UTD system can be found in [7], where pair-wise matching precisions of 39% and 69% are reported on English and Xitsonga, respectively. The optimal hyperparameters determined on the English validation data are used directly on Xitsonga without any alteration, representing the realistic setting where we have a zero-resource language without any labelled validation data. Since validation data is also not available for early stopping,

the Xitsonga models are trained for the number of epochs that was found to be optimal on the English validation data.

We use cepstral mean and variance normalised Mel-frequency cepstral coefficients (MFCCs) with first- and second-order derivatives to represent the speech audio. The discovered word pairs are DTW-aligned to obtain frame-level pairs for training the different networks. For the Triamese network, which also requires negative examples, we randomly choose a frame from a third discovered word that is spoken by the same speaker as the first word in the pair, but with a different predicted type. The CTriamese network requires a fourth item for the final CAE branch's output. We randomly choose a fourth word that matches the predicted type of the third word in the triplet, and DTW-align it with the third word to form the input-output frame-pair for the third network branch.

B. Neural network architectures

1) *Correspondence autoencoder*: The CAE (§II-A) consists of two halves: an encoder and a decoder. In our case, the encoder and decoder both consist of six 100-unit ReLU layers with a 39-dimensional middle bottleneck layer from which features are extracted. The network is trained using the Adadelata optimiser [38] with a learning rate of 0.001, as was done in [28].

2) *Triamese network*: A single branch in our Triamese network (§II-B) consists of six densely connected 100-unit ReLU layers followed by a 39-unit ReLU embedding layer. [21] used four hidden layers of 1000 units each and an embedding layer of 100 units. We use a different architecture to keep embeddings comparable (39 dimensional), but include development results using the architecture of [21]. As in [21], the Triamese network is trained using stochastic gradient descent with a learning rate of 0.01 and a decay of 10^{-6} to optimise the contrastive loss in (1) with a margin of $m = 0.15$.

3) *Correspondence-Triamese network*: Each branch of the CTriamese network (§II-C) mirrors the structure of a single CAE, consisting of six 100-unit ReLU encoding layers, a 39-unit bottleneck embedding layer, and another six 100-unit ReLU decoding layers. For the contrastive loss, a margin of $m = 0.15$ was found to perform best on the English validation data. For adding speaker information in the CAE and CTriamese models, we use a speaker-embedding dimensionality of 100. The CTriamese network is trained using Adadelata [38] with a learning rate of 0.001. On our data, each model took less than an hour to train on a single GeForce GTX 1080 GPU.

C. Evaluation

The multi-speaker word discrimination task developed in [39], called the *same-different task*, quantifies how similar words of the same type are relative to other words of different types based on a given speech representation. This task has been developed specifically for evaluating speech features without requiring downstream modelling assumptions. By calculating the DTW distance between every pair in a set of isolated test words (represented using the features under evaluation) we classify two words as being the same if the distance between them is less than some threshold. The threshold is then varied from zero to the maximum to produce a precision-recall curve.

The area under this curve is called the average precision (AP), which we use as our first evaluation metric.

We also consider ABX discrimination, a task inspired by human perceptual studies asking whether the distance between an anchor item A and another item X of the same type is smaller than the distance between the anchor and an imposter B according to the representation under evaluation [27]. We specifically use the cross-speaker triphone discrimination task [7], where triphone A and B are minimal pairs differing in the middle phone (e.g. ‘bag’ and ‘bug’), with A and B from the same speaker and X from a different speaker. ABX is reported as an error rate averaged over all minimal triphone pairs in the corpus.¹

IV. EXPERIMENTAL RESULTS

Table I gives results on the English and Xitsonga validation and test sets for unsupervised models trained on UTD-discovered terms. We first evaluated the AP of all the models on the English validation set, and then evaluated only the most representative models on the English and Xitsonga test data (which is why test results are missing for three model variants). Speaker conditioning is indicated with the ‘w/ speaker’ label. A larger AP indicates that the features allow for better word discrimination while a lower ABX error rate indicates better minimal-pair triphone discrimination. We also include a baseline where the same-different and ABX tasks are performed directly using DTW over MFCCs.

Table I indicates that the CAE outperforms the Triamese networks on both languages. While the CAE outperforms the MFCC baseline in all cases, the Triamese model with a 39-dimensional embedding layer fares worse than MFCCs on the same-different task (AP) on the English test data. To show that this is due to the lower dimensionality used here, we include development results where a 100-dimensional embedding is used with the architecture of [21]. This model achieves a 13% relative improvement in AP over the MFCCs. Moreover, despite the lower dimensionality, the 39-dimensional Triamese model still outperforms MFCCs in the ABX task on both languages (as reported in previous work). Future work could consider architectural improvements and more advanced Triamese-pair sampling schemes [22], but our specific goal here is a fair like-for-like comparison between the different models, which is why we use the 39-dimensional Triamese model for testing.

Considering the results of the CAE and CTriamese networks with and without speaker conditioning, we see that speaker conditioning slightly worsens the performance of the CAE while the CTriamese benefits from speaker conditioning. In [21], a method for incorporating speaker information in the Triamese model was also proposed, but this did not consistently improve discrimination performance.

The best performance across all sets and languages are achieved by the hybrid CTriamese network, outperforming the MFCC baseline and both the CAE and Triamese networks in terms of AP and ABX error rate. This shows that the CAE and Triamese network’s methodologies are complementary.

¹We use the most recent version of the ABX evaluation code [40]. Because of changes, this means that results might not be directly comparable to previous studies, but performance across the models here can be compared directly.

TABLE I
AVERAGE PRECISION (AP) AND ABX ERROR RATE (BOTH IN %) FOR UNSUPERVISED NEURAL NETWORK FEATURE EXTRACTORS TRAINED ON TERMS FROM AN UNSUPERVISED TERM DISCOVERY (UTD) SYSTEM.

Model	English val.	English test		Xitsonga test	
	AP	AP	ABX	AP	ABX
MFCC	36.8	35.9	21.4	28.1	19.0
CAE	47.4	46.0	16.9	57.4	12.1
CAE w/ speaker	47.2	-	-	-	-
Triamese (39 dim.)	38.2	34.8	20.7	45.9	17.6
Triamese (100 dim.) [18]	41.5	-	-	-	-
CTriamese	49.3	-	-	-	-
CTriamese w/ speaker	50.4	47.2	16.5	60.7	12.0

TABLE II
AVERAGE PRECISION (%) ON ENGLISH DATA WHEN TRAINING THE MODELS ON GROUND TRUTH WORDS INSTEAD OF UTD-DISCOVERED TERMS.

Model	Validation	Test
MFCC	36.8	35.9
CAE	51.9	45.5
Triamese (39 dim.)	46.3	38.8
CTriamese w/ speaker	53.0	50.8

In Table II we also present the same-different performance of the models when trained on ground truth English words instead of discovered terms in order to quantify the performance loss induced by the use of UTD terms in the truly unsupervised setting (Table I). While the trends are the same as in Table I, the Triamese network sees a larger relative increase in performance when trained on the ground truth words. This seems to indicate that the soft reconstructive loss of the CAE is less sensitive to the quality of the training data than the contrastive loss of the Triamese network. The difference in performance between the CAE and CTriamese approaches are also smaller in Table II (when using ground truth words) than in Table I (the unsupervised case).

V. CONCLUSION

We have compared the performance of the correspondence autoencoder (CAE) and Triamese networks on word and minimal-pair triphone discrimination tasks. We found that the CAE outperforms our implementation of the Triamese network. More importantly, we showed that the methodologies of the Triamese and CAE networks are complementary: we combine the two into a new CTriamese hybrid network, which outperforms both the individual approaches in intrinsic evaluations on English and Xitsonga data. In future work we plan to use the CTriamese approach in downstream speech processing tasks, and to incorporate some of the more advanced Triamese-pair sampling strategies that have been proposed [22].

ACKNOWLEDGEMENTS

This work is based on research supported in part by the National Research Foundation of South Africa (Grant Numbers: 120409). HK is supported by a Google Faculty Award.

REFERENCES

- [1] L. Besacier, E. Barnard, A. Karpov, and T. Schultz, "Automatic speech recognition for under-resourced languages: A survey," *Speech Commun.*, vol. 56, pp. 85–100, 2014.
- [2] Y. Zhang and J. R. Glass, "Unsupervised spoken keyword spotting via segmental DTW on Gaussian posteriorgrams," in *Proc. ASRU*, 2009.
- [3] K. Levin, A. Jansen, and B. Van Durme, "Segmental acoustic indexing for zero resource keyword search," in *Proc. ICASSP*, 2015.
- [4] S. Settle, K. Levin, H. Kamper, and K. Livescu, "Query-by-example search with discriminative neural acoustic word embeddings," in *Proc. Interspeech*, 2017.
- [5] Y.-H. Wang, H.-y. Lee, and L.-s. Lee, "Segmental audio word2vec: Representing utterances as sequences of vectors with applications in spoken term detection," in *Proc. ICASSP*, 2018.
- [6] A. Jansen, E. Dupoux, S. Goldwater, M. Johnson, S. Khudanpur, K. Church, N. Feldman, H. Hermansky, F. Metze, R. Rose, M. Seltzer, P. Clark, I. McGraw, B. Varadarajan, E. Bennett, B. Borschinger, J. Chiu, E. Dunbar, A. Fourtassi, D. Harwath, C. Lee, K. Levin, A. Norouzzian, V. Peddinti, R. Richardson, T. Schatz, and S. Thomas, "A summary of the 2012 JHU CLSP workshop on zero resource speech technologies and models of early language acquisition," in *Proc. ICASSP*, 2013.
- [7] M. Versteegh, R. Thiollire, T. Schatz, X.-N. Cao Kam, X. Anguera, A. Jansen, and E. Dupoux, "The Zero Resource Speech Challenge 2015," in *Proc. Interspeech*, 2015.
- [8] M. Versteegh, X. Anguera, A. Jansen, and E. Dupoux, "The Zero Resource Speech Challenge 2015: Proposed approaches and results," in *Proc. SLTU*, 2016.
- [9] E. Dunbar, X. N. Cao, J. Benjumea, J. Karadayi, M. Bernard, L. Besacier, X. Anguera, and E. Dupoux, "The Zero Resource Speech Challenge 2017," in *Proc. ASRU*, 2017.
- [10] M. Heck, S. Sakti, and S. Nakamura, "Feature optimized DPGMM clustering for unsupervised subword modeling: A contribution to ZeroSpeech 2017," in *Proc. ASRU*, 2017.
- [11] M. Heck, S. Sakti, and S. Nakamura, "Learning supervised feature transformations on zero resources for improved acoustic unit discovery," *IEICE T. Inf. Syst.*, vol. 101, no. 1, pp. 205–214, 2018.
- [12] E. Dunbar, R. Algayres, J. Karadayi, M. Bernard, J. Benjumea, X.-N. Cao, L. Miskic, C. Dugrain, L. Ondel, A. W. Black, L. Besacier, S. Sakti, and E. Dupoux, "The Zero Resource Speech Challenge 2019: TTS without T," in *Proc. Interspeech*, 2019.
- [13] L. Badino, C. Canevari, L. Fadiga, and G. Metta, "An auto-encoder based approach to unsupervised learning of subword units," in *Proc. ICASSP*, 2014.
- [14] L. Badino, A. Mereta, and L. Rosasco, "Discovering discrete subword units with binarized autoencoders and hidden-Markov-model encoders," in *Proc. Interspeech*, 2015.
- [15] D. Renshaw, H. Kamper, A. Jansen, and S. J. Goldwater, "A comparison of neural network methods for unsupervised representation learning on the Zero Resource Speech Challenge," in *Proc. Interspeech*, 2015.
- [16] T. Tsuchiya, N. Tawara, T. Ogawa, and T. Kobayashi, "Speaker invariant feature extraction for zero-resource languages with adversarial learning," in *Proc. ICASSP*, 2018.
- [17] R. Eloff, A. Nortje, B. L. Van Niekerk, A. Govender, L. Nortje, A. Pretorius, E. Van Biljon, E. Van der Westhuizen, L. Van Staden, and H. Kamper, "Unsupervised acoustic unit discovery for speech synthesis using discrete latent-variable neural networks," in *Proc. Interspeech*, 2019.
- [18] G. Synnaeve, T. Schatz, and E. Dupoux, "Phonetics embedding learning with side information," in *Proc. SLT*, 2014.
- [19] R. Thiollire, E. Dunbar, G. Synnaeve, M. Versteegh, and E. Dupoux, "A hybrid dynamic time warping-deep neural network architecture for unsupervised acoustic modeling," in *Proc. Interspeech*, 2015.
- [20] N. Zeghidour, G. Synnaeve, M. Versteegh, and E. Dupoux, "A deep scattering spectrum-deep Siamese network pipeline for unsupervised acoustic modeling," in *Proc. ICASSP*, 2016.
- [21] N. Zeghidour, G. Synnaeve, N. Usunier, and E. Dupoux, "Joint learning of speaker and phonetic similarities with siamese networks," *Proc. Interspeech*, 2016.
- [22] R. Riad, C. Dancette, J. Karadayi, N. Zeghidour, T. Schatz, and E. Dupoux, "Sampling strategies in Siamese networks for unsupervised speech representation learning," in *Proc. Interspeech*, 2018.
- [23] J. Bromley, J. W. Bentz, L. Bottou, I. Guyon, Y. LeCun, C. Moore, E. Säckinger, and R. Shah, "Signature verification using a "Siamese" time delay neural network," *Int. J. Pattern Rec.*, vol. 7, no. 4, pp. 669–688, 1993.
- [24] A. Jansen, S. Thomas, and H. Hermansky, "Weak top-down constraints for unsupervised acoustic model training," in *Proc. ICASSP*, 2013.
- [25] A. S. Park and J. R. Glass, "Unsupervised pattern discovery in speech," *IEEE Trans. Audio, Speech, Language Process.*, vol. 16, no. 1, pp. 186–197, 2008.
- [26] A. Jansen and B. Van Durme, "Efficient spoken term discovery using randomized algorithms," in *Proc. ASRU*, 2011.
- [27] T. Schatz, V. Peddinti, F. Bach, A. Jansen, H. Hermansky, and E. Dupoux, "Evaluating speech features with the minimal-pair ABX task: Analysis of the classical MFC/PLP pipeline," in *Proc. Interspeech*, 2013.
- [28] H. Kamper, M. Elsner, A. Jansen, and S. J. Goldwater, "Unsupervised neural network based feature extraction using weak top-down constraints," in *Proc. ICASSP*, 2015.
- [29] I. Goodfellow, Y. Bengio, and A. Courville, *Deep Learning*. MIT Press, 2016, <http://www.deeplearningbook.org>.
- [30] D. Renshaw, "Representation learning for unsupervised speech processing," Master's thesis, University of Edinburgh, 2016.
- [31] H. Kamper, A. Jansen, and S. Goldwater, "A segmental framework for fully-unsupervised large-vocabulary speech recognition," *Comput. Speech Lang.*, vol. 46, pp. 154–174, 2017.
- [32] E. Hermann and S. J. Goldwater, "Multilingual bottleneck features for subword modeling in zero-resource languages," in *Proc. Interspeech*, 2018.
- [33] R. Menon, H. Kamper, J. Quinn, and T. Niesler, "Almost zero-resource ASR-free keyword spotting using multilingual bottleneck features and correspondence autoencoders," in *Proc. Interspeech*, 2019.
- [34] H. Kamper, "Truly unsupervised acoustic word embeddings using weak top-down constraints in encoder-decoder models," in *Proc. ICASSP*, 2019.
- [35] K. M. Hermann and P. Blunsom, "Multilingual distributed representations without word alignment," in *Proc. ICLR*, 2014.
- [36] M. A. Pitt, K. Johnson, E. Hume, S. Kiesling, and W. Raymond, "The Buckeye corpus of conversational speech: Labeling conventions and a test of transcriber reliability," *Speech Commun.*, vol. 54, no. 1, pp. 89–95, 2005.
- [37] N. J. De Vries, M. H. Davel, J. Badenhorst, W. D. Basson, F. De Wet, E. Barnard, and A. De Waal, "A smartphone-based ASR data collection tool for under-resourced languages," *Speech Commun.*, vol. 56, pp. 119–131, 2014.
- [38] M. D. Zeiler, "ADADELTA: An adaptive learning rate method," *arXiv preprint arXiv:1212.5701*, 2012.
- [39] M. A. Carlin, S. Thomas, A. Jansen, and H. Hermansky, "Rapid evaluation of speech representations for spoken term discovery," in *Proc. Interspeech*, 2011.
- [40] T. Schatz, "ABX-discriminability measures and applications," Ph.D. dissertation, Université Paris 6 (UPMC), 2016.