



Multilingual Acoustic Word Embeddings For Zero-Resource Languages

by
Christiaan Jacobs

*Dissertation presented for the degree of
Doctor of Philosophy (Electronic Engineering) in the Faculty of
Engineering at Stellenbosch University*

Supervisor: Prof. Herman Kamper

December 2023



ACKNOWLEDGEMENTS

I am profoundly grateful for the privilege of having Prof. Herman Kamper as my supervisor. His guidance, wisdom, and unwavering support have not only shaped this dissertation but have also left an indelible mark on various facets of my life beyond the realms of academia. I appreciate the opportunities I had to learn and grow under his supervision.

To all those individuals whom I encountered during this journey, your kindness and patience have been a source of inspiration. I appreciate the time you dedicated to sharing your insights and expertise.

A special appreciation goes to my friends and family who supported me through the highs and lows. Your encouragement, understanding, and belief in me have been invaluable.

This endeavour would not have been possible without the collective support of all those around me. Thank you for being an integral part of this journey.



DECLARATION

By submitting this dissertation electronically, I declare that the entirety of the work contained therein is my own, original work, that I am the sole author thereof (save to the extent explicitly otherwise stated), that reproduction and publication thereof by Stellenbosch University will not infringe any third party rights and that I have not previously in its entirety or in part submitted it for obtaining any qualification.

December 2023

ABSTRACT

Developing speech applications with neural networks require large amounts of transcribed speech data. The scarcity of labelled speech data therefore restricts the development of speech applications to only a few well-resourced languages. To address this problem, researchers are taking steps towards developing speech models for languages where no labelled data is available. In this *zero-resource* setting, researchers are developing methods that aim to learn meaningful linguistic structures from unlabelled speech alone.

Many zero-resource speech applications require speech segments of different durations to be compared. Acoustic word embeddings (AWEs) are fixed-dimensional representations of variable-duration speech segments. Proximity in vector space should indicate similarity between the original acoustic segments. This allows fast and easy comparison between spoken words.

To produce AWEs for a zero-resource language, one approach is to use unlabelled data from the target language. Another approach is to exploit the benefits of supervised learning by training a single multilingual AWE model on data from multiple well-resourced languages, and then applying the resulting model to an unseen target language. Previous studies have shown that the supervised multilingual transfer approach outperforms the unsupervised monolingual approach. However, the multilingual approach is still far from reaching the performance of supervised AWE approaches that are trained on the target language itself.

In this thesis, we make five specific contributions to the development of AWE models and their downstream application. First, we introduce a novel AWE model called the CONTRASTIVERNN. We compare this model against state-of-the-art AWE models. On a word discrimination task, we show that the CONTRASTIVERNN outperforms all existing models in the unsupervised monolingual setting with an absolute improvement in average precision ranging from 3.3% to 17.8% across six evaluation languages. In the multilingual transfer setting, the CONTRASTIVERNN performs on par with existing models.

As our second contribution, we propose a new adaptation strategy. After a multilingual model is trained, instead of directly applying it to a target language, we first fine-tune the model using unlabelled data from the target language. The CONTRASTIVERNN, although performing on par with multilingual variants, showed the highest increase after adaptation, giving an improvement of roughly 5% in average precision on five of the six evaluation

ABSTRACT

languages.

As our third contribution, we take a step back and question the effect a particular set of training languages have on a target language. We specifically investigate the impact of training a multilingual model on languages that belong to the same language family as the target language. We perform multiple experiments on African languages which show the benefit of using related languages over unrelated languages. For example, a multilingual model trained on one-tenth of the data from a related language outperforms a model trained on all the available training data from unrelated languages.

As our fourth contribution, we showcase the applicability of AWEs by applying them to a real downstream task: we develop an AWE-based keyword spotting system (KWS) for hate speech detection in radio broadcasts. We validate performance using actual Swahili radio audio extracted from radio stations in Kenya, a country in Sub-Saharan Africa. In developmental experiments, our system falls short of a speech recognition-based KWS system using five minutes of annotated target data. However, when applying the system to real in-the-wild radio broadcasts, our AWE-based system (requiring less than a minute of template audio) proves to be more robust, nearly matching the performance of a 30-hour speech recognition model.

In the fifth and final contribution, we introduce three novel semantic AWE models. The goal here is that the resulting embeddings should not only be similar for words from the same type but also for words sharing contextual meaning, similar to how textual word embeddings are grouped together based on semantic relatedness. For instance, spoken instances of “football” and “soccer”, although acoustically different, should have similar acoustic embeddings. We specifically propose leveraging a pre-trained multilingual AWE model to assist semantic modelling. Our best approach involves clustering word segments using a multilingual AWE model, deriving soft pseudo-word labels from the cluster centroids, and then training a classifier model on the soft vectors. In an intrinsic word similarity task measuring semantics, this multilingual transfer approach outperforms all previous semantic AWE methods. We also show—for the first time—that AWEs can be used for downstream semantic query-by-example search.

UITTREKSEL

Die ontwikkeling van spraaktoepassings met die gebruik van neurale netwerke vereis groot hoeveelhede geannoteerde spraakdata. Die tekort aan geannoteerde data beperk dus die ontwikkeling van spraaktoepassings tot slegs 'n paar tale waarvoor daar wel data beskikbaar is. Om hierdie probleem aan te spreek is navorsers besig om stappe te neem om spraakmodelle te ontwikkel vir tale waarvoor daar geen geannoteerde spraakdata beskikbaar is nie. In hierdie *nul-hulpbron* omgewing, ontwikkel navorsers metodes wat daarop gerig is om betekenisvolle taalkundige strukture uit ongeannoteerde spraak te leer.

Baie nul-hulpbron spraaktoepassings vereis dat spraaksegmente van verskillende lengtes met mekaar vergelyk word. Akoestiese woordvektor (AWV) modelle projekteer spraaksegmente van arbitrêre lengte na 'n enkele vektor. Nabyheid in hierdie vektorruimte moet 'n aanduiding gee van hoe soortgelyk die oorspronklike akoestiese segmente is. Dit laat toe dat gesproke woorde vinnig en maklik met mekaar vergelyk kan word.

Een benadering om AWVe vir 'n nul-hulpbrontaal te produseer is om ongeannoteerde spraakdata van die teikentaal te gebruik. 'n Alternatiewe benadering is om die voordele van getoesigde leer te benut deur 'n enkele veeltalige AWV model af te rig op data van veelvuldige goedbefondsde tale, en dan gevolglik die model op 'n nul-hulpbron teikentaal toe te pas. Vorige studies het getoon dat die veeltalige oordragbenadering onder toesig beter resultate lewer as die eentalige benadering sonder toesig. Daar is egter 'n groot gaping tussen die kwaliteit van die veeltalige oordragbenadering AWVe in vergelyking met die kwaliteit van AWVe wat onder toesig afgerig is op data van die teikentaal.

In hierdie tesis lewer ons vyf spesifieke bydraes in die ontwikkeling van AWV modelle vir nul-hulpbrontale. Eerstens stel ons 'n nuwe AWV model bekend genaamd die CONTRASTIVERNN. Ons vergelyk hierdie model met bestaande AWV modelle. Op 'n woorddiskriminasietoets wys ons dat die CONTRASTIVERNN beter presteer as alle bestaande AWV modelle in die eentalige omgewing sonder toesig met 'n absolute verbetering in gemiddelde presisie wat wissel tussen 3.3% tot 17.8% oor ses evalueringstale. In die veeltalige oordragomgewing presteer die CONTRASTIVERNN op gelyke voet met bestaande modelle.

Vir ons tweede bydrae stel ons 'n nuwe afrigstrategie voor. Nadat 'n veeltalige model afgerig is, in plaas daarvan om dit direk op 'n teikentaal toe te pas, verfyn ons eers die model deur gebruik te maak van ongeannoteerde data van die teikentaal. Alhoewel die

CONTRASTIVERNN op gelyke voet met veeltalige variante presteer, toon dit die hoogste toename na aanpassing, met verbeteringe van ongeveer 5% in presisie op vyf van die ses evalueringstale.

Vir ons derde bydrae neem ons 'n tree terug en bevraagteken die effek wat 'n bepaalde stel afrigtale op 'n teikentaal het. Ons ondersoek spesifiek die impak wat die afrigtale in 'n veeltalige AWV model op die teikentaal het wanneer die afrigtale in dieselfde familie as die teikentaal is. Ons doen verskei eksperimente op Afrika-tale wat die voordeel van die gebruik van verwante tale bo onverwante tale toon. 'n Veeltalige model wat op een tiende van die data van 'n verwante taal afgerig is, presteer byvoorbeeld beter as 'n model wat afgerig is op al die beskikbare afrigdata van onverwante tale.

Vir ons vierde bydrae ontwikkel ons 'n sleutelwoordopsporingstelsel wat gebruik maak van AWVe om haatspraak in radiouitsendings te identifiseer. Ons toets ons stelsel op regte radiouitsendings wat versamel is van radiostasies in Kenia. In ontwikkelings eksperimente presteer 'n automatiese spraakherkenningsstelsel, wat slegs vyf minute afrigdata van die teikentaal gebruik, beter as ons AWV gebaseerde stelsel. Nietemin bewys ons AWV gebaseerde stelsel (wat minder as 'n minuut se templatwoorde nodig) sy robuustheid in praktiese toepassings op werklike radiouitsendings deur soortgelyk te presteer as 'n 30-uur spraakherkenningsstelsel.

Vir ons vyfde en finale bydrae stel ons drie nuwe semantiese AWV modelle voor. Hierdie woordvektore behoort nie net soortgelyk te wees vir woorde van dieselfde tipe nie maar ook vir woorde wat kontekstuele betekenis deel, soortgelyk aan hoe teks gebaseerde woordvektore semantiese verhoudings weerspieël. Gesproke voorbeelde van byvoorbeeld, “voetbal” en “sokker”, alhoewel hulle akoesties verskillend is, behoort soortgelyke woordvektore te hê. Ons beste benadering behels die trosvorming van spraaksegmente deur gebruik te maak van 'n veeltalige AWV model. Ons gebruik dan die trossentroïede om sagte pseudowoordetikette af te lei, wat daarna gebruik word om 'n klassifiseerdermodel af te rig. In 'n intrinsieke woordgelykheidstaak wat semantiek meet, presteer hierdie veeltalige oordragbenadering beter as alle vorige semantiese AWV metodes. Ons toon ook vir die eerste keer dat AWVe gebruik kan word vir semantiese soek-dmv-voorbeeld soektogte.

TABLE OF CONTENTS

Acknowledgements	i
Declaration	ii
Abstract	iii
Uittreksel	v
Table of contents	vii
1. Introduction	1
1.1. Motivation	1
1.2. Goals and methodology	2
1.2.1. Acoustic word embeddings	2
1.2.2. Acoustic word embeddings in a zero-resource setting	4
1.2.3. Downstream application	6
1.2.4. Semantic acoustic word embeddings	7
1.3. Contributions	8
1.3.1. Publications	9
1.4. Thesis overview	10
2. Background	12
2.1. Neural networks	12
2.1.1. Recurrent neural network	13
2.1.2. Autoencoder	14
2.1.3. Correspondence autoencoder	14
2.1.4. Siamese neural network	15
2.2. Speech processing	16
2.2.1. Speech features	16
2.2.2. Dynamic time warping	17
2.2.3. Query-by-example speech search	18
2.2.4. Unsupervised term discovery	19

TABLE OF CONTENTS

2.3.	Extracting acoustic word embeddings	20
2.4.	Unsupervised acoustic word embedding training strategies	21
2.4.1.	Unsupervised monolingual models	22
2.4.2.	Supervised multilingual models	22
2.4.3.	Unsupervised adaptation of multilingual models	22
2.5.	Evaluation of acoustic word embeddings	23
2.6.	Chapter summary	24
3.	Contrastive learning for acoustic word embeddings	25
3.1.	Related work	25
3.2.	Baseline acoustic word embedding models	27
3.2.1.	CAE-RNN	27
3.2.2.	SiameseRNN	28
3.3.	ContrastiveRNN	29
3.3.1.	Contrastive learning	29
3.3.2.	Contrastive learning in acoustic word embeddings	30
3.4.	Experimental setup	31
3.4.1.	Data	32
3.4.2.	Acoustic word embedding models	32
3.5.	Experiments	33
3.5.1.	Word discrimination	33
3.5.2.	Speaker classification	34
3.6.	Chapter summary	35
4.	Multilingual adaptation	36
4.1.	Intuition and related work	36
4.2.	Unsupervised adaptation of multilingual models	37
4.3.	Experimental setup	38
4.4.	Experiments	39
4.4.1.	Word discrimination	39
4.4.2.	Speaker classification	40
4.4.3.	Qualitative analysis	41
4.5.	Chapter summary	42
5.	Impact of language choice in a multilingual transfer setting	43
5.1.	Related work	43
5.2.	Languages of South Africa	45
5.3.	Query-by-example speech search	47
5.4.	Experimental setup	48
5.4.1.	Data	48

TABLE OF CONTENTS

5.4.2. Models	50
5.4.3. Evaluation	50
5.4.3.1. Word discrimination	51
5.4.3.2. Query-by-example	51
5.5. Experiments	52
5.5.1. Cross-lingual evaluation	52
5.5.2. Multilingual evaluation	54
5.5.3. Adding more languages	55
5.6. Chapter summary	57
6. Hate speech detection using multilingual acoustic word embeddings	59
6.1. Background	60
6.2. Related work	61
6.3. Overview of methodology	63
6.4. Low-resource keyword spotting	63
6.4.1. Keyword spotting with ASR	64
6.4.2. Keyword spotting with multilingual acoustic word embeddings . . .	65
6.5. Experimental setup	66
6.5.1. Data	66
6.5.1.1. Development	66
6.5.1.2. In-the-wild	67
6.5.2. Models	68
6.5.2.1. ASR model	68
6.5.2.2. Multilingual acoustic word embedding model	68
6.5.3. Evaluation	69
6.5.3.1. Controlled experiments	69
6.5.3.2. In-the-wild	70
6.6. Experiments	71
6.6.1. KWS results in a controlled environment	71
6.6.2. KWS results in the wild	72
6.7. Further analysis	74
6.7.1. XLSR vs MFCCs	74
6.7.2. Threshold	75
6.8. Chapter summary	76
7. Leveraging multilingual transfer for unsupervised semantic acoustic word embeddings	77
7.1. Background	78
7.2. Skip-gram	80
7.3. Related work	81

NOMENCLATURE

7.4.	Semantic AWEs trained from scratch	83
7.4.1.	Speech2Vec	83
7.4.2.	Semantic ContrastiveRNN	85
7.5.	Our approach: Semantic AWEs using multilingual transfer	85
7.5.1.	ContrastiveRNN with multilingual initialisation	85
7.5.2.	Projecting multilingual AWEs	86
7.5.3.	Cluster+Skip-gram	87
7.6.	Semantic query-by-example	90
7.7.	Setup	90
7.7.1.	Data	90
7.7.2.	Semantic models trained from scratch	91
7.7.3.	Semantic models using multilingual transfer	92
7.7.4.	Evaluation	93
7.7.4.1.	Intrinsic semantic AWE evaluation: Word similarity . . .	93
7.7.4.2.	Semantic QbE	95
7.8.	Experimental results	95
7.8.1.	Intrinsic evaluation: Word similarity	95
7.8.2.	Extrinsic evaluation: Semantic QbE	96
7.8.3.	Qualitative analysis	98
7.9.	Chapter summary	99
8.	Summary and conclusions	100
8.1.	Main findings	100
8.2.	Future work	102
	References	104

CHAPTER 1

INTRODUCTION

Over the last few years, we have seen great strides in the advancement of automatic speech recognition systems. Most state-of-the-art speech applications rely on neural networks with millions or even billions of parameters. However, with the increase in network sizes, the amount of required training data also increases.

Developing speech systems requires large amounts of transcribed speech data. For most languages, a sufficient amount of labelled speech data does not exist; some languages do not even have a written form [1]. To put this into perspective, Google Assistant only supports twelve languages¹; yet there are roughly 7 000 languages spoken throughout the world [2]. Clearly, we are in need of methods that rely on less transcribed speech data in order to accommodate the majority of under-resourced languages.

Most existing technologies rely on supervised training techniques. Collecting labelled data for under-resourced languages is time-consuming and expensive to the extent that it is unfeasible in many cases. Consequently, much work has been done on developing training strategies using speech data with no labels, known as the *zero-resource* setting [3]. In this setting, the goal is to develop methods that can discover linguistic structures and representations directly from unlabelled speech data [3, 4]. This research direction has strong links with studies on language acquisition in infants since infants also acquire their native language without explicit supervision [5].

1.1 MOTIVATION

Although full speech recognition is not possible in most zero-resource settings, researchers have proposed methods for applications such as speech search [6–8], word discovery [9–12], and segmentation and clustering [13–15]. Many of these applications require speech segments of different durations to be compared. This is conventionally done using alignment, for example with dynamic time warping (DTW), but this is computationally expensive and can be inaccurate [16].

¹<https://support.google.com/googlenest/answer/7550584>

1.2. GOALS AND METHODOLOGY

Acoustic word embeddings (AWEs) emerged as an alignment-free alternative for measuring the similarity between two speech segments [17]. AWEs are fixed-dimensional vector representations of variable-length speech segments where instances of the same word type should have similar embeddings. Given an accurate AWE model, the similarity between two speech segments can easily be determined by simply calculating the distance between their embeddings.

Currently, the most successful AWE approaches employ deep neural networks. For the zero-resource setting, many unsupervised AWE approaches have been explored, mainly relying on autoencoder-based neural models trained on unlabelled data in the target language [18–21]. However, there still exists a large performance gap between these unsupervised models and their supervised counterparts [17, 21], where word labels and word boundaries are available. A recent alternative for obtaining AWEs in a zero-resource language is to use multilingual transfer learning [22–25]. The goal is to have the benefits of supervised learning by training a model on labelled data from multiple well-resourced languages, but to then apply the model to an unseen target zero-resource language without fine-tuning it—a form of transductive transfer learning [26]. This multilingual transfer approach has been shown to outperform unsupervised monolingual AWE models [23].

The universal embeddings resulting from multilingual transfer are of great interest, as it allows fast and easy development of speech applications in a new language. In addition to the direct use of these embeddings, for example in search applications, these representations could also potentially be used as an additional signal in other upstream tasks.

1.2 GOALS AND METHODOLOGY

This thesis mainly focuses on recent advancements in AWE modelling following a multilingual transfer approach. In the sections that follow we highlight our new contributions in AWE-base modelling and their application in practical speech systems. But before we delve into the contributions, we first familiarise the reader with the fundamentals of AWE representations.

1.2.1 ACOUSTIC WORD EMBEDDINGS

To start with, we want to make a distinction between textual word embeddings (TWEs) and AWEs. TWEs are fixed-dimensional vector representations of written words. The use of TWEs are ubiquitous in modern natural language processing applications and research, where the first use of TWEs dates back to 1986 [27]. State-of-the-art word embedding methods [28, 29] exploit the co-occurrence statistics of words in a text corpora to learn word representations such that words with the same meaning have similar representations.

1.2. GOALS AND METHODOLOGY

Consequently, for TWEs, words that appear in a similar context will appear close to each other in vector space. Furthermore, every instance of the same word will be represented by only one embedding.

On the other hand, AWEs are fixed-dimensional vector representations of *spoken* words. AWEs are learned speech representations that capture the acoustic properties of spoken words. In contrast to TWEs, since every instance of a spoken word will be different due to the continuous nature of speech, every unique realisation of a word also has a unique vector representation. Ideally, instances of the same word type should appear close to each other in vector space and further away from instances of different word types.

In Figure 1.1 we illustrate the fundamental differences between TWEs and AWEs. In Figure 1.1a, multiple embeddings are displayed in a two-dimensional vector space, obtained through training a TWE model and then projecting the embeddings to two dimensions using principal components analysis. Note that each word type is assigned to a single embedding. Here we see embeddings of words that share contextual meaning are positioned close to each other. For example, “little” and “young” (orange), and “road” and “street” (purple).

Figure 1.1b shows the AWEs derived from the corresponding audio of the text data used for training the TWEs shown in Figure 1.1a. Here, spoken word segments are embedded such that each instance has a unique embedding that appears close to embeddings of other instances from the same word type. For example, all the acoustic realisations of “road” (dark purple) and “street” (light purple) are close to each other. We also see how phonetic

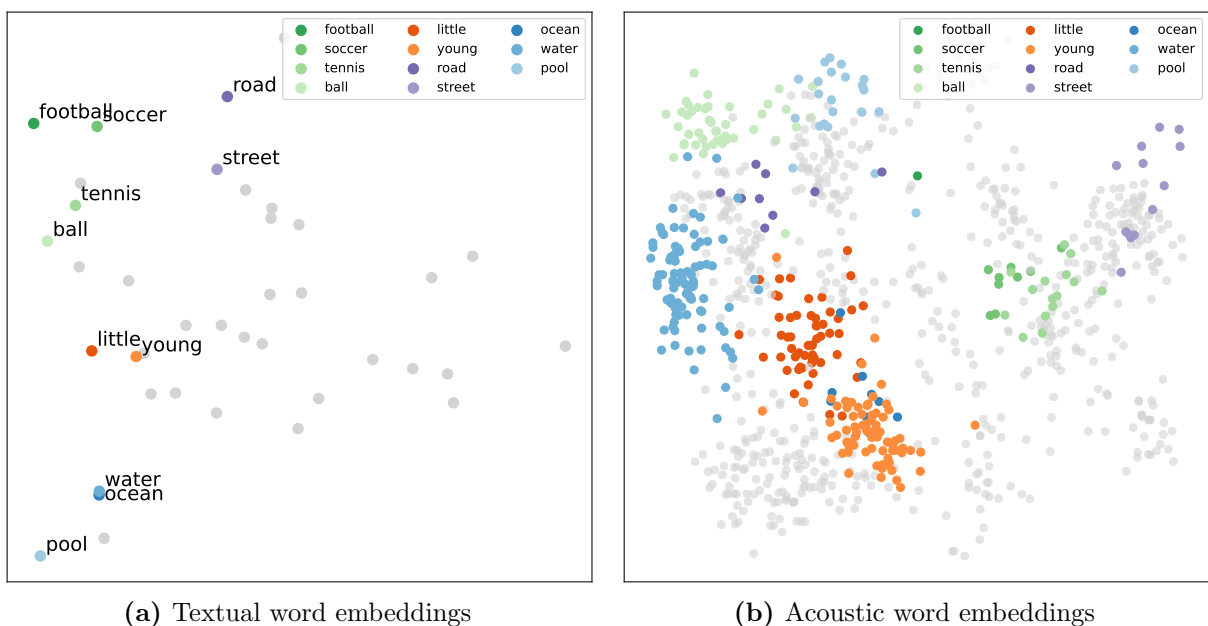


Figure 1.1: Illustrating the difference between (a) TWEs and (b) AWEs. In a TWE space, words sharing contextual meaning end up close to each other (different shades of the same colour). For AWEs, each spoken instance maps to a unique embedding (multiple of the same colour) with acoustically similar segments positioned close to each other.

similarities are captured, for example, instances of “ball” and “pool” are also close, given that they end with the same phone. It is clear that AWEs only store acoustic similarity and no information related to word meaning. For example, instances of “water”, “pool” and “ocean”, although semantically similar (Figure 1.1a), are spread out in the AWE space.

The main difference between AWEs and TWEs is summarised as follows: TWE models focus on learning embeddings that reflect the semantic relationship among words where instances of the same word are represented by the same embedding vector. Conversely, AWE models focus on learning embeddings that encapsulate the acoustic properties of spoken words where instances of the same spoken word are represented by unique (but similar) embeddings. Throughout this thesis we only consider AWEs (in Section 1.2.4 we introduce a different form of AWEs that share similarities with TWEs, but they are also derived from spoken word segments without using word labels).

1.2.2 ACOUSTIC WORD EMBEDDINGS IN A ZERO-RESOURCE SETTING

In this section we present an overview of AWE modelling for zero-resource languages. This is necessary to understand the contributions of this thesis, which we start to present at the end of this section.

In the zero-resource setting, we face the challenge of developing AWE models for languages without labelled training data. Two training strategies have been considered for AWE modelling in this scenario (more detail in Section 2.4). One approach is to use unlabelled audio data available in the target language [18–21]. In the absence of labelled data, training word pairs are obtained from an unsupervised term discovery (UTD) system, which automatically finds recurring word-like patterns in an unlabelled speech collection. Using discovered pairs from the target language enables the development of *unsupervised monolingual* AWE models.

A recent alternative for obtaining embeddings on a zero-resource language is to use multilingual transfer learning [22–25]. The idea is to train a *supervised multilingual* AWE model jointly on a number of well-resourced languages for which labelled data is available, but to then apply the model to an unseen zero-resource language. This multilingual transfer approach was found to outperform monolingual unsupervised learning approaches in [23–26, 30]

Various neural-based AWE models have been explored to learn latent representations from frame-level features. The most effective ones leverage recurrent neural networks (RNNs) [19, 21, 23, 25, 31, 32]. We briefly introduce two existing AWE models which have been successfully applied to both the unsupervised monolingual and supervised multilingual training strategies (more detail in Chapter 3): the *correspondence autoencoder* (CAE-RNN) [21] and SIAMESERNN [32]. The CAE-RNN optimise an autoencoder-like

reconstruction loss through an encoder-decoder RNN structure. Instead of reconstructing an input segment directly from the final encoder RNN hidden state, the CAE-RNN attempts to reconstruct another speech segment of the same type as the input. Unlike the reconstruction loss used in the CAE-RNN, the SIAMESERNN model explicitly optimises relative distances between embeddings [32]. The objective is to minimise the distance between the embedding of an input word and a word of the same type while at the same time maximising the distance between the input and an embedding from a different type [23, 32].

Based on the overview provided, we identify three key elements to consider for potentially improving AWEs in a zero-resource setting: the learning objective of the AWE model, the training strategy taking into account the scarcity of labelled data, and the choice of languages used during multilingual training. We now concretely lay out our approach to addressing these aspects.

AWE models and learning objectives. Recently, self-supervised contrastive learning has gained significant attention. This approach involves using proxy tasks to automatically obtain target labels from the data [33, 34]. Originally proposed for vision problems [35–37], it has since also been used as an effective pre-training step for supervised speech recognition [38–42]. A number of loss functions have been introduced in the context of self-supervised learning which have not been considered for AWEs. We specifically consider the contrastive loss of [43, 44] in a new AWE model which we call the CONTRASTIVERNN. While a Siamese AWE model [20, 32] optimises the relative distance between one positive and one negative pair, our contrastive AWE model jointly embeds a number of speech segments and then attempts to select a positive item among several negative items.

AWE training strategies. While a monolingual AWE model is designed to capture language-specific nuances, a multilingual AWE model learns universal linguistic properties that are invariant across languages. We now ask whether unsupervised learning and multilingual transfer are complementary. More specifically, can multilingual transfer further benefit from incorporating unsupervised learning? To answer this question we propose to adapt a multilingual AWE model to a target zero-resource language. This involves fine-tuning a pre-trained multilingual model’s parameters using discovered word pairs.

Language choice in multilingual transfer. Although there is a clear benefit in applying multilingual AWE models to an unseen zero-resource language, it is still unclear how the particular choice of training languages affects subsequent performance. A careful selection of training languages, based on language family, proved successful in language identification [45] and automatic speech recognition (ASR) [46, 47], but this has not been considered in AWE modelling. Preliminary experiments [23] show improved scores when training a monolingual model on one language and applying it to another from the

1.2. GOALS AND METHODOLOGY

same family. But this has not been investigated systematically and there are still several unanswered questions: Does the benefit of training on related languages diminish as we train on more languages (which might or might not come from the same family as the target zero-resource language)? When training exclusively on related languages, does performance suffer when adding an unrelated language? Should we prioritise data set size or language diversity when collecting data for multilingual AWE transfer? We address these questions by performing several experiments where we add data from different language families, and also control for the amount of data per language.

Above we established three aspects this thesis considers aimed at improving the quality of AWEs for zero-resource language. We now turn our attention to a practical downstream speech task that benefits from the application of these AWEs.

1.2.3 DOWNSTREAM APPLICATION

For zero-resource languages, speech applications are developed without the need for labelled training data. Among these applications is *query-by-example* (QbE), a speech retrieval task. In QbE, the goal is to use a spoken query to search an unlabelled audio corpus, aiming to retrieve utterances that contain instances of the query type.

DTW is typically used to match the speech features of a spoken query to the search utterances [9, 48–50]. However, DTW is slow and has some limitations [16, 51]. AWEs have been proposed as an alternative for matching the query segment and search segments by jointly mapping them to the same vector space [6, 8, 24, 52–54]. This allows fast comparison between the query segment and search segments. Although AWE-based QbE systems have proven successful in controlled experiments, there has been limited work investigating the effectiveness of these systems beyond the experimental environment, where training and testing data come from different domains.

As far as we know, only the works of Saeb et al. [55] and Menon et al. [56, 57, 58] consider speech retrieval performance on real in-the-wild data, applied to radio broadcast audio. These approaches rely on labelled data from the target language. Annotating audio (even in small quantities) is time-consuming and requires specialised linguistic knowledge which is not feasible when these systems need to be rapidly deployed.

In this thesis, we explore a real-world speech retrieval task of significant importance: hate speech detection in radio broadcasts [59] through keyword spotting (KWS). KWS typically relies on an ASR model trained on labelled data from the target language [60, 61]. To address this issue of label scarcity, we instead propose to extend QbE for KWS. To perform KWS with QbE only requires a small number of spoken templates to serve as queries for the keywords of interest [57, 62].

We approach this problem by first performing experiments in a controlled environment where training and test data come from the same domain. Then, we put our systems to a

1.2. GOALS AND METHODOLOGY

real-life test by evaluating performance on radio broadcast audio, specifically, Swahili audio (a low-resource language) scraped from radio stations in Kenya, a country in Sub-Saharan Africa.

1.2.4 SEMANTIC ACOUSTIC WORD EMBEDDINGS

In AWE modelling, the goal is to map variable-duration speech segments to fixed-dimensional vectors such that different acoustic realisations of the same word type have similar embeddings (Figure 1.1b). These embeddings have proven to be useful in applications requiring matching word segments from the same type (Section 1.2.3). The TWEs described in Section 1.2.1 have a distinct property where written words are mapped to similar representations if they share contextual meaning.

We now introduce a new form of representation learning called *semantic AWEs* that should uphold both the properties of the AWEs and TWEs. In semantic AWE modelling the goal would be to map speech segments to vector representations that not only capture whether two segments are instances of the same word, but also the semantic relationship between words. Figure 1.2 shows such a desired acoustic-semantic space. Here we see all instances of the same type end up close to each other (as for the AWEs in Figure 1.1b),




Figure 1.2: Semantic AWEs capture acoustic similarity among instances of the same word type (represented by the same colour) while also preserving word meaning (reflected in the shade of the same colour).

1.3. CONTRIBUTIONS

for example “road” (dark purple) and “street” (light purple). But now, the instances of “road” and “street” also end up close to each other (as the TWEs in Figure 1.1a).

Only a limited number of studies have tried to address the problem of learning semantic AWEs from unlabelled speech alone [63, 64]. These methodologies centre around training exclusively on unlabelled audio data from the target language. In doing so, a model is challenged to distinguish between acoustic properties related to word type while simultaneously untangling word semantics.

To overcome this challenge, we propose leveraging the recent improvements in multilingual transfer in AWEs. Specifically, we propose using a pre-trained multilingual AWE model as an additional signal to improve semantic AWEs in a target language where we only have unlabelled speech. Since the multilingual model already captures acoustic properties, this should simplify the semantic learning problem. In this thesis, we introduce three new semantic AWE models that incorporate multilingual AWEs.

1.3 CONTRIBUTIONS

To take the above together, this thesis makes the following specific contributions to advancing zero-resource speech application via AWEs:

- We propose a new AWE model, the CONTRASTIVERNN. The model architecture is similar to existing AWE models but optimises a different objective function, not previously considered for AWE modelling. We show this model outperforms existing AWE models in the unsupervised monolingual training strategy on six evaluation languages; the CONTRASTIVERNN performs on par with existing models in the multilingual transfer setting.
- To our knowledge, we present the first unsupervised adaptation of multilingual AWE models. Previous work performed adaptation of multilingual AWE models using speech segments with class labels [25]. Instead of using true word segments, we use unknown word-like pairs. These are obtained from applying a UTD system—itself unsupervised—to an unlabelled speech corpus in the target language. The discovered word pairs are then used to fine-tune the multilingual model to the target language. We show that unsupervised adaptation is beneficial, with the new model we propose (CONTRASTIVERNN) giving the best performance after adaptation.
- To our knowledge, we are the first to perform extensive analysis on the choice of training languages in the multilingual transfer setting, specifically for AWEs. We experiment with different training languages and amounts of data on real low-resource languages. We show the benefit of using languages that belong to the same language family as the target zero-resource language when training a multilingual AWE model.

1.3. CONTRIBUTIONS

- To our knowledge, we are the first to develop a keyword spotting system for hate speech detection in radio broadcasts using multilingual AWEs. We perform our main experiment using real broadcast in a low-resource language, Swahili, scraped from radio stations in Kenya. We show our AWE-based KWS system is more robust to a domain mismatch compared to existing ASR-based KWS systems, especially when large amounts of labelled data from the target language are unavailable.
- For the first time, we propose leveraging multilingual AWEs to learn representations of whole-word speech segments that capture the meaning of words (rather than acoustic properties) in a low-resource setting. Only a handful of studies have looked at this problem [63, 64]. We introduce three new semantic models, with one of them showing large improvements over previous approaches. We are also the first to apply these semantic acoustic embeddings in a downstream semantic speech retrieval task.

1.3.1 PUBLICATIONS

The contributions above is summarised in the following publications.

Research paper 1: (Conference)

C. Jacobs, Y. Matuselych, and H. Kamper, “Acoustic word embeddings for zero-resource languages using self-supervised contrastive learning and multilingual adaptation,” *IEEE Spoken Language Technology (SLT)*, 2021.

Research paper 2: (Conference)

C. Jacobs and H. Kamper, “Multilingual transfer of acoustic word embeddings improves when training on languages related to the target zero-resource language,” *Interspeech*, 2021.

Research paper 3: (Conference)

C. Jacobs, N. C. Rakotonirina, E. A. Chimoto, B. A. Bassett, and H. Kamper, “Towards hate speech detection in low-resource languages: Comparing ASR to acoustic word embeddings on Wolof and Swahili,” *Interspeech*, 2023.

Research paper 4: (Submitted Journal)

C. Jacobs, and H. Kamper, “Leveraging multilingual transfer for unsupervised semantic acoustic word embeddings,” submitted to *IEEE Signal Processing Letters*, 2023.

1.4 THESIS OVERVIEW

Chapter 2: Background. The thesis starts by familiarising the reader with concepts to follow in subsequent chapters. We describe various neural network models relevant to producing AWEs. The processing of speech data is described. A few zero-resource speech applications relevant to AWEs are introduced to the reader. We provide a detailed description of AWE approaches in the zero-resource setting, including the training strategies we consider in this thesis. Lastly, we discuss how we measure the quality of AWEs in both intrinsic and extrinsic evaluation tasks.

Chapter 3: Contrastive learning for acoustic word embeddings. This chapter presents a detailed description of existing AWE models that we reimplement as baselines for the rest of the thesis. We then introduce a new embedding model we called the CONTRASTIVERNN. We compare the CONTRASTIVERNN in both the unsupervised monolingual and supervised multilingual training strategies. Following the unsupervised monolingual approach, the CONTRASTIVERNN outperforms the existing models on all six evaluation languages, with an absolute increase in average precision ranging from 3.3% to 17.8%. In the multilingual transfer setting, the CONTRASTIVERNN shows a marginal increase compared to existing models, mostly performing on par.

Chapter 4: Multilingual adaptation. This chapter introduces the unsupervised adaptation of multilingual AWE models, a new training strategy. A multilingual AWE model is trained using labelled data from multiple-well-resourced languages; the model is then fine-tuned using unlabelled data from the target zero-resource language before applying it. We perform unsupervised adaptation on all the AWE models considered in Chapter 3. The CONTRASTIVERNN model shows the highest performance increase after adaptation with roughly up to 5% absolute increase in average precision on five out of six evaluation languages.

Chapter 5: Impact of language choice in a multilingual transfer setting. This chapter investigates the impact the choice of training languages have in a transfer learning setting. Here we assume a model trained on one set of languages will perform differently from a model trained on another set of languages. We specifically consider the effect of training on languages that belong to the same language family as the target language. We perform multiple experiments using data from South African languages. We show the benefit of including related languages in multilingual modelling, even in small quantities. This chapter also provides advice for practitioners who want to develop speech applications using AWEs. Here we also apply multilingual models in a downstream query-by-example speech search task.

Chapter 6: Hate speech detection using multilingual acoustic word embeddings. This chapter addresses the problem of hate speech detection in low-resource languages through keyword spotting (KWS). Our goal is to compare existing ASR-based KWS systems to a multilingual AWE-based approach on real radio broadcast audio. We first develop and test our systems using experimental datasets: we use two low-resource languages, Swahili and Wolof. In this in-domain setting, an ASR model fine-tuned using as little as five minutes of labelled data outperforms the AWE-based KWS system. However, in a real-life scenario, when applying the systems to real radio broadcasts, the AWE system proves to be more robust by almost reaching the performance of an ASR model fine-tuned on 30 hours of data.

Chapter 7: Leveraging multilingual transfer for unsupervised semantic acoustic word embeddings. This chapter introduces a novel approach to produce semantic AWEs. Specifically, we propose using multilingual AWEs to assist in this learning task. Our best semantic AWE approach involves clustering word segments using the multilingual AWE model, deriving soft pseudo-word labels from the cluster centroids, and then training a classifier model on the soft vectors. In an intrinsic word similarity task measuring semantics, our multilingual transfer approach for semantic modelling outperforms all previous semantic AWE methods. We also apply these semantic AWEs to a downstream semantic query-by-example search.

Chapter 8: Summary and conclusion. This chapter highlights the main findings of this thesis and provides recommendations for future work.

CHAPTER 2

BACKGROUND

This chapter introduces concepts the reader needs to be familiar with to follow subsequent chapters. We give some background information on neural networks relevant to AWE modelling. We describe the processing of speech data and different speech applications relevant to a zero-resource setting. We then go on to describe different AWE models and how they are trained, and explain how we evaluate AWEs.

2.1 NEURAL NETWORKS

In this section, we provide the reader with the necessary neural network fundamentals to follow subsequent chapters.

Feedforward neural networks (FFNNs) and convolutional neural networks (CNNs) are suitable for fixed-length inputs to create fixed-length outputs. Recurrent neural networks (RNNs) are appropriate for handling variable-length sequences of inputs to produce fixed- or variable-length outputs. Settle and Livescu [32] showed that RNNs outperform FFNNs and CNNs in AWE modelling (more on this in Section 2.3); we therefore only consider AWE approaches built on RNNs.

We use the notation $g_{\Phi}(\mathbf{x})$ to represent a network where Φ are trainable weight parameters and \mathbf{x} is an input to the network. We use the notation $f_{\Theta}(\mathbf{x})$ to represent an entire model which can be composed of multiple network types for example:

$$f_{\Theta_{Model}}(\mathbf{x}) = g_{\Phi_{FFNN}}(g_{\Phi_{RNN}}(\mathbf{x})), \quad (2.1)$$

where Θ is the trainable parameters of all sub-networks. In the following sections, we briefly explain the behaviour of RNNs and two network configurations used for AWE modelling.

2.1. NEURAL NETWORKS




Figure 2.1: The structure of an RNN with one recurrent layer.

2.1.1 RECURRENT NEURAL NETWORK


Recurrent neural networks (RNNs) are suitable for processing sequential data. Unlike FFNNs and CNNs, RNNs contain feedback loops, allowing the network to predict an output using information from previous inputs and the current input [65].

The network consists of one or more recurrent layers that share its parameters throughout. Each recurrent layer has trainable weight parameters $\mathbf{W}_1 \in \mathbb{R}^{(N \times N)}$, $\mathbf{W}_2 \in \mathbb{R}^{(N \times M)}$ and $\mathbf{b} \in \mathbb{R}^{(N \times 1)}$. Given an input sequence $\mathbf{u} = (\mathbf{u}_1, \mathbf{u}_2, \dots, \mathbf{u}_\eta)$, the recurrent layer produces an output sequence $\mathbf{v} = (\mathbf{v}_1, \mathbf{v}_2, \dots, \mathbf{v}_\eta)$. For a single input $\mathbf{u}_i \in \mathbb{R}^M$ the recurrent layer produces an output $\mathbf{v}_i \in \mathbb{R}^N$ using the following recursive function:

$$\mathbf{v}_i = \sigma(\mathbf{W}_2 \mathbf{u}_i + \mathbf{W}_1 \mathbf{v}_{i-1} + \mathbf{b}), \quad (2.2)$$

where σ is some activation function and \mathbf{b} a bias parameter. The network structure of an RNN is illustrated in Figure 2.1. The weight parameter \mathbf{W}_1 is updated to preserve relevant information from the initial input \mathbf{u}_1 that is relevant to produce output \mathbf{v}_i with current input \mathbf{u}_i . Weight parameter \mathbf{W}_2 filters the relevant information from input \mathbf{u}_i to produce \mathbf{v}_i . These weight parameters control the state of the network over time depending on the inputs.

In practice RNNs do not learn long-term dependencies [66] due to exploding or vanishing gradients [67]. A special type of recurrent unit, long short-term memory (LSTM), has been introduced to combat these difficulties [68].

**Figure 2.2:** The AE and CAE structure.

2.1.2 AUTOENCODER

An autoencoder (AE) neural network is a model that has a target output identical to the input [69]. The idea is that the model should learn a lower-dimensional representation of a given input vector, containing enough meaningful information, that would allow the model to reconstruct the input. This model is useful for unsupervised learning since no knowledge regarding the input is needed. The AE model consists of an encoder and decoder structure. Given an input \mathbf{x} , the model's encoder $f_{\Phi_{ENC}}(\cdot)$ processes the input to produce a latent embedding \mathbf{z} . The latent embedding is then given as input to the model's decoder $f_{\Phi_{DEC}}(\cdot)$ to produce an output $\hat{\mathbf{y}}$, illustrated in Figure 2.2

The trainable parameters of the encoder and decoder together create the AE model $f_{\Theta_{AE}}(\cdot)$. With input \mathbf{x} , the model $f_{\Theta_{AE}}(\cdot)$ is trained to minimise the reconstruction loss between the desired output \mathbf{y} and the models output $\hat{\mathbf{y}}$:

$$\begin{aligned} \ell(\mathbf{x}, \mathbf{y}) &= \|\mathbf{y} - f_{\Theta_{AE}}(\mathbf{x})\|^2 \\ &= \|\mathbf{y} - \hat{\mathbf{y}}\|^2 \end{aligned} \quad (2.3)$$


For the AE the target output \mathbf{y} is equal to the input \mathbf{x} , the loss function can be written as:

$$\begin{aligned} \ell(\mathbf{x}) &= \|\mathbf{x} - f_{\Theta_{AE}}(\mathbf{x})\|^2 \\ &= \|\mathbf{x} - \hat{\mathbf{y}}\|^2 \end{aligned} \quad (2.4)$$

2.1.3 CORRESPONDENCE AUTOENCODER

The correspondence autoencoder (CAE) network structure is identical to the AE as illustrated in Figure 2.2. The only difference is that the target output of the CAE is

2.1. NEURAL NETWORKS

**Figure 2.3:** The Siamese network structure.

not exactly the input. Instead, given an input the CAE tries to produce an output that belongs to the same class as the input [70].

The CAE loss is given by setting the target output \mathbf{y} in Equation 2.3 to \mathbf{x}_p :

$$\begin{aligned} \ell(\mathbf{x}, \mathbf{x}_p) &= \|\mathbf{x}_p - f_{\Theta_{CAE}}(\mathbf{x})\|^2 \\ &= \|\mathbf{x}_p - \hat{\mathbf{y}}\|^2, \end{aligned} \quad (2.5)$$

where \mathbf{x}_p represents an instance from the the same class as input \mathbf{x} .

2.1.4 SIAMESE NEURAL NETWORK

Instead of reconstructing an input representation, the Siamese network measures the distance between representations of two inputs [71]. The network consists of two identical sub-networks with tied parameters. The model is trained to optimise the relative distance between representations of two data instances by mirroring parameter updates across both sub-networks. Formally, the model $f_{\Theta_{SIA}}(.)$ produces representations \mathbf{z}_1 and \mathbf{z}_2 from inputs \mathbf{x}_1 and \mathbf{x}_2 respectively, illustrated in Figure 2.3. Ideally, embeddings \mathbf{z}_1 and \mathbf{z}_2 from same class instances \mathbf{x}_1 and \mathbf{x}_2 should be similar, and embeddings \mathbf{z}_1 and \mathbf{z}_2 from different class instances \mathbf{x}_1 and \mathbf{x}_2 should be unlike.

In other words, embeddings from the same class should be close to each other in vector space and far away from embeddings from other classes. By introducing negative samples during training we can better accomplish this goal [72–75]. For a single training instance the model needs to minimise the distance between an input \mathbf{x}_a and another instance from the same class \mathbf{x}_p while simultaneously maximising the distance between the input \mathbf{x}_a and an instance from a different class \mathbf{x}_n (subscripts refer to anchor, positive and negative) as illustrated in Figure 2.4. To optimise the relative distance, the Siamese network optimises a triplet hinge-like loss:

$$l_{\text{triplet}}(\mathbf{z}_a, \mathbf{z}_p, \mathbf{z}_n) = \max\{0, m + d(\mathbf{z}_a, \mathbf{z}_p) - d(\mathbf{z}_a, \mathbf{z}_n)\}, \quad (2.6)$$

2.2. SPEECH PROCESSING




Figure 2.4: Visualisation of the relative distance between embeddings (a) before training and (b) after training, showing how representations from the same class ($\mathbf{z}_a, \mathbf{z}_p$) are pulled together and pushed away from a different class \mathbf{z}_n by optimising the triplet loss.

where

$$d(\mathbf{z}_1, \mathbf{z}_2) = 1 - \frac{\mathbf{z}_1 \cdot \mathbf{z}_2}{\|\mathbf{z}_1\| \|\mathbf{z}_2\|} \quad (2.7)$$

is the cosine distance between embeddings \mathbf{z}_1 and \mathbf{z}_2 extracted from inputs \mathbf{x}_1 and \mathbf{x}_2 with m a margin parameter. The triplet hinge-like loss is at a minimum when all embedding pairs ($\mathbf{z}_a, \mathbf{z}_p$) of the same type are a distance m closer than embedding pairs ($\mathbf{z}_a, \mathbf{z}_n$) of different types. Various studies showed the benefit of optimising the relative distances instead of only minimising the absolute distance between data instances from the same class [72, 73].

2.2 SPEECH PROCESSING

We now describe the feature representation of speech segments that are presented to the neural networks (described above in Section 2.1) as inputs and outputs. We then go on to describe speech tasks relevant to AWEs in a zero-resource setting.

2.2.1 SPEECH FEATURES

Throughout this thesis we use mel-frequency cepstrum coefficients (MFCCs) to make a reliable comparative study between AWE models considered in previous work. All speech segments are parametrised as 13-dimensional MFCCs using a window size of 25 ms and a 10 ms frame shift. Additionally, speaker normalisation is performed per utterance.

All AWE models are trained on MFCC features from isolated word segments. In the supervised setting, word boundaries are obtained from forced alignments. In the unsupervised setting, we use an unsupervised term discovery system (Section 2.2.4) to obtain unknown word-like segments from unlabelled data in the target language.

2.2. SPEECH PROCESSING




Figure 2.5: Illustration of the mapping between two signals of different durations that yield the lowest alignment cost using DTW.

For dynamic time warping (described below in Section 2.2.2) baseline experiments, delta and double-delta MFCCs are included.

Throughout the course of this thesis, significant advancements have been made in the field of self-supervised speech representation learning (SSRL) using transformer networks [76–79]. The primary objective of SSRL is to train a network using unlabelled speech data to generate frame-level representations that offer more contextual information compared to traditional rule-based MFCCs. In our experiments, we use the Wav2Vec2.0 XLSR model, which has been trained on unlabelled audio from multiple languages [77]. Using the XLSR model, we extract speech features from 25 ms of speech with a frame shift of 20 ms. We use the output of the 12th transformer layer producing input features of 1024 dimensions. From Chapter 6 we include these features in our experiments.

2.2.2 DYNAMIC TIME WARPING

Dynamic time warping (DTW) is used to calculate the dissimilarity between two temporal sequences. This method is closely related to the *edit distance* used in natural language processing. Edit distance is a method used to quantify the difference between two strings by identifying the minimum number of operations required to transform one of the strings into the other.¹ Assume there are three possible operations to transform one string into the other: insertion, deletion and substitution. Now consider the following two words, “time” and “camel”. To transform “time” into “camel” we need to substitute the ‘t’ with a ‘c’, substitute the ‘i’ with an ‘a’ and insert a ‘l’ at the end. This transformation requires three operations, therefore the edit distance between the two words is three.

DTW is used to calculate the dissimilarity between two temporal sequences similar in the way edit distance is used to calculate the dissimilarity between strings. DTW is an algorithm that finds the optimal alignment between two time series through a dynamic programming method. Consider the example signals in Figure 2.5 (a). DTW can be used to calculate how similar the blue signal is to the red signal by calculating the distance between each point from the blue signal to every point in the red signal. This results in

¹https://en.wikipedia.org/wiki/Edit_distance

2.2. SPEECH PROCESSING




Figure 2.6: Illustration of the optimal alignment of MFCCs extracted from two speech segments using DTW.

a matrix of distances known as a cost matrix from which we can identify a particular mapping of points from the blue signal to the red signal that results in the lowest alignment cost. The mapping of points from the blue signal to points in the red signal that yield the lowest alignment cost is displayed in Figure 2.5 (b).

For speech signals, DTW can be applied to a set of extracted speech features for example MFCCs. Similarity between two “points” is calculated as some distance between the MFCC vectors such as the cosine distance. Figure 2.6 shows an example of the mapping of one speech segment to another using extracted MFCC frames.

DTW is found in most state-of-the-art unsupervised term discovery (Section 2.2.4) and query-by-example (Section 2.2.3) systems as described below. AWEs were introduced as an alignment-free method to measure the similarity between word segments operating at linear time complexity compared to the polynomial time complexity of DTW. We use DTW as a baseline to compare AWEs in a word discrimination task (Section 2.5) and in a query-by-example speech search (Section 2.2.3) task.

2.2.3 QUERY-BY-EXAMPLE SPEECH SEARCH

Query-by-example (QbE) speech search is the task of retrieving utterances from a speech collection related to a given a spoken query. Early QbE systems implemented with large vocabulary continuous speech recognition (LVCR) systems proved to be successful [80, 81]. However, training an LVCR requires large amounts of labelled data and has limitations to handling out-of-vocabulary words. Studies focussing on implementing systems following a phonetic approach alleviate these impediments [82, 83]. These systems still rely on language-specific knowledge and labelled data which is impractical for the zero-resource setting.

Recently, several studies focused on developing QbE speech search systems, specifically for the zero-resource setting [84–88]. In the zero-resource setting, a typical approach is to learn efficient frame-level feature representations for both the query and all possible

segments from the speech collection then perform DTW to find matching representations [48, 49]. However, DTW is computationally expensive, restricting QbE on large-scale speech collections. Several studies considered less expensive DTW QbE alignment methods [89, 90]. Moreover, Jansen van Durme [50] introduced a QbE system where raw speech frames QbE hashed to bit signatures followed by applying a QbE nearest neighbour similarity search algorithm, reducing the computational time to logarithmic time.

Levin et al. [6] introduced AWEs to replace the frame-based representations used to perform QbE in [50], to alternatively use whole-word segments representations, using Laplacian eigenmaps. This approach showed a reduction in computation time and improved accuracies given that the AWEs have better lexical discrimination (more in Section 2.3).

Settle et al. [52] replaced the Laplacian eigenmaps with deep neural network AWEs, specifically, they implement a SIAMSEERNN (Section 2.1.4) AWE model to perform QbE search which showed significant improvement. Recently, several studies proposed performing QbE using AWEs with a neural approach [8, 24, 53].

In Chapter 5 we discuss the implementation of a simplified QbE approach used by Kamper et al. [91] for measuring AWE performance in a downstream application.

2.2.4 UNSUPERVISED TERM DISCOVERY

The goal of QbE is to retrieve one or more utterances within an unlabelled speech collection given a spoken query. Unsupervised term discovery (UTD) instead aims to extract and group unknown word-like segments from an unlabelled speech collection (i.e. no query is given). Throughout this thesis, we use UTD to find word-like speech segments from an unlabelled speech collection that allows us to train AWE models in a fully unsupervised way.

Park and Glass [9] were the first to introduce a UTD system. The UTD system aims to identify repeating subsequences within a speech signal using only the signal itself. The system of Park and Glass tries to find recurring speech patterns by applying a segmental DTW algorithm to audio data. The segmental DTW algorithm tries to find subsequence alignments between the acoustic features of pairs of continuous spoken utterances.

Jansen and van Durme [10] presented a UTD system that reduces the computation complexity from quadratic in [9] to linearithmic. They did this by applying a random projection algorithm that maps sequences in the acoustic feature space to fixed-length bit signatures, allowing efficient similarity calculation through an approximate nearest neighbour search algorithm.

This task is of great interest as it allows matching speech segments predicted to be of the same word type from untranscribed speech, which is useful in the zero-resource setting. Many unsupervised AWEs have been trained using word pairs obtained from UTD systems [21, 23, 30].

2.3 EXTRACTING ACOUSTIC WORD EMBEDDINGS

We now explore previous work addressing AWEs, focussing on zero-resource languages for which transcribed audio data is not available.

Levin et al. [17] were the first to propose representing whole word speech segments as fixed-dimensional vectors that contain linguistic meaning. They explored various embedding techniques. For each, different levels of information are assumed to be available. In a setting where the only available data is unknown isolated speech segments (word boundaries are available), they perform downsampling on extracted speech features to create fixed-dimensional representations. They considered uniform and non-uniform downsampling (using hidden Markov models), both methods yielding fixed-dimensional embeddings. Using downsampling as a means to create AWEs was also used as a baseline for comparing more sophisticated approaches considered in [18, 21].

Another approach Levin et al. [17] presented involves reference vectors and Laplacian eigenmaps. Assuming a reference set of distinct unknown speech segments is available, the DTW cost between a speech segment and each reference segment is concatenated to create a reference vector. The reference set should contain enough segments to form a basis for all possible speech segments, consequently, producing very high-dimensional reference vectors. They apply a linear dimensionality reduction technique, principal component analysis, to reduce the dimensionality of the reference vectors, showing a minimal decrease in performance. They also proposed an AWE approach using a non-linear graph embedding technique. They evaluated all their methods on a word discrimination task where only the Laplacian eigenmaps approach reached the performance of DTW. Although, some of these unsupervised approaches were successfully implemented in [13, 52, 92, 93] they are still far from reaching the performance of the supervised approaches (also considered by Levin et al.) where word labels are available.

Many studies followed, aiming to improve AWEs in the zero-resource setting, mostly using deep neural networks. Kamper et al. [20] follow a CNN-based approach that makes use of weak supervision. Here speech segments are grouped into unknown word types using available word class labels. They use these word pairs to train an AWE model called the SIAMESECNN using a hinge-like contrastive loss function, where the model tries to minimise the distance between the unknown word pairs while maximising the distance between words from different unknown types. This loss is described in Section 2.1. They show that the SIAMESECNN outperforms the Laplacian eigenmaps approach of Levin et al. [17] and, more importantly, the SIAMESECNN, trained with weak supervision, performs on par with a supervised classification model (CLASSIFIERCNN) where word labels are available.

Chung et al. [19] were the first to apply an RNN to produce fixed-dimensional representations from variable-duration speech segments without any supervision. They present

a sequence-to-sequence autoencoder AWE model. An encoder RNN processes an input sequence, a decoder RNN then tries to reconstruct the input sequence taking the final layer of the encoder RNN as input. This reconstruction loss is described in Section 2.1. After training the model, the final hidden layer of the encoder RNN is taken as the embedding of an input sequence. They did not explicitly compare embedding quality to previous methods, however, they show applying the autoencoder AWE outperforms DTW in a QbE speech search task.

Settle and Livescu [32] reimplemented the CNN-based models of Kamper et al. [20] using RNNs. In a similar experimental setup as Kamper et al. [20] they show that the Siamese network build on RNNs (SIAMESERNN) outperform the SIAMESECNN of Kamper et al. [20].

Finally, Kamper [21] propose a method for embeddings in a truly zero-resource setting. In this unsupervised setting, no labelled speech data or word boundaries are available. A UTD (Section 2.2.4) system—itself unsupervised—is used to find word-like pairs predicted to be of the same unknown type. These pairs are presented to an autoencoder-like network where the target to reconstruct is not identical to the input (like the autoencoder from Chung et al. [19]) but rather an instance predicted to be of the same type. This model, the correspondence autoencoder RNN (CAE-RNN), outperforms all previous unsupervised AWE approaches and delivers similar results as DTW executing at a much lower run-time.

Although these unsupervised methods are useful for the zero-resource setting, there still exists a large performance gap between the supervised AWE methods [17, 21]. Recently, Kamper et al. [23] exploits labelled data from multiple well-resourced languages to train a single supervised multilingual AWE model and then apply the model to an unseen zero-resource language. In this transfer learning setting, they train multilingual versions for three existing models, CLASSIFIERRNN [32], SIAMESERNN [32] and CAE-RNN [21]. They compare these multilingual models to their monolingual unsupervised counterparts, where the multilingual models outperform all of them. Others have since also considered a multilingual approach to AWE modelling [24, 25, 94]. The multilingual transfer approach is currently the best approach. However, the performance compared to the supervised models still lack by a large margin.

In the following section we summarise some of the training strategies identified in this section.

2.4 UNSUPERVISED ACOUSTIC WORD EMBEDDING TRAINING STRATEGIES

In this section we explicitly mention the training strategies used for unsupervised AWE modelling that we use in subsequent chapters. Firstly, we summarise two training strategies

identified in the preceding section. One option is to train unsupervised monolingual models directly on unlabelled data (Section 2.4.1), introduced by Kamper [21]. Another option is to train a supervised multilingual model on labelled data from well-resourced languages and then apply the model to a zero-resource language (Section 2.4.2), introduced by Kamper et al. [22]. In Section 2.4.3 we describe a new AWE training strategy, which we include here in the background chapter for the sake of completeness—we really propose this new setting only in Chapter 4. In this strategy, multilingual models are fine-tuned to a zero-resource language using unsupervised adaptation. All the AWE models considered throughout this thesis can be used in all three settings, as explained below.

2.4.1 UNSUPERVISED MONOLINGUAL MODELS

For any of the AWE models in Chapter 3, we need pairs of segments containing words of the same type; for the SIAMESE-RNN and CONTRASTIVE-RNN we additionally need negative examples. In a zero-resource setting there is no transcribed speech to construct such pairs. But pairs can be obtained automatically by applying a UTD (Section 2.2.4) system to an unlabelled speech collection from the target zero-resource language. This system discovers pairs of word-like segments, predicted to be of the same unknown type.

The discovered pairs can be used to sample positive and negative examples for any of the three models in Chapter 3. Since the UTD system has no prior knowledge of the language or word boundaries within the unlabelled speech data, the entire process can be considered unsupervised. Using this methodology, we consider purely unsupervised monolingual versions of each AWE model in Chapter 3.

2.4.2 SUPERVISED MULTILINGUAL MODELS

Instead of relying on discovered words from the target zero-resource language, we can exploit labelled data from well-resourced languages to train a single multilingual supervised AWE model [23, 25]. This model can then be applied to an unseen zero-resource language. Since a supervised model is trained for one task and applied to another, this can be seen as a form of *transfer learning* [26, 95]

Experiments in [23] showed that multilingual versions of the CAE-RNN and SIAMESE-RNN outperform unsupervised monolingual variants. A multilingual CONTRASTIVE-RNN hasn't been considered in a previous study, as far as we know. We consider supervised multilingual variants of all three models in Chapter 3.

2.4.3 UNSUPERVISED ADAPTATION OF MULTILINGUAL MODELS

While previous studies have found that multilingual AWE models (Section 2.4.2) are superior to unsupervised AWE models (Section 2.4.1), one question is whether multilingual

models could be tailored to a particular zero-resource language in an unsupervised way.

We propose to adapt a multilingual AWE model to a target zero-resource language by fine-tuning the multilingual model’s parameters using discovered pairs from the target zero-resource language. These discovered segments are obtained by applying a UTD system to unlabelled data from the target zero-resource language. The idea is that adapting the multilingual AWE model to the target language would allow the model to learn aspects unique to that language. We consider the adaptation of multilingual versions of all three AWE models in Chapter 4.

As far as we know, we are the first to perform unsupervised adaptation of multilingual AWE models for the zero-resource setting.

2.5 EVALUATION OF ACOUSTIC WORD EMBEDDINGS

Ultimately, we want to evaluate the performance of AWEs in downstream speech applications. However, we would like to measure the quality of the embeddings without being tied to a specific system architecture for easy and fast comparison between different embedding approaches. We therefore use a word discrimination task that measures the intrinsic quality of the embeddings. The *same-different* task, introduced by Carlin et al. [96], is specifically designed to quantify the separability of same-word and different-word representations across a range of word types and speakers.

In the same-different task, we are given a pair of acoustic segments, each a true word, and we must decide whether the segments are examples of the same or different words. To evaluate a particular embedding method, a set of isolated test words are first embedded. Two words can then be classified as being of the same or different type based on some distance threshold, and a precision-recall curve is obtained by varying the threshold. The area under this curve is used as the final evaluation metric, referred to as the average precision (AP). Carlin et al. [96] show that this AP correlates with the phone recognition accuracy in a spoken term discovery system, indicating its potential for predicting performance in downstream tasks.

In our implementation, we use a cosine distance to measure the similarity between two vectors that Levin et al. [17] found to generally perform better than Euclidean distance. Concretely, let’s consider two test word embeddings, \mathbf{z}_i and \mathbf{z}_j . Two words are declared to be the same if their cosine distance is less or equal to some distance threshold τ , as follows:

$$1 - \frac{\mathbf{z}_i \cdot \mathbf{z}_j}{\|\mathbf{z}_i\| \|\mathbf{z}_j\|} \leq \tau \quad (2.8)$$

Moreover, we are particularly interested in obtaining embeddings that are speaker-invariant. We therefore calculate AP by only taking the recall over instances of the same word spoken by different speakers. In a test set of all possible word pairs, there exist a

2.6. CHAPTER SUMMARY

subset C_{SWDP} containing pairs of words that are the same but said by different speakers (SWDP). If the number of word pairs declared to be the same based on a threshold τ from this subset is $N_{SWDP}(\tau)$, we calculate the recall as follows:

$$R_{SWDP}(\tau) = \frac{N_{SWDP}(\tau)}{|C_{SWDP}|} \quad (2.9)$$

In other words, we consider the more difficult setting where a model does not get credit for recalling the same word if it is said by the same speaker.

However, Algayres et al. [97] have shown that this same-different evaluation task is not always indicative of downstream system performance specifically for AWEs. Here they compared the results of the intrinsic word-discrimination task against a downstream task: the unsupervised estimation of frequencies of speech segments in a given corpus. In general, they found the AP to correlate with frequency estimation. However, some inconsistencies appear in fine-grained distinctions across embedding approaches and languages. Moreover, Abdullah et al. [98] analyse the correlation between the word discrimination task and word phonological similarity. In their experiments, they show that AWEs optimising contrastive objectives (such as the SIAMESE-RNN and CONTRASTIVE-RNN) yield strong discriminative embeddings but for the most part, fail to reflect the phonological distance between word forms.

Given these inconsistencies in the evaluation of the quality of AWEs, we perform QbE speech search (Section 2.2.3) as an additional evaluation of AWEs. In contrast to the word discrimination task, this test does not assume a set of isolated words, but instead operates on full unsegmented utterances.

2.6 CHAPTER SUMMARY

We looked at AWEs for zero-resource languages and identified the multilingual transfer strategy as state-of-the-art, specifically when implemented using the CAE-RNN and SIAMESE-RNN models of [21] and [32], respectively. We introduced two zero-resource speech applications, QbE and UTD. We use QbE as a downstream evaluation task to evaluate the quality of AWEs in Chapter 5, 6, and 7 (semantic QbE is performed and explained in the latter). We use UTD to obtain word pairs for training unsupervised monolingual AWEs in Chapter 3. Details regarding the implementation of AWEs were covered by exploring different neural network models, as well as the preprocessing of audio and how the quality of AWEs is measured.

CHAPTER 3

CONTRASTIVE LEARNING FOR ACOUSTIC WORD EMBEDDINGS

In this chapter we first describe two state-of-the-art AWE models previously mentioned in Section 2.3, namely the CAE-RNN and SIAMESERNN. We then present a new model that has not been used for AWE modelling, which relies on contrastive learning. We call this model the CONTRASTIVERNN. A direct comparison is made between the two existing models and the new model by reproducing an experimental setup used in previous work. We evaluate each AWE model following both the unsupervised monolingual approach (Section 2.4.1) and the multilingual transfer approach (Section 2.4.2). In the first, an unsupervised AWE model is trained on unlabelled data in the target language. In the second, a supervised multilingual model is trained on labelled data from multiple well-resourced languages and then only applied to an unseen target language. Here we describe the training setup and discuss the results obtained for each of the three AWE models.

Parts of this chapter were presented at IEEE Spoken Language Technology (SLT):

C. Jacobs, Y. Matuskevych, and H. Kamper, “Acoustic word embeddings for zero-resource languages using self-supervised contrastive learning and multilingual adaptation,” in *Proceedings of SLT*, 2021.

3.1 RELATED WORK

Many unsupervised AWE modelling approaches have been proposed as an alternative to DTW (Section 2.2.2) for calculating the similarity between speech segments of different durations. Here we remind the reader of the state-of-the-art approaches identified in Section 2.3, all of which employ deep neural networks built on RNNs (Section 2.1.1).

An RNN-based AWE model sequentially processes a word segment $X = (\mathbf{x}_1, \mathbf{x}_2, \mathbf{x}_3, \dots, \mathbf{x}_T)$, where each \mathbf{x}_t is a frame-level acoustic feature, to produce a fixed-dimensional vector that represents segment X . The AE-RNN of Chung et al. [19] consists

3.1. RELATED WORK

of an encoder RNN and decoder RNN structure. A single recurrent layer processes a word segment X . The output of the recurrent layer (encoder RNN), after processing the final acoustic feature \mathbf{x}_T , is given as input to another single recurrent layer (decoder RNN) to reconstruct the original input sequence. After training, the final output of the recurrent layer is taken as the AWE of an input segment. In their experiments, the AE-RNN is trained on unknown isolated speech segments. The speech segments are obtained by segmenting unlabelled speech data in the zero-resource language using forced-alignments. Although the AE-RNN models itself is trained in an unsupervised fashion, using word boundaries for segmentation relies on transcriptions of the zero-resource language.

The SIAMESERNN of Settle and Livescu [32] explicitly optimise for the distance between embeddings instead of using a reconstruction loss as the AE-RNN. They show the SIAMESERNN is especially better at organising embeddings of words that were not seen during training compared to other methods they considered. Here they implement an encoder RNN with multiple recurrent layers, instead of a single recurrent layer. Additionally, they apply a set of fully connected layers to the output of the last recurrent layer, projecting the output to a lower-dimensional vector. The output of the RNN might be larger than the desired learned representation needed to learn intermediate information while processing the sequential input sequence. Therefore, the final layer of the encoder RNN might contain redundant information not contributing to the discriminative characteristics of the learned representation. By adding a set of fully connected layers, the representation can be transformed to be more discriminative. They train the SIAMESERNN using a weak form of supervision where word segments are grouped into unknown word types. Again, this approach is not completely unsupervised since word boundaries are used for segmentation and words are grouped using class labels.

The CAE-RNN model of Kamper [21] is an extension of the AE-RNN. In the CAE-RNN, unlike the AE-RNN, the target output is not identical to the input, but rather an instance of the same word type. They were the first to use this correspondence learning technique in an encoder-decoder structure operating on whole-word speech segments. Similar to Settle and Livescu [32], they extended the encoder RNN structure by adding a single fully connected feedforward layer to the output of the final output of the last recurrent layer. Most importantly, the CAE-RNN is trained in a complete zero-resource setting. In this setting, no word boundaries or class labels are assumed to be available. Isolated word-like pairs are obtained by applying a UTD system (Section 2.2.4) to the unlabelled speech data. The UTD system is in itself unsupervised making the whole process unsupervised. This training setting is the unsupervised monolingual setting as we described in Section 2.4.1. They showed the CAE-RNN outperforms the AE-RNN trained in the fully zero-resource setting.

Recently, Kamper et al. [23] presented the multilingual transfer training strategy as described in Section 2.4.2. Here they implement multilingual variants of the CAE-RNN

3.2. BASELINE ACOUSTIC WORD EMBEDDING MODELS

and SIAMESERNN. They show that applying a multilingual model to an unseen target language, outperforms an unsupervised monolingual model trained on unlabelled data from the target language.

3.2 BASELINE ACOUSTIC WORD EMBEDDING MODELS

In this section we describe the CAE-RNN and SIAMESERNN AWE models as implemented in Kamper et al. [23] which we reimplement. We compare these two models to the CONTRASTIVERNN, a new model we introduce in Section 3.3.

3.2.1 CAE-RNN

The CAE-RNN AWE model of Kamper [21], Kamper et al. [23] uses the CAE encoder-decoder network structure (Section 2.1.3) built on RNNs. The CAE-RNN is trained on pairs of speech segments (X, X') , with $X = \mathbf{x}_1, \dots, \mathbf{x}_T$ and $X' = \mathbf{x}_1, \dots, \mathbf{x}_{T'}$, containing different instances of the same word type, with each \mathbf{x}_t an acoustic feature vector. Given an input segment X , the encoder RNN produces a latent embedding \mathbf{z} from projecting the output of the final hidden layer to a lower-dimensional representation. The embedding \mathbf{z} is given as input to the decoder RNN to reconstruct an output sequence $Y = (\mathbf{y}_1, \mathbf{y}_2, \mathbf{y}_3, \dots, \mathbf{y}_{T'})$, where \mathbf{y}_t is the projection of the final hidden layer at time t conditioned on \mathbf{z} . The CAE-RNN model is trained using the reconstruction loss,

$$\begin{aligned} \ell(X, X') &= \|X' - f_{\Theta_{CAE}}(X)\|^2 \\ &= \|X' - Y\|^2, \end{aligned} \tag{3.1}$$

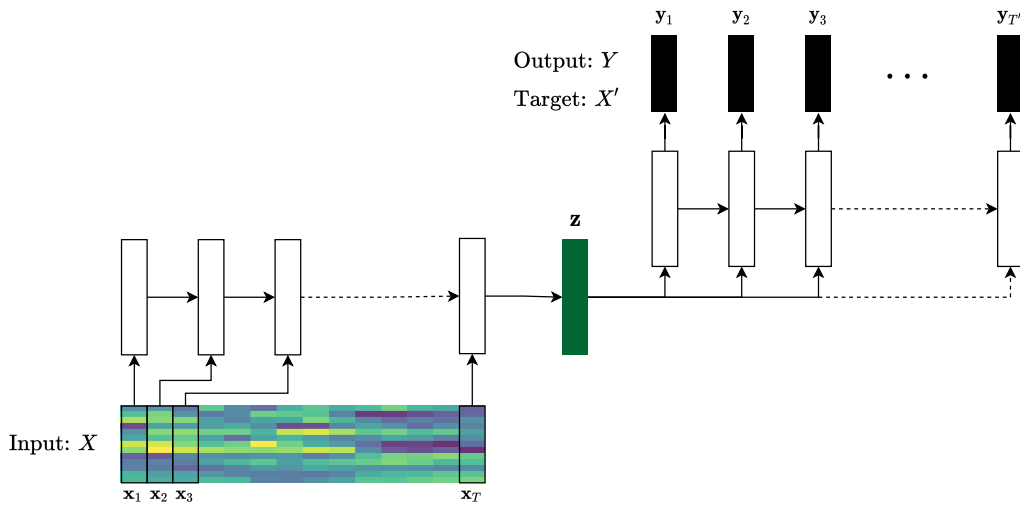


Figure 3.1: The CAE-RNN is trained to reconstruct an instance X' of the same word type as the input sequence X . T' and T are the lengths of X' and X , respectively.

3.2. BASELINE ACOUSTIC WORD EMBEDDING MODELS

with X' and Y the target and output, respectively. Figure 3.1 illustrates this model. After training, the projection \mathbf{z} is taken as the AWE of input segment X .

By forcing the model to reconstruct a word segment from a different instance of the same word type, the AWEs should be invariant to properties not common to two segments such as speaker, gender and channel, but capture the aspects that are, for instance, word type. As in [21], we first pretrain the CAE-RNN as an autoencoder using the loss in Equation 2.3 and then switch to the loss function for correspondence training.

3.2.2 SIAMESERNN

The SIAMESERNN of Kamper et al. [23], Settle and Livescu [32], uses the Siamese network configuration (Section 2.1.4) built on RNNs. The model consists of an encoder RNN that processes a word segment X to produce an embedding \mathbf{z} , similar to the encoder of the CAE-RNN. Instead of reconstructing a target output segment, the SIAMESERNN model explicitly optimises the relative distances between embeddings.

Given input sequences X_a , X_p , X_n , the model produces embeddings \mathbf{z}_a , \mathbf{z}_p , \mathbf{z}_n , as illustrated in Figure 3.2. Inputs X_a and X_p are from the same word type (subscripts indicate anchor and positive) and X_n is from a different word type (negative). For a single triplet of inputs, the model is trained using the triplet loss:

$$l_{\text{triplet}}(\mathbf{z}_a, \mathbf{z}_p, \mathbf{z}_n) = \max\{0, m + d(\mathbf{z}_a, \mathbf{z}_p) - d(\mathbf{z}_a, \mathbf{z}_n)\}, \quad (3.2)$$

with m a margin parameter and $d(\mathbf{u}, \mathbf{v}) = 1 - \mathbf{u}^T \mathbf{v} / \|\mathbf{u}\| \|\mathbf{v}\|$ denoting the cosine distance between two vectors \mathbf{u} and \mathbf{v} . This loss is at a minimum when all embedding pairs $(\mathbf{z}_a, \mathbf{z}_p)$ of the same type are more similar by a margin m than pairs $(\mathbf{z}_a, \mathbf{z}_n)$ of different types.

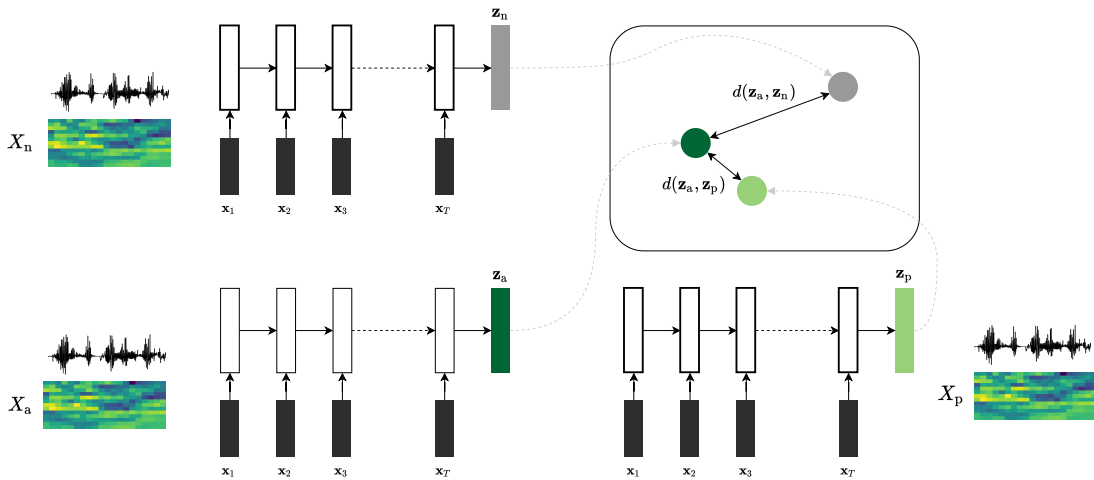


Figure 3.2: In the SIAMESERNN, three encoder RNNs use the same set of parameters to produce embeddings \mathbf{z}_a , \mathbf{z}_p , \mathbf{z}_n from input segments X_a , X_p , X_n . The model is trained to minimise the distance between the anchor and the positive item while maximising the distance between the anchor and negative item.

3.3. CONTRASTIVERNN

Generating all possible triplets would lead to slow convergence given that many of them would easily satisfy the triplet constraint in Equation 3.2. Therefore, we select triplets that violate the triplet constraint. The SIAMESERNN of Kamper et al. [23] uses the online semi-hard mining strategy of Schroff et al. [99] to sample negative examples. They sample hard negatives and positives from a single training batch. In a single training batch, they sample all anchor-positive pairs and select the hardest negative for each. In our implementation we use an online hard mining strategy [74]: for each item (anchor) in the batch we select the hardest positive and hardest negative example. For each item in the batch, the hardest positive is an instance of the same word type whose embedding is currently the furthest away out of all the other positive items in the batch. Conversely, the hardest negative is an item from a different word type whose embedding is the closest out of all the negative items in the batch.

3.3 CONTRASTIVERNN

This section introduces a new AWE model, namely the CONTRASTIVERNN. We will compare this model to two existing state-of-the-art AWE models, namely the CAE-RNN and SIAMESERNN. We first describe the concept of contrastive learning. We then describe how we implement contrastive learning in AWE modelling.

3.3.1 CONTRASTIVE LEARNING

In broad terms, contrastive learning is a technique to learn representations from data inputs such that representations from the same class are close in vector space and representations from different classes are separated. A contrastive loss function calculates the distance between two outputs from the same class and contrasts that with the distance between the output of one or more different classes. The triplet loss optimised by the SIAMESERNN is a form of contrastive learning, but in this thesis, we refer to contrastive loss where a positive example is compared to multiple negative examples.

Distance metric learning frameworks, such as the triplet loss, often suffer from slow convergence and poor local optima [44, 100]. The triplet loss only compares an input example to one negative example without considering examples from the remaining negative classes. Consequently, for a single training instance, an input example is only guaranteed to be pushed away from one negative example. Hence, samples are only separated from limited negative examples, still appearing close to several other classes in vector space. In practice, after enough iterations of randomly selecting negative examples, the triplet loss should be optimised for optimal class separability. However, even with curated negative example mining strategies, the triplet loss often struggles to organise the embedding space adequately.

3.3. CONTRASTIVERNN

To alleviate this problem, Sohn [44] proposed a *multi-class N-pair* loss that optimises to identify a positive example from $N - 1$ negative examples. This loss extends the triplet loss by incorporating multiple negative examples for each positive pair: an input example is being compared to a positive example and negative examples from multiple classes that it needs to discriminate from at the same time.

Sohn [44] formalised the N -pair loss as follows: given a $(N+1)$ -tuple training example $\{X_a, X_p, X_{n_1}, X_{n_2}, \dots, X_{n_{N-1}}\}$, with X_a the input example, X_p a positive example and $\{X_i\}_{i=1}^{N-1}$ negative examples, the $(N+1)$ -tuple loss is defined as follows:

$$l_{(N+1)\text{-tuple}}(\{X_a, X_p, X_{n_1}, X_{n_2}, \dots, X_{n_{N-1}}\}; f_\Theta) = \log \left(1 + \sum_{i=1}^{N-1} \exp(\mathbf{z}_a^T \mathbf{z}_{n_i} - \mathbf{z}_a^T \mathbf{z}_p) \right) \quad (3.3)$$

where f_Θ is an embedding function producing embeddings \mathbf{z} from segments X . Sohn showed that for the scenario when $N = 2$, the $(2+1)$ -tuple loss closely resembles the triplet loss in Equation 2.6:

$$l_{(2+1)\text{-tuple}}(\{X_a, X_p, X_n\}; f_\Theta) = \log(1 + \exp(\mathbf{z}_a^T \mathbf{z}_n - \mathbf{z}_a^T \mathbf{z}_p)) \quad (3.4)$$

where the embedding function f_Θ minimises the $l_{(2+1)\text{-tuple}}$ if and only if it minimises l_{triplet} . Rewriting the loss function in Equation 3.3 to

$$\log \left(1 + \sum_{i=1}^{N-1} \exp(\mathbf{z}_a^T \mathbf{z}_{n_i} - \mathbf{z}_a^T \mathbf{z}_p) \right) = -\log \frac{\exp(\mathbf{z}_a^T \mathbf{z}_p)}{\exp(\mathbf{z}_a^T \mathbf{z}_p) + \sum_{i=1}^{N-1} \exp(\mathbf{z}_a^T \mathbf{z}_{n_i})} \quad (3.5)$$

looks similar to the multi-class logistic loss where including more negative examples gives a better approximation. This shows the advantage of the $(N+1)$ -tuple loss over the triplet loss.

Previous work successfully implemented this multi-class negative mining loss for visual representation learning [43, 101]. To the best of our knowledge, this loss has not been implemented in whole-word speech representation learning in AWE modelling. In the following section we describe how we apply contrastive learning to AWE modelling.

3.3.2 CONTRASTIVE LEARNING IN ACOUSTIC WORD EMBEDDINGS

We specifically implement the contrastive loss of [43] using the same embedding function as the SIAMESERNN (Section 3.2.2); we call this model the CONTRASTIVERNN. Concretely, given inputs X_a and X_p and multiple negative examples X_{n_1}, \dots, X_{n_K} , the CONTRASTIVERNN produces embeddings $\mathbf{z}_a, \mathbf{z}_p, \mathbf{z}_{n_1}, \dots, \mathbf{z}_{n_K}$. Let $\text{sim}(\mathbf{u}, \mathbf{v}) = \mathbf{u}^\top \mathbf{v} / \|\mathbf{u}\| \|\mathbf{v}\|$ denote the cosine similarity between two vectors \mathbf{u} and \mathbf{v} . The loss given a positive pair (X_a, X_p) and the set of negative examples is then defined as [43]:

$$\ell(\mathbf{z}_a, \mathbf{z}_p, \mathbf{z}_{n_1}, \dots, \mathbf{z}_{n_K}) = -\log \frac{\exp\{\text{sim}(\mathbf{z}_a, \mathbf{z}_p)/\tau\}}{\sum_{j \in \{p, n_1, \dots, n_K\}} \exp\{\text{sim}(\mathbf{z}_a, \mathbf{z}_j)/\tau\}}, \quad (3.6)$$

3.4. EXPERIMENTAL SETUP




Figure 3.3: A visualisation of the difference in the optimisation of (a) the SIAMESERNN and (b) the CONTRASTIVERNN for a single positive pair $(\mathbf{z}_a, \mathbf{z}_p)$ in the embedding space.

where τ is a temperature parameter. The temperature parameter does not directly affect the accuracy but helps gradients to be propagated more easily.

The difference between this loss and the triplet loss used in the SIAMESERNN is illustrated in Figure 3.3. Ideally, we would like to include a negative example of each class for an input example. However, this is not computationally efficient and could lead to poor generalisation with class imbalances. We follow an offline batch construction process where we randomly choose N distinct positive pairs. Given a positive pair (X_a, X_p) , the remaining $2(N - 1)$ items are then treated as negative examples. The final loss is calculated as the sum of the loss over all N positive pairs within the batch.

The CONTRASTIVERNN can be trained in both the unsupervised monolingual (Section 2.4.1) and supervised multilingual setting (Section 2.4.2).

3.4 EXPERIMENTAL SETUP

We follow the same setup as Kamper et al. [23] to train AWE models in the monolingual unsupervised (Section 2.4.1) and supervised multilingual (Section 2.4.2) settings. In the first, an unsupervised AWE model is trained on unlabelled data in the target language. In the second, a supervised multilingual model is trained on labelled data from multiple well-resourced languages and then only applied to an unseen target language. Here we describe the training and evaluation data, and details regarding the configuration of the AWE models (CAE-RNN, SIAMESERNN, CONTRASTIVERNN).

3.4. EXPERIMENTAL SETUP

3.4.1 DATA

We perform experiments using the GlobalPhone corpus of read speech [102]. This corpus contains 20 languages covering a variety of speech peculiarities. For each language, 100 adult native speakers were recorded reading approximately 100 sentences selected from national newspaper articles available on the web. As in [23], we treat six languages as our target zero-resource languages: Spanish (ES), Hausa (HA), Croatian (HR), Swedish (SV), Turkish (TR) and Mandarin (ZH). Each language has on average 16 hours of training, 2 hours of development and 2 hours of test data. We apply the UTD system of [10] to the training set of each zero-resource language and use the discovered pairs to train unsupervised monolingual embedding models (Section 2.4.1). The UTD system discovers around 36k pairs for each language, where pair-wise matching precisions vary between 32% (SV) and 79% (ZH). Training conditions for the unsupervised monolingual CAE-RNN, SIAMESERNN and CONTRASTIVERNN models are determined by doing validation on the Spanish development data. The same hyperparameters are then used for the five remaining zero-resource languages.

For training supervised multilingual embedding models (Section 2.4.2), six other GlobalPhone languages are chosen as well-resourced languages: Czech, French, Polish, Portuguese, Russian and Thai. Each well-resourced language has on average 21 hours of labelled training data. We pool the data from all six well-resourced languages and train a multilingual CAE-RNN, a SIAMESERNN and a CONTRASTIVERNN. Instead of using the development data from one of the zero-resource languages, we use another well-resourced language, German, for validation of each model before applying it to the zero-resource languages. We only use 300k positive word pairs for each model, as further increasing the number of pairs did not give improvements on the German validation data.

All speech audio is parametrised as $D = 13$ dimensional static Mel-frequency cepstral coefficients (MFCCs).

3.4.2 ACOUSTIC WORD EMBEDDING MODELS

All our models use a similar architecture as implemented in [23]: encoders and decoders consist of three unidirectional RNNs with 400-dimensional hidden vectors, and all models use an embedding size of 130 dimensions. Models are optimised using Adam optimisation [103], with a learning rate of 0.001 for the CAE-RNN and CONTRASTIVERNN, and $5 \cdot 10^{-4}$ for the SIAMESERNN. The margin parameter m in Section 3.2.2 and temperature parameter τ in Section 3.2.1 are set to 0.25 and 0.1, respectively. We train the CAE-RNN and SIAMESERNN with a batch size of 300, and set the batch size for the CONTRASTIVERNN to 600.

Learning rates and hyperparameters are set by experimenting on Spanish development data for the unsupervised monolingual models, and German development data for the

supervised multilingual models.

3.5 EXPERIMENTS

We start in Section 3.5.1 by evaluating the different AWE models using the word discrimination task described in Section 2.5. Instead of only looking at word discrimination results, it is useful to also use other methods to try and better understand the organisation of AWE spaces [104], especially in light of recent results showing that AP has limitations as described in Section 2.5. We therefore also look at speaker classification performance in Section 3.5.2. Speaker classification is not necessarily a task in which we want our models to be good at. Since we are mostly interested in producing AWEs that are speaker invariant, a lower speaker classification is desirable. This probing experiment is useful to get insight on the AWEs produced by different AWE models. For example, two AWEs might show similar AP in the word discrimination task, while one may capture more speaker information.

3.5.1 WORD DISCRIMINATION

We first consider purely unsupervised monolingual models (Section 2.4.1). We are particularly interested in the performance of the CONTRASTIVERNN, which has not been considered in previous work. The top section in Table 3.1 shows the performance for the unsupervised monolingual AWE models applied to the test data from the six zero-resource languages.¹ As a baseline, we also give the results where DTW is used directly on the MFCCs to perform the word discrimination task. We see that the CONTRASTIVERNN consistently outperforms the CAE-RNN and SIAMESERNN approaches on all six zero-resource languages. The CONTRASTIVERNN is also the only model to perform better than DTW on all six zero-resource languages, which is noteworthy since DTW has access to the full sequences for discriminating between words.

Next, we consider the supervised multilingual models (Section 2.4.2). The bottom section of Table 3.1 shows the performance for the supervised multilingual models applied to the six zero-resource languages. By comparing these supervised multilingual models to the unsupervised monolingual models (top), we see that in almost all cases the multilingual models outperform the purely unsupervised monolingual models, as also in [22, 23]. However, on Mandarin (ZH), the unsupervised monolingual CONTRASTIVERNN model outperforms all three multilingual models. Comparing the three multilingual models, we do not see a consistent winner between the CONTRASTIVERNN and CAE-RNN, with one performing better on some languages while the other performs better on others.

¹We note that the results for the CAE-RNN and SIAMESERNN here are slightly different to that of [22, 23], despite using the same test and training setup. We believe this is due to the different negative sampling scheme for the SIAMESERNN and other small differences in our implementation.

3.5. EXPERIMENTS

Table 3.1: AP (%) on test data for the six zero-resource languages. The purely unsupervised monolingual models are trained on discovered word segments (top). The multilingual models are trained on true words by pooling labelled training data from six well-resourced languages (bottom). The best approach in each subsection is shown in bold.

Model	ES	HA	HR	SV	TR	ZH
<i>Unsupervised models:</i>						
DTW	36.2	23.8	17.0	27.8	16.2	35.9
CAE-RNN	52.7	18.6	24.5	28.0	14.2	33.7
SIAMESERNN	56.6	16.8	21.1	31.8	22.8	52.0
CONTRASTIVERNN	70.6	36.4	27.8	37.9	31.3	57.1
<i>Multilingual models:</i>						
CAE-RNN	72.4	49.3	44.5	52.7	34.4	53.9
SIAMESERNN	70.3	45.3	40.6	47.5	27.7	49.9
CONTRASTIVERNN	73.3	50.6	45.1	46.4	34.6	53.2

The multilingual SIAMESERNN generally performs worst, although it outperforms the CONTRASTIVERNN on Swedish (SV).

3.5.2 SPEAKER CLASSIFICATION

We now perform a probing experiment by considering the extent to which different AWE models capture speaker information. Performing well in a speaker classification task is not always desirable. A lower speaker classification score indicates that the representations capture little speaker information, which is what we want.

To measure speaker invariance, we use a linear classifier to predict a word’s speaker identity from its AWE. Specifically, we train a multi-class logistic regression model on 80% of the development data and test it on the remaining 20%.

The top section of Table 3.2 shows speaker classification results on development data for the three types of monolingual unsupervised models (Section 2.4.1). Since we are interested in how well models abstract away from speaker information, we consider lower accuracy as better (shown in bold). The CONTRASTIVERNN achieves the lowest speaker classification performance across all languages, except on Croatian where it performs very similarly to the SIAMESERNN. This suggests that among the unsupervised monolingual models, the CONTRASTIVERNN is the best at abstracting away from speaker identity (at the surface level captured by a linear classifier).

Next, we consider speaker classification performance for the multilingual models (Section 2.4.2). Comparing the middle and top sections of Table 3.2, we see that for each multilingual model (bottom) the speaker classification performance drops from its

3.6. CHAPTER SUMMARY

Table 3.2: Speaker classification accuracy (%) on development data for the six zero-resource languages using purely unsupervised monolingual models (top) and multilingual models (bottom).

Model	ES	HA	HR	SV	TR	ZH
<i>Unsupervised models:</i>						
CAE-RNN	52.8	49.5	53.8	47.8	48.9	62.0
SIAMESE-RNN	38.9	38.4	37.2	36.4	38.6	44.3
CONTRASTIVERNN	34.6	33.1	37.8	33.6	35.5	39.9
<i>Multilingual models:</i>						
CAE-RNN	40.6	44.7	45.0	42.3	42.0	49.4
SIAMESE-RNN	32.6	30.3	31.9	28.6	30.3	32.5
CONTRASTIVERNN	27.6	28.7	26.9	29.2	32.1	31.9

corresponding unsupervised monolingual version (top) across all six languages, again indicating an improvement in speaker invariance. Comparing the three multilingual models to each other (bottom), the CONTRASTIVERNN has the lowest speaker classification performance on four out of the six evaluation languages.

We conclude that the multilingual models are better at producing speaker-invariant AWEs compared to their monolingual variants. Also, the CONTRASTIVERNN is better at abstracting away from speaker information in both the unsupervised monolingual and multilingual transfer setting compared to the CAE-RNN and SIAMESE-RNN.

3.6 CHAPTER SUMMARY

In this chapter we reimplemented two existing AWE models (CAE-RNN and SIAMESE-RNN). We reproduced an experimental setup to compare these two models to a new model we presented (CONTRASTIVERNN), in both the unsupervised monolingual and supervised multilingual settings. We show that the CONTRASTIVERNN outperforms the CAE-RNN and SIAMESE-RNN in the unsupervised monolingual setting and performs on par with the CAE-RNN in the supervised multilingual setting. In the following chapter, we determine whether some available data (unlabelled) in the zero-resource setting can be used to adapt a multilingual model to cater for a specific zero-resource language.

CHAPTER 4

MULTILINGUAL ADAPTATION

In the previous chapter, we concluded that applying a supervised multilingual AWE model, trained on labelled data from multiple languages, to an unseen target zero-resource language outperforms an unsupervised monolingual AWE model trained on unlabelled data from the target language. In this chapter, we explore whether the performance of a multilingual AWE model for a specific zero-resource language can be improved by adapting the model using unlabelled data available in that particular language.

Parts of this chapter were presented at IEEE Spoken Language Technology (SLT): C. Jacobs, Y. Matusevych, and H. Kamper, “Acoustic word embeddings for zero-resource languages using self-supervised contrastive learning and multilingual adaptation,” in Proceedings of SLT, 2021.

4.1 INTUITION AND RELATED WORK

The supervised multilingual transfer approach of Kamper et al. [23] as described in Section 2.4.2 and implemented in Chapter 3, is a form of transfer learning categorised as *transductive transfer learning*.

Transductive transfer learning [95]: Given a source domain \mathcal{D}_S and a corresponding learning task \mathcal{T}_S , a target domain \mathcal{D}_T and a corresponding learning task \mathcal{T}_T , transductive transfer learning aims to improve the learning of the target predictive function $f_{T(\cdot)}$ in \mathcal{D}_T using the knowledge in \mathcal{D}_S and \mathcal{T}_S , where $\mathcal{D}_S \neq \mathcal{D}_T$ and $\mathcal{T}_S = \mathcal{T}_T$. In addition, some unlabelled target domain data must be available at training time.

The scenario of training a model in one domain and then transferring it to a domain of interest, is especially useful in the unsupervised setting where a sufficient amount of labelled data is available in the source domain and only unlabelled data in the target domain are assumed to be available. In the supervised multilingual approach, we identify

the source domain \mathcal{D}_S as the setting in which labelled training data from multiple well-resourced languages are available and the target domain \mathcal{D}_T as the zero-resource setting where only unlabelled speech data from an unseen target language is available. We identify the target predictive function as an embedding function $f_{\theta}(\mathbf{x})$ that produces a fixed-dimensional embedding given an arbitrary-length speech segment \mathbf{x} . The task (producing fixed-dimensional embeddings from variable-length speech segments) across both domains \mathcal{T}_S and \mathcal{T}_T are the same, irrespective of the language. We confirm the multilingual transfer approach fits the definition of transductive transfer learning as given in the definition. Note, the transfer learning method implemented by Kamper et al. [23] (Section 2.4.2) only partially performs transductive transfer learning by not considering any data from the target domain. An AWE model was simply trained on labelled data from multiple languages (source domain) and then directly applied to an unseen zero-resource language (target domain).

In this chapter, we extend this transfer learning technique by using available unlabelled data from a zero-resource language to modify the embedding function. This transfer learning technique can simply be treated as a regular semi-supervised problem, ignoring the domain difference by considering the source and target data as labelled and unlabelled, respectively.

Recently, Hu et al. [25] presented adaptation of acoustically grounded word embeddings (AWEs are trained jointly on speech and text) by using labelled data (limited) in the target domain, showing improved results on the target language. However, to the best of our knowledge, we are the first to perform adaptation of AWEs with unlabelled data in the target domain.

4.2 UNSUPERVISED ADAPTATION OF MULTILINGUAL MODELS

We aim to improve the performance of a multilingual AWE model on a specific zero-resource language by manipulating the learned parameters of a multilingual model using unlabelled data from that language. In other words, we adapt a multilingual model by fine-tuning the pre-trained multilingual model directly on the available unlabelled data in the target domain. The idea is that adapting the multilingual AWE model to the target language would allow the model to learn aspects unique to the language.

We consider the unsupervised adaptation of all three multilingual AWE models (CAE-RNN, SIAMESERNN, CONTRASTIVERNN) in Chapter 3. We again use the UTD system of [10] to obtain discovered word pairs from the target language. We then use these discovered pairs to adapt the multilingual models in the same way we trained the unsupervised monolingual models in Chapter 3. We need to update a multilingual model's parameters

4.3. EXPERIMENTAL SETUP

in such a way that it doesn't unlearn its current "knowledge". Using a learning rate lower than the learning rate used for training the source multilingual model is one way to prevent parameters from deviating too much. When fine-tuning a model, instead of allowing all model parameters to update, freezing some parameters is another way to preserve some of the learned information.

4.3 EXPERIMENTAL SETUP

The experimental setup is the same as described in Section 3.4.1, except for the differences we describe here.

We adapt each multilingual model to every zero-resource language using the same discovered word pairs as for the unsupervised monolingual models in Chapter 3. We performed validation experiments to determine which parameters to update and which to keep fixed from the source multilingual models using Spanish development data. The same hyperparameters are used to adapt a multilingual model to each of the five remaining zero-resource languages. For the CAE-RNN (described in Section 3.2.1), we found that it is best to freeze the multilingual encoder RNN weights and only update the weights between the final encoder RNN hidden state and the embedding \mathbf{z} , as illustrated in Figure 4.1. We also found that it is best to re-initialise the decoder RNN weights randomly before




Figure 4.1: The same CAE-RNN network architecture as displayed in Figure 2.2. When adapting the CAE-RNN we freeze the encoder RNN weights (highlighted in grey) only allowing the projection weights \mathbf{W}_{proj} between the final hidden layer of the encoder RNN and latent embedding \mathbf{z} to update. The decoder RNN is re-initialised with random weights before training on the target data (highlighted in white).

4.4. EXPERIMENTS

training on the target data.

For the SIAMESERNN (Section 3.2.2) and CONTRASTIVERNN (Section 3.3) we update all weights during adaptation. The learning rate for the CAE-RNN, SIAMESERNN and CONTRASTIVERNN were adjusted to $1 \cdot 10^{-4}$, $1 \cdot 10^{-5}$ and $1 \cdot 10^{-4}$, respectively.

4.4 EXPERIMENTS

We repeat the experiments in Chapter 3, this time evaluating the multilingual adapted models on the word discrimination and speaker classification evaluation tasks. We specifically, compare the results of the adapted multilingual models to the original multilingual models in Table 3.1. In addition to the quantitative evaluations, we also provide some qualitative analysis to illustrate the effect of adaptation in vector space.

4.4.1 WORD DISCRIMINATION

The results after adapting each multilingual model (CAE-RNN, SIAMESERNN, CONTRASTIVERNN) to each of the zero-resource languages are shown in the bottom section of Table 4.1. In the top section of Table 4.1, we again show the performance of the source multilingual models on each language (bottom section of Table 3.1). Comparing the adapted models (bottom) to the original multilingual models (top) in Table 4.1, we see that most of the adapted models outperform their corresponding source multilingual models, with the CONTRASTIVERNN and SIAMESERNN improving substantially after adaptation on some of the languages. The adapted CONTRASTIVERNN models outperform

Table 4.1: The AP (%) on test data for all six zero-resource languages. The top section is a duplicate of the multilingual models in Table 3.1. The bottom section is the results achieved by adapting each multilingual model (CAE-RNN, SIAMESERNN, CONTRASTIVERNN) to every zero-resource language using discovered words from the target zero-resource language. The model delivering the highest score per language is highlighted in bold.

Model	ES	HA	HR	SV	TR	ZH
<i>Multilingual models:</i>						
CAE-RNN	72.4	49.3	44.5	52.7	34.4	53.9
SIAMESERNN	70.3	45.3	40.6	47.5	27.7	49.9
CONTRASTIVERNN	73.3	50.6	45.1	46.4	34.6	53.2
<i>Multilingual adapted:</i>						
CAE-RNN	74.2	49.4	45.9	53.4	34.5	53.9
SIAMESERNN	74.5	44.7	37.6	50.3	30.3	57.3
CONTRASTIVERNN	76.6	56.7	44.4	54.4	40.5	60.4

4.4. EXPERIMENTS

the adapted CAE-RNN and SIAMESE-RNN models on five out of the six zero-resource languages. The results in Table 4.1 are some of the best-reported results on these data from the same-different AP. From the results in Table 4.1 we conclude that unsupervised adaptation of multilingual models to a target zero-resource language is an effective AWE approach, especially when coupled with the self-supervised contrastive loss.

One question is whether adapted models close the gap between the zero-resource setting and the best-case scenario where we have labelled data available in a target language. To answer this, Table 4.2 compares adapted multilingual models (bottom) to “oracle” supervised monolingual models trained on labelled data from the six evaluation languages (top) on development data. Although Table 4.1 shows that adaptation greatly improves performance in the zero-resource setting, Table 4.2 shows that multilingual adaptation still does not reach the performance of supervised monolingual models.

4.4.2 SPEAKER CLASSIFICATION

As in Section 3.5.2, we again consider speaker classification probing experiments to see to what extent the resulting features normalise out speaker characteristics after adaptation. As a reminder, the resulting AWEs should avoid capturing speaker information, in other words, a lower speaker classification score is desired.

The impact of unsupervised adaptation (Section 2.4.3) on speaker invariance, is shown at the bottom of Table 4.3. We again show the original multilingual speaker classification results (bottom section of Table 3.2) at the top of Table 4.3. After adaptation (bottom) we see that speaker classification results improve consistently compared to their corresponding source multilingual model (top). Although this seems to indicate that the adapted AWEs capture more speaker information, these embeddings still lead to better word discrimination

Table 4.2: AP (%) on development data. The supervised monolingual models are trained on ground truth words from the target languages to determine an upper bound on performance. The adapted models are the same as those in Table 4.1 but applied to development data here for the purpose of analysis.

Model	ES	HA	HR	SV	TR	ZH
<i>Supervised monolingual:</i>						
CAE-RNN	70.2	79.7	63.0	55.8	65.6	84.2
SIAMESE-RNN	78.6	85.2	79.3	68.6	77.6	93.1
CONTRASTIVERNN	81.8	82.4	80.3	70.9	80.5	92.1
<i>Multilingual adapted:</i>						
CAE-RNN	51.7	59.5	44.3	38.6	40.9	52.2
SIAMESE-RNN	51.5	52.7	38.4	40.3	33.0	56.5
CONTRASTIVERNN	58.0	60.5	40.5	43.8	46.9	60.2

4.4. EXPERIMENTS

Table 4.3: Speaker classification accuracy (%) on development data for the zero-resource languages using multilingual models after adaptation.

Model	ES	HA	HR	SV	TR	ZH
<i>Multilingual models:</i>						
CAE-RNN	40.6	44.7	45.0	42.3	42.0	49.4
SIAMESE-RNN	32.6	30.3	31.9	28.6	30.3	32.5
CONTRASTIVERNN	27.6	28.7	26.9	29.2	32.1	31.9
<i>Multilingual adapted:</i>						
CAE-RNN	40.3	45.3	45.0	43.1	41.0	49.7
SIAMESE-RNN	33.6	33.9	30.6	29.4	33.7	37.9
CONTRASTIVERNN	30.3	32.1	26.9	29.5	36.5	35.1

performance (Table 4.1). A similar trend was observed in [23]: a model leading to better (linear) speaker classification performance does not necessarily give worse AP. Note, in Chapter 3 we emphasised the desire for producing AWEs that capture less speaker information. Although this is the case for many tasks, capturing speaker information may be favourable for some downstream tasks.

We conclude that that unsupervised adaptation using unlabelled data in a target zero-resource language leads to representations allowing better word discrimination while retaining more speaker information. Capturing speaker information might be beneficial for some downstream tasks. For practical implementation, the choice of applying adaptation will depend on the downstream application at hand.

4.4.3 QUALITATIVE ANALYSIS

Given the high-dimensional embedding space it is not easy to gain an intuitive understanding of how these embedding spaces are organised. We can however apply dimensionality reduction techniques to visualise the high-dimensional data in a lower-dimensional space. t-SNE [105] is a non-linear dimensionality reduction technique for the purpose of this task, specifically designed for data exploration and visualisation of high-dimensional data. The goal is to find a lower-dimensional mapping of higher-dimensional data while preserving relative distances between data points. Figure 4.2 shows a t-SNE visualisation of the AWEs produced by the CONTRASTIVERNN on Hausa data before and after adaptation. In this curated example, we see how some of the words that are clustered together by the multilingual model (e.g. “amfani” and “hankali”) are separated after adaptation. This observation supports the results in Table 4.1 that shows, exposing the multilingual model to the target zero-resource language benefits the model by learning some peculiarities from the specific language to improve the organisation of embeddings in vector space.

4.5. CHAPTER SUMMARY

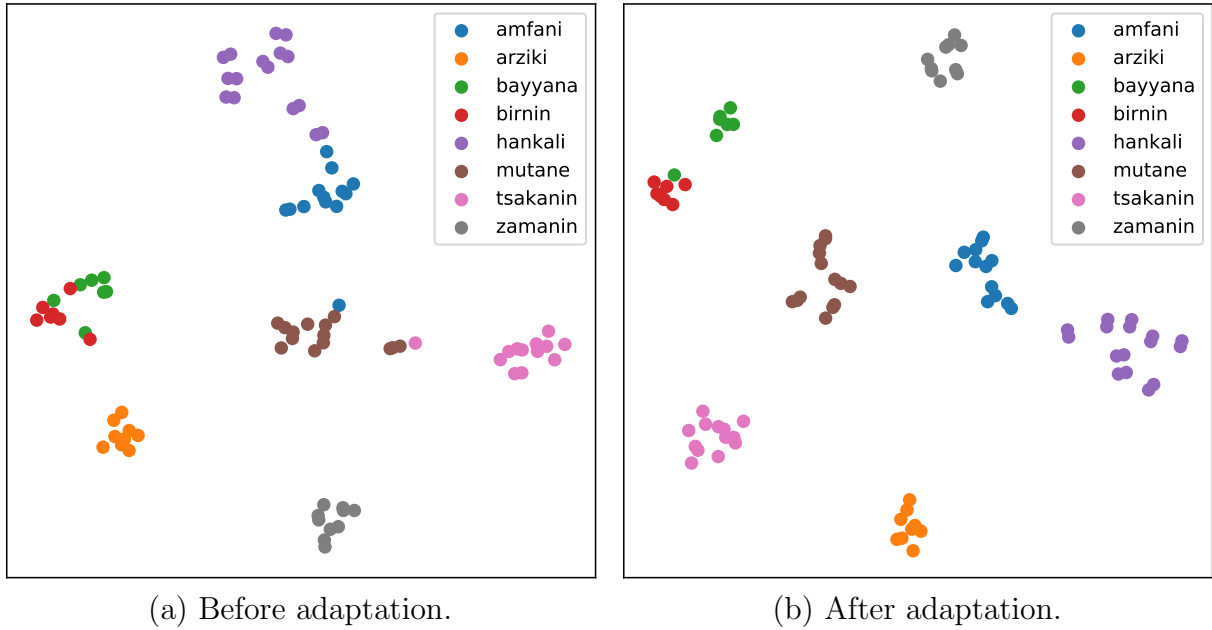


Figure 4.2: t-SNE visualisations of acoustic embeddings for the most frequent words in the Hausa test data, produced by (a) the multilingual CONTRASTIVERNN model (Table 3.1) and (b) the multilingual CONTRASTIVERNN model adapted to Hausa (Table 4.1).

4.5 CHAPTER SUMMARY

In this chapter we explored the benefit of using some unlabelled speech data available from a zero-resource language to cater a pretrained multilingual AWE model to that specific language. Our results indicate that the multilingual transfer approach (Section 2.4.2) along with unsupervised adaptation (proposed in this chapter) is beneficial, especially when implemented with the CONTRASTIVERNN AWE model which showed the highest performance increase after adaptation compared to the CAE-RNN and SIAMESERNN. In the following chapter, we further investigate multilingual models, where we investigate the choice of multilingual training language for specific zero-resource languages.

CHAPTER 5

IMPACT OF LANGUAGE CHOICE IN A MULTILINGUAL TRANSFER SETTING

In the previous chapter, we showed how the performance of an AWE model can improve on a zero-resource language by adapting the model to that particular language using unlabelled data from that language. In this chapter, we again try to improve the performance of a multilingual model on a zero-resource language, this time considering the choice of training languages in the multilingual transfer setting. We perform our experiments on a speech corpus containing under-resourced languages spoken in South Africa. We hypothesise that applying a multilingual model trained on one set of languages to a target language would achieve different results from a model trained on another set of languages. Specifically, we investigate the difference in performance when training on languages that are more related to the target than others, and the amount of training data required to achieve competitive results. With such knowledge, practitioners could save time and money on collecting resources when implementing a multilingual AWE model for a specific target zero-resource language.

The work in this chapter was presented at Interspeech:

C. Jacobs and H. Kamper, “Multilingual transfer of acoustic word embeddings improves when training on languages related to the target zero-resource language,” in *Proceedings of Interspeech*, 2021.

5.1 RELATED WORK

The impact training languages have on a particular target language, in a multilingual transfer setting, has been observed by several studies in general speech processing (i.e. unrelated to AWEs). Recent work in van der Westhuizen et al. [47] considers the choice of training languages in a multilingual ASR model. They perform their experiments in a similar setting to the multilingual transfer setting we explored in the previous chapter.

5.1. RELATED WORK

However, in contrast to having no available labelled data from the target zero-resource languages, they have access to a limited amount of transcribed speech data for each of the target languages. With a total of six training languages, they train various combinations of multilingual models for each target language, always including the training data from the target language. They show that any combination of two or more languages (where one is the target language) outperform a monolingual model trained on the target language alone. This is similar to our findings in Table 3.1, showing a model trained on labelled data from multiple languages outperforms a model trained on unlabelled data (instead of limited labelled data) from the target language itself. Importantly, although adding all languages improves over a monolingual model, they show a significant improved performance can be achieved by a more careful selection of training languages: best results are achieved for both target languages when only training a multilingual model with one additional language to the target language. While some of the languages are phonetically more similar to the target languages than others, the impact thereof was not observed. In broad terms, languages are said to be phonetically similar when they share similar contrastive sound units (phonemes), which are the individual sounds that make up all words in a language. Although the benefit of using related languages were not observed by van der Westhuizen et al. [47], we identify a more careful selection of training languages in a multilingual transfer setting might be beneficial.

In contrast to the results of van der Westhuizen et al. [47], not observing the benefit of related language in a multilingual ASR model, Yi et al. [46] however shows performance increase when adding languages related to the target in a multilingual ASR model. They also show the benefit of adding data from more languages, where a multilingual ASR model trained with up to twelve languages outperform all models trained on fewer languages on all the target languages. Importantly, different to van der Westhuizen et al. [47], they do not include any data from the target languages when training the multilingual models.

In another task, spoken language identification (LID)—identifying the language spoken in a speech recording—van der Merwe [45] and Bedyakin and Mikhaylovskiy [106] show that more confusion occurs between the identity of languages that belong to the same language family. Interestingly to us, the experiments in van der Merwe [45] are performed on the same dataset we use to perform the experiments in this chapter to train multilingual AWE models (see Section 5.4.1). Van der Merwe [45] shows that most confusion occurs between languages from the same family, for example, Zulu and Xhosa, belonging to the Nguni language family, and Afrikaans and English, belonging to the West-Germanic family (see Section 5.2), with the least confusion observed between language from far related languages, for example, Zulu and Afrikaans. This might be related to words shared among the vocabularies of two closely related languages, but also the linguistic characteristics of related languages since for example, there doesn't exist much overlap in words between Afrikaans and English. We further investigate this matter in Section 5.4.1.

5.2. LANGUAGES OF SOUTH AFRICA

Although the impact of the family to which a language belong have been observed in other speech tasks, it has not been investigated extensively for AWEs. Preliminary experiments in Kamper et al. [23] showed improved scores when training a monolingual AWE model on one language and applying it to another from the same language family. But this has not been investigated thoroughly or with multilingual models, leaving several unanswered questions.

In this chapter, we answer some of the following questions: Does the benefit of training on related languages diminish as we train on more languages (which might or might not come from the same family as the target zero-resource language)? When training exclusively on related languages, does performance suffer when adding an unrelated language? Should we prioritise data set size or language diversity when collecting data for multilingual AWE transfer? To answer these questions we use a corpus of under-resourced languages spoken in South Africa. In the following section, we describe these languages and how they can be used to benefit the performance of a specific language.

5.2 LANGUAGES OF SOUTH AFRICA

The majority of languages in Africa are considered under-resourced [107]. This includes the eleven official languages of South Africa, ten of which are indigenous to South Africa: isiZulu, isiXhosa, isiNdebele, siSwati, Sepedi, Setswana, Sesotho, Xitsonga, Tshivenda and Afrikaans. The eleventh language, English, is in contrast, the most well-resourced language in the world. It is the country's lingua franca and the primary language of government, business, and commerce, bridging the language barrier between native languages among different groups in South Africa. Figure 5.1 shows the geographical regions where each language is predominantly spoken as first language.

These languages can be grouped into different language families based on the linguistic links they share. A language family is a group of different languages related through descent from a common ancestral language or parental language, called the protolanguage [108]. Although some languages do not have a protolanguage, most languages spoken throughout the world belong to a language family. The eleven official languages of South Africa are grouped into two primary groups, Southern Bantu and Indo-European, from where the languages are further grouped into sub-families as illustrated in Figure 5.2. From the Southern Bantu family there exist two principal branches, Nguni-Tsonga and Sotho-Makua-Venda, to which nine of the languages belong.

The Nguni-Tsonga group further branches into two groups, Nguni and Tswa-Ronga where isiZulu, isiXhosa, siSwati, and isiNdebele belong to the former and Xitsonga to the latter. The Nguni languages use a five-vowel system and are characterised by click

¹Images show statistical data from the 2011 South Africa's Census retrieved and adapted from Wikipedia, distributed under the CC-BY 2.0 license.

5.2. LANGUAGES OF SOUTH AFRICA

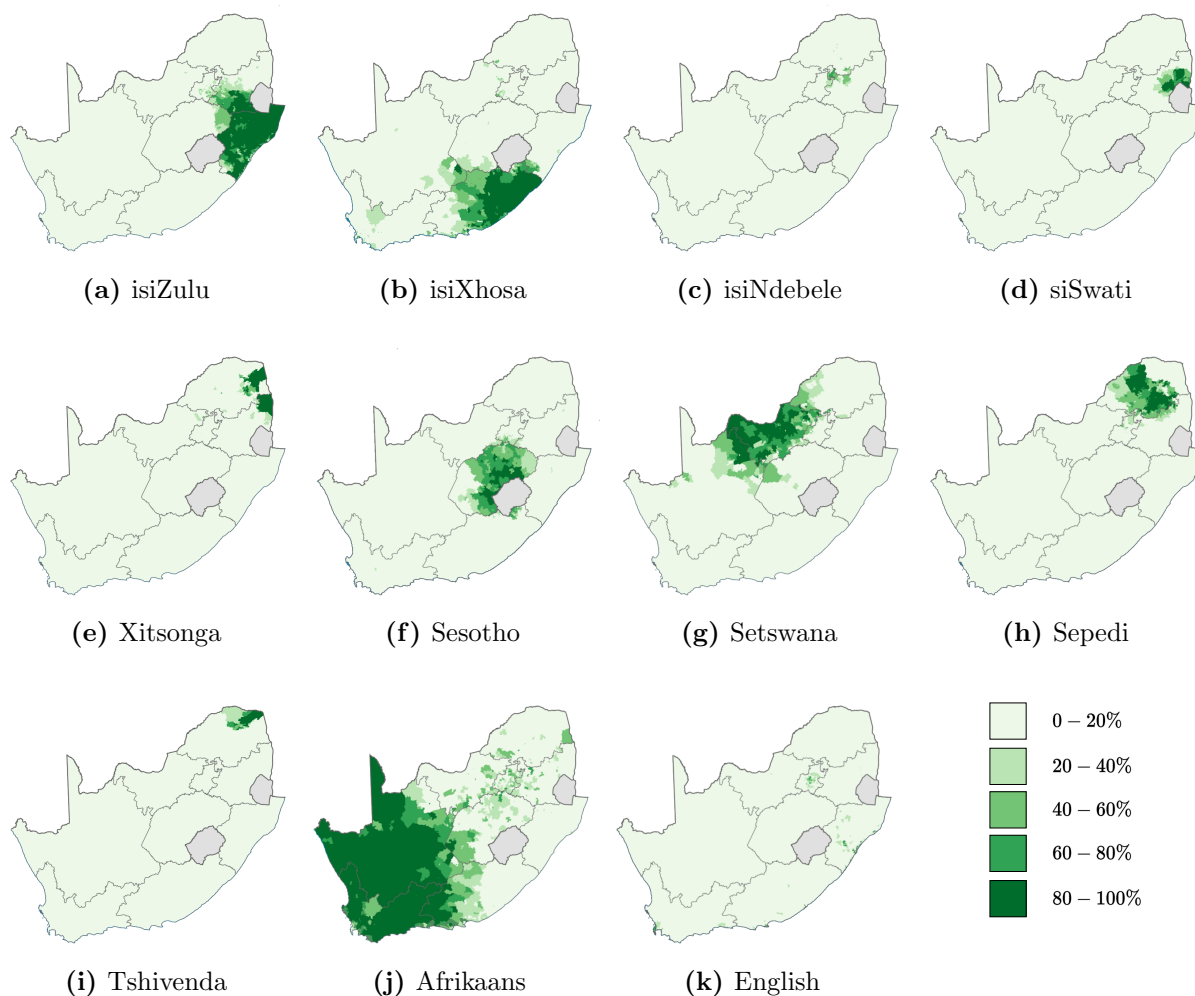


Figure 5.1: Geographical distributions showing the preferred language of communication in households in South Africa.¹

consonants. Xitsonga shares similar properties to the Nguni languages but has very few words with click consonants.

From the Sotho-Makua-Venda family, three languages belong to the Sotho-Tswana group: Sepedi, Setswana and Sesotho. In the Sotho-Makua-Venda subfamily, Venda is somewhat of a standalone. They use a 9-vowel system and unlike the Nguni languages, Sepedi and Setswana do not have clicks but Sesotho utilises click consonants in some words.

All the languages from the Southern Bantu group are known as tonal languages, where a word said with different tonal inflexions will convey different meanings. Moreover, for the Southern Bantu branch, languages grouped into families at the lowest level would (to an extent) be intelligible to a native speaker of another language in the same group.

The other two languages, Afrikaans and English, belong to the Indo-European family and are grouped at a higher level to belong to the set of West-Germanic languages. Although they are not as similar as some of the Southern Bantu languages, they are far less similar to each other than any of the languages in the Southern Bantu family.

5.3. QUERY-BY-EXAMPLE SPEECH SEARCH

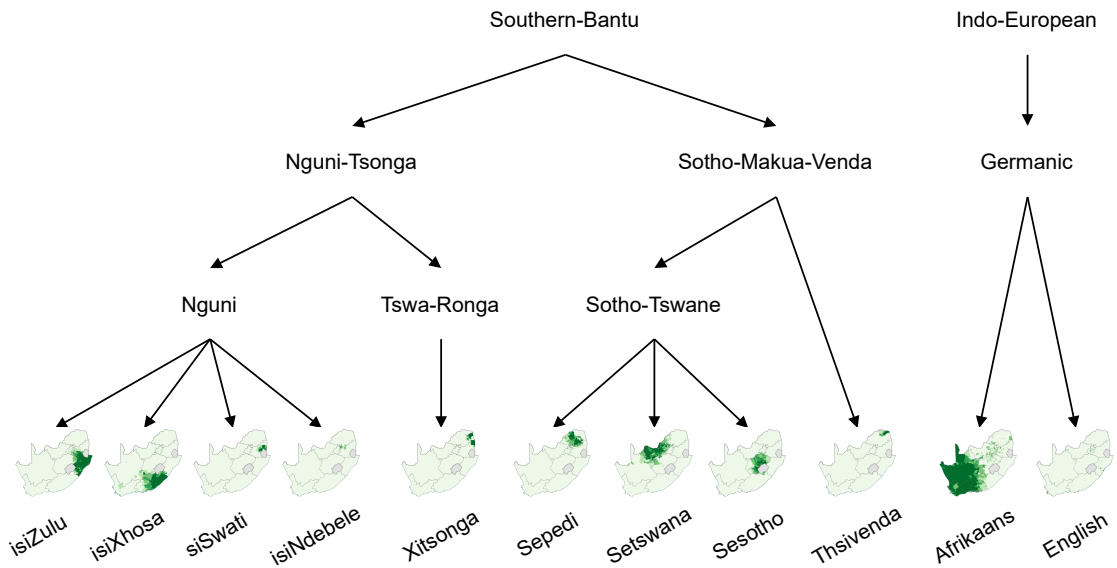


Figure 5.2: A family tree for the official South African languages.

5.3 QUERY-BY-EXAMPLE SPEECH SEARCH

In addition to the word discrimination task measuring the intrinsic quality of the AWEs (Section 2.5), we apply AWE models to a speech task, Query-by-Example (QbE) (Section 2.2.3), to measure an AWE model’s performance in an actual zero-resource downstream application. In contrast to the word discrimination task, QbE speech search does not assume a test set of isolated words but instead operates on full unsegmented utterances. Our goal is not to deliver competitive QbE results compared to existing work. Instead, we use this task as a supplementary evaluation task to support the word discrimination results as discussed in Section 2.5. We hope that the QbE results would correlate with the AP performance of the same-different word discrimination task (Section 2.5).

We implement a simplified QbE system used in Kamper et al. [91]. This system aims to retrieve utterances containing a given spoken query word from a speech collection containing multiple speech utterances. If we knew the word boundaries in the search collection, we could embed each of the words in an utterance and simply look up the closest embeddings to the query. Instead, because we do not have word boundaries in the zero-resource setting, each utterance is split into overlapping segments from some minimum to some maximum duration as illustrated in Figure 5.3. Each segment from each utterance is then embedded separately using the same AWE model. The spoken query word is also embedded using the same AWE model. To do the QbE task, the query embedding is compared to each of the utterance sub-segment embeddings (using cosine distance), and the minimum distance over the utterance is then taken as the score for whether the utterance contains the given query.

5.4. EXPERIMENTAL SETUP




Figure 5.3: A search utterance is split into overlapping segments which are embedded individually using the same AWE model.

5.4 EXPERIMENTAL SETUP

Here we describe the setup for training multilingual AWE models for the purpose of investigating the impact of related languages. We also give the implementation details for the QbE system (described above in Section 5.3) and how we evaluate the system retrieval accuracy.

5.4.1 DATA

We perform all our experiments on the NCHLT speech corpus [109], which provides wide-band speech from each of the eleven official South African languages (Section 5.2). We use a version of the corpus where all repeated utterances are removed, leaving roughly 56 hours of speech from around 200 speakers in each language with an equal split between males and females. We use the default training, validation and test sets from which word segments are extracted using forced alignments. After preliminary experiments and inspecting the extracted word segments, we added some constraints on the duration of word segments to be considered for this experiment. Nguni languages follow a conjunctive writing system and the Sotho languages follow a disconjunctive writing system such that the average word-length of Bantu languages are much longer in duration than Sotho languages [110]. We modified all the sets to only contain words with a minimum number of characters and frames for each language as shown in Table 5.1. We observe that, although the number of training items for each language is somewhat equal, the number of unique words in each

5.4. EXPERIMENTAL SETUP

Table 5.1: The number of word segments for each language in every set (train, development, test) after applying word constraints.

Language	Constraints		No. items			No. types		
	Min. chars	Min. frames	Train	Val	Test	Train	Val	Test
Xhosa	5	50	95k	7k	6.4k	26.4k	4.9k	4.2k
isiZulu	5	50	103k	8k	7.6k	22.9k	6.3k	4.9k
siSwati	5	50	97k	8k	7.2k	11.3k	4.1k	3.8k
isiNdebele	5	50	110k	9k	9.3k	13.8k	5k	4.8k
Sepedi	3	30	86k	5.2k	4.9k	10.9k	1.7k	2.3k
Setswana	4	40	110k	6.4k	5.7k	5.2k	2.2k	1.1k
Sotho	3	30	100k	5.9k	5.4k	10.2k	2.4k	2.5k
Afrikaans	5	50	98k	6.7k	6.2k	7.5k	3.1k	2.8k
English	5	50	116k	5k	5k	6.7k	2.2k	2.5k
Xitsonga	3	30	116k	7.6k	8.4k	5.8k	2.1k	2.1k
Tshivenda	3	30	117k	8k	7.1k	7.4k	2.2k	3.3k

language is quite different. For example, Xhosa and Afrikaans have similar number of training items, 95k and 98k, but the number of unique word types within the set is 26.4k compared to 7.5k. We will observe this effect in our experiments in the following section.

We recognise that languages from the same language family, especially the ones related at the lowest level (Figure 5.2), might share words among their vocabulary. As an initial investigation of language similarity, we therefore examine the amount of overlap between the training sets of all languages. Figure 5.4 shows the percentage of unique words for each language that also appear in the training set of every other language. As expected, we observe the most overlap between languages that belong to the same language family, illustrated by the darker shade blocks around the diagonal. For most of the language, with the exception of English and Afrikaans, limited text resources are available on the internet. Therefore, a large part of the text was obtained from South Africa’s government website, containing content in all the official languages. Given the nature of the prompts, many English words are present in those recordings and are noticed by the consistent overlap between English words in the training sets of all the other languages.

Of the eleven languages, we treat six of them as well-resourced in our experiments that follow: Xhosa (Xho*), siSwati (Ssw*), isiNdebele (Nbl*), Sepedi (Nso[†]), Setswana (Tsn[†]) and English (Eng[‡]).² We use labelled data from the well-resourced languages to train a single supervised multilingual model and then apply the model to the target zero-resource languages. More specifically, we extract 100k true positive word pairs using forced alignments for each training language. We select Zulu (Zul*), Sotho (Sot[†]), Afrikaans (Afr[‡]), and Xitsonga (Tso) as our zero-resource languages and use another language, Tshivenda (Ven), for validation of each model. Training and evaluation languages are

²We use superscripts to indicate the different language families: *Nguni, [†]Sotho-Tswana, [‡]Germanic.

5.4. EXPERIMENTAL SETUP

	<i>Xho</i> [*]	<i>Zul</i> [*]	<i>Ssw</i> [*]	<i>Nbl</i> [*]	<i>Nso</i> [†]	<i>Tsn</i> [†]	<i>Sot</i> [†]	<i>Afr</i> [‡]	<i>Eng</i> [‡]	<i>Tso</i>	<i>Ven</i>
<i>Xho</i> [*]	100.0	12.28	2.929	6.943	0.7769	0.5571	0.8982	0.1251	1.038	0.595	0.8262
<i>Zul</i> [*]	14.14	100.0	4.53	9.736	0.8248	0.707	1.017	0.1527	1.261	0.6546	0.8292
<i>Ssw</i> [*]	6.841	9.187	100.0	7.31	1.381	1.221	1.814	0.3098	2.469	1.425	1.372
<i>Nbl</i> [*]	13.24	16.12	5.969	100.0	1.409	1.293	1.922	0.383	3.288	1.676	1.481
<i>Nso</i> [†]	1.873	1.727	1.425	1.781	100.0	12.64	17.8	0.5116	3.526	2.348	3.983
<i>Tsn</i> [†]	2.845	3.135	2.671	3.464	26.79	100.0	28.47	0.929	7.741	3.368	5.825
<i>Sot</i> [†]	2.302	2.263	1.991	2.584	18.93	14.29	100.0	0.6993	6.303	2.923	5.012
<i>Afr</i> [‡]	0.4405	0.4672	0.4672	0.7074	0.7475	0.6407	0.961	100.0	2.082	0.2803	0.5873
<i>Eng</i> [‡]	4.09	4.313	4.164	6.791	5.761	5.97	9.687	2.328	100.0	2.582	5.045
<i>Tso</i>	2.697	2.576	2.765	3.985	4.414	2.989	5.17	0.3607	2.971	100.0	5.582
<i>Ven</i>	2.928	2.552	2.082	2.754	5.856	4.043	6.931	0.591	4.54	4.365	100.0

Figure 5.4: The percentage of words in the training set of each language (row) that appear in the training set of every other language (column).

carefully selected such that for each evaluation language at least one language from the same family is part of the training languages, except for Tso which is in a group of its own.

5.4.2 MODELS

We use the CONTRASTIVERNN presented in Chapter 3 as the embedding function for all experiments in this chapter. The CONTRASTIVERNN configuration and hyperparameter selection is the same as the multilingual CONTRASTIVERNN models trained in Chapter 3. The encoder unit consists of three unidirectional RNNs with 400-dimensional hidden vectors, with an embedding size of $M = 130$ dimensions. All models are optimised using Adam optimisation [103] with a learning rate of 0.001. The temperature parameter τ in Equation 3.6 is set to 0.1.

5.4.3 EVALUATION

We again use the same-different word discrimination task (Section 2.5) to measure the quality of the AWEs. Additionally, we perform QbE analysis to measure the performance of the AWEs in a downstream zero-resource task (Section 5.3).

5.4. EXPERIMENTAL SETUP

5.4.3.1 WORD DISCRIMINATION

As a recap from Section 2.5, to evaluate a particular AWE model, a set of isolated test word segments is embedded. The number of word pairs in the test set for each language is shown in Table 5.1. For every word pair in this set, the cosine distance between their embeddings is calculated. Two words can then be classified as being of the same or different type based on some distance threshold, and a precision-recall curve is obtained by varying the threshold. The area under this curve is used as the final evaluation metric, referred to as the average precision (AP).

5.4.3.2 QUERY-BY-EXAMPLE

We evaluate a particular AWE in the QbE speech system described in Section 5.3. For each evaluation language, we use approximately two hours of test utterances as the search collection. Sub-segments for the utterances in the speech collection are obtained by embedding windows stretching from 20 to 60 frames with a 3-frame overlap. For each evaluation language, we randomly draw instances of 15 spoken query word types from a disjoint speech set (the development set—which we never use for any validation experiments) where we only consider query words with at least 5 characters for Afr and Zul and 3 for Sot. There are between 6 and 51 occurrences of each query word and appears 110 to 49 times in the search collection. For each QbE test, we ensure that the relevant multilingual AWE model has not seen any of the search or query data during training or validation.

To evaluate the QbE retrieval accuracy, we calculate the precision at ten ($P@10$), which is the fraction of the ten top-scoring retrieved utterances from the search collection that contains the given query. For example, assume the utterances retrieved by a QbE system that contains the query are at rank 1,2,4,7,8 among the ten top-ranked utterances. For that single query, the $P@10$ will be $\frac{5}{10}$, given that only five utterances contain the occurrence of the query, while the other five were false retrievals. The reported $P@10$ values are the averages calculated over all query types,

$$\text{average } P@10 = \frac{1}{K} \sum_{k=1}^K (P@10)_k, \quad (5.1)$$

where K is the number of query types and $(P@10)_k$ indicates the average precision at ten for all query items from type k .

As a baseline, we also give the results where DTW (Section 2.2.2) is used directly on the MFCCs of the segmented speech collection.

5.5. EXPERIMENTS

5.5 EXPERIMENTS

5.5.1 CROSS-LINGUAL EVALUATION

Before looking at multilingual modelling, we first consider a cross-lingual evaluation where we treat each language as a training language, train a supervised monolingual AWE model, and then apply it to every other language. This allows us to see the effect of training on related languages in a pairwise fashion. The results are shown in Figure 5.5. The result for training and evaluating a model on the same language is given along the diagonal, showing the top-line performance. For each evaluation language excluding Tso and Ven, which are in family groups of their own, the best results are achieved from models trained on a language from the same family. For example, on Zul, Xho is the best training language giving an AP of 58.5%. Eng is the only exception where the model trained on Sot performs better than using Afr, the other Germanic language. Although Ven is in its own group at the lowest layer of the family tree in Figure 5.2, some of the best results when evaluating on Ven are obtained using models trained on Sotho-Tswana languages (Sot, Tsn, Nso), which

	<i>Xho</i> [*]	<i>Zul</i> [*]	<i>Ssw</i> [*]	<i>Nbl</i> [*]	<i>Nso</i> [†]	<i>Tsn</i> [†]	<i>Sot</i> [†]	<i>Afr</i> [‡]	<i>Eng</i> [‡]	<i>Tso</i>	<i>Ven</i>
<i>Xho</i> [*]	84.0	58.5	62.3	55.2	24.2	55.1	43.0	25.8	50.5	55.6	32.3
<i>Zul</i> [*]	62.9	79.7	62.6	52.8	20.2	47.3	35.8	19.9	34.4	46.8	25.2
<i>Ssw</i> [*]	45.2	50.0	84.4	51.0	26.0	56.9	44.1	20.1	28.6	53.5	26.0
<i>Nbl</i> [*]	61.5	58.1	60.7	78.1	22.8	56.9	44.9	23.2	34.8	53.9	32.2
<i>Nso</i> [†]	43.6	23.2	46.6	31.6	72.4	72.3	61.9	31.8	46.7	60.2	44.9
<i>Tsn</i> [†]	37.8	27.0	45.5	38.2	43.5	88.8	65.6	30.6	36.9	54.1	38.2
<i>Sot</i> [†]	33.2	28.0	45.4	38.8	48.5	81.6	83.4	36.4	57.2	56.5	40.7
<i>Afr</i> [‡]	14.9	10.8	27.4	16.9	12.5	39.5	28.8	90.5	42.7	30.7	14.4
<i>Eng</i> [‡]	16.9	12.1	26.8	17.5	15.4	45.6	28.9	47.2	95.7	35.7	17.0
<i>Tso</i>	43.0	25.8	44.3	34.2	32.4	57.4	49.8	26.2	31.6	92.1	44.9
<i>Ven</i>	36.3	19.1	33.6	29.9	39.1	62.8	50.5	21.2	29.0	67.5	82.0

Figure 5.5: AP (%) when training a monolingual supervised CONTRASTIVERNN on each language (rows) and then evaluating it on each of the other languages (columns). Heatmap colours are normalised for each evaluation language (i.e. per column).

5.5. EXPERIMENTS

	<i>Xho</i> [*]	<i>Zul</i> [*]	<i>Ssw</i> [*]	<i>Nbl</i> [*]	<i>Nso</i> [†]	<i>Tsn</i> [†]	<i>Sot</i> [†]	<i>Afr</i> [†]	<i>Eng</i> [†]	<i>Tso</i>	<i>Ven</i>
<i>Xho</i> [*]	67.2	55.8	61.0	54.3	21.9	51.0	41.9	17.7	25.5	54.6	30.6
<i>Zul</i> [*]	40.4	61.5	49.7	45.6	18.8	42.6	27.2	16.7	37.0	45.3	21.0
<i>Ssw</i> [*]	45.5	43.8	73.9	36.8	21.6	50.2	35.3	20.8	25.9	47.3	24.5
<i>Nbl</i> [*]	37.4	47.0	53.4	61.8	21.2	54.4	35.3	23.5	25.6	46.3	25.5
<i>Nso</i> [†]	34.5	21.0	44.7	34.6	53.3	73.2	54.9	27.4	30.7	54.1	38.7
<i>Tsn</i> [†]	33.9	21.6	37.7	26.6	42.8	73.1	48.5	27.3	27.6	47.0	25.9
<i>Sot</i> [†]	27.4	19.1	43.6	22.5	32.8	67.5	52.9	25.1	26.7	43.6	26.6
<i>Afr</i> [†]	14.3	9.73	27.2	15.0	13.2	38.7	24.6	89.0	48.1	32.5	16.3
<i>Eng</i> [†]	14.7	13.2	30.1	15.3	14.3	45.0	29.2	52.1	83.4	33.6	16.2
<i>Tso</i>	34.0	21.2	39.5	27.1	30.3	54.0	43.8	25.0	26.9	61.8	32.5
<i>Ven</i>	38.0	16.9	36.7	27.9	30.3	65.0	41.1	24.6	26.3	55.9	54.2

Figure 5.6: AP (%) when training a monolingual supervised CONTRASTIVERNN on each language (rows) and then evaluating it on each of the other languages (columns) with overlapping words removed from training and evaluation data. Heatmap colours are normalised for each evaluation language (i.e. per column).

are in the same family at a higher level. We also see that for all nine Bantu evaluation languages, the worst performance is obtained from the two Germanic models (Afr, Eng).

Comparing Figure 5.5 to Figure 5.4, we notice a high correlation between the AP and the percentage overlap between the training data of two languages. For example, monolingual models trained on Zul and Nbl with 14.14% and 13.24% overlap gives an AP of 62.9% and 61.5%, respectively. On the other hand, a model trained on Ssw and evaluated on Xho with only 6.84% overlap gives an AP of 45.2%. Although all three languages belong to the same language family as Xho, the two with a higher percentage overlap outperform the one with less overlap. Moreover, we suspect that the model trained on Sot outperforms the model trained on Afr (from the same language family as Eng) when evaluating on Eng because of the high number of Eng words present in the training set of Sot, showing a 9.69% overlap.

This begs the question of whether the increased performance of training on a language related to the target language (Figure 5.5) is related to the phonetic properties shared among them, or because a model is exposed to training examples from the target language. We investigate this uncertainty by reproducing the experiment yielding the results in

5.5. EXPERIMENTS

Figure 5.5, this time removing the overlapping word segments from the training and validation sets. The results without overlapping data are displayed in Figure 5.6. A monolingual model is trained on data on each language where all words that overlap with any other language is removed, then applied to every other language where words from the training language are removed from the evaluation set.

First, we observe a performance decrease when training and evaluating on the same language (shown along the diagonal), especially for the Bantu languages. Secondly, for most languages, the AP considerably decreased when training a model on one language and applying it to another from the same language family. Importantly, although the performance decreased, the best results are still achieved by training on related languages, except for the one model where training on Ven outperforms a model trained on Nbl by 0.6% when evaluated on Xho. Finally, for most languages, the AP decreased on Eng after the Eng words were removed from the training data of all the other languages. For example, the AP of Sot dropped from 57.2% to 26.7%, with Afr (same language family as Eng) now achieving the best AP (11.1% higher than second best). This experiment indicates the potential of using related languages in multilingual AWE modelling which we explore in the following section.

5.5.2 MULTILINGUAL EVALUATION

The cross-lingual experiment above was in large part our inspiration for the subsequent analysis of multilingual models. Concretely, we hypothesise that even better performance can be achieved by training on multiple languages from the same family as the target zero-resource language and that this would be superior to multilingual models trained on unrelated languages. Focusing on two evaluation languages, Zul and Sot from distinct language families, we investigate this hypothesis by training multilingual models with different language combinations, as shown in Table 5.2. Here we use the original data without disregarding any particular words since it better simulates a real-world scenario. Firstly, we see the best result on a language when all the training languages come from the same family as the target. Secondly, we see how the performance gradually decrease as the number of training languages related to the target drop. Furthermore, notice the performance boost from including even just one training language related to the target compared to not including any. For instance, on Sot we see an increase of more than 12% absolute when adding just one related language (from 52.5% and 51.9% to 64.8%).

To further demonstrate the benefit of using training languages from the same family, we train a multilingual model for each evaluation language on all its related languages using a 10% subset of the original data. For both Zul and Sot, the subset models outperform the models where no related languages are used. E.g. on Zul, the Xho+Ssw+Nbl subset model outperforms the full Tsn+Nso+Eng model (no related languages) by more than 20% in

5.5. EXPERIMENTS

Table 5.2: AP (%) on test data for multilingual models trained on different combinations of well-resourced languages. Models are applied to two zero-resource languages from different language families, Nguni and Sotho-Tswana. For each training language, 100k word pairs were extracted.

Multilingual model	Zul*	Sot [†]
<u><i>Nguni:</i></u>		
Xho* + Ssw* + Nbl*	68.6	—
Xho* + Ssw* + Eng [‡]	60.9	—
Xho* + Nso [†] + Eng [‡]	55.7	—
Tsn* + Nso [†] + Eng [‡]	37.5	—
Xho* + Ssw* + Nbl* (10% subset)	58.6	—
<u><i>Sotho-Tswana:</i></u>		
Nso [†] + Tsn [†]	—	76.7
Nso [†] + Eng [‡]	—	64.8
Xho* + Ssw*	—	51.9
Xho* + Eng [‡]	—	52.5
Nso [†] + Tsn [†] (10% subset)	—	58.4

AP. Moreover, this subset model (58.6% in AP) even outperforms the Xho+Nso+Eng model (55.7%) where all the training data from one related language are included and almost matches the AP when using two related languages (Xho+Ssw+Eng, 60.9%).

These comparisons do more than just show the benefit of training on related languages: they also show that it is beneficial to train on a diverse set of related languages.

5.5.3 ADDING MORE LANGUAGES

In the above experiments, we controlled for the amount of data per language and saw that training on languages from the same family improves multilingual transfer. But this raises a question: will adding additional unrelated languages harm performance? This question is of practical relevance in cases where related data is either unavailable or the system developer is unsure for which language AWEs will be extracted. To answer this, we systematically train two sequences of multilingual models on all six well-resourced languages, evaluating each target language as a new training language is added.

Same-different results for all five evaluation languages are shown in Figure 5.7 and QbE results for Zul, Sot and Afr are shown in Figure 5.8. Notice how the trends in the same-different and QbE results track each other closely for corresponding evaluation languages. In the first sequence of multilingual models (green), we start by adding the three Nguni languages (Xho, Ssw, Nbl), followed by the two Sotho-Tswana languages (Nso, Tsn), and lastly the Germanic training language (Eng). The second sequence (orange) does not follow a systematic procedure.

5.5. EXPERIMENTS

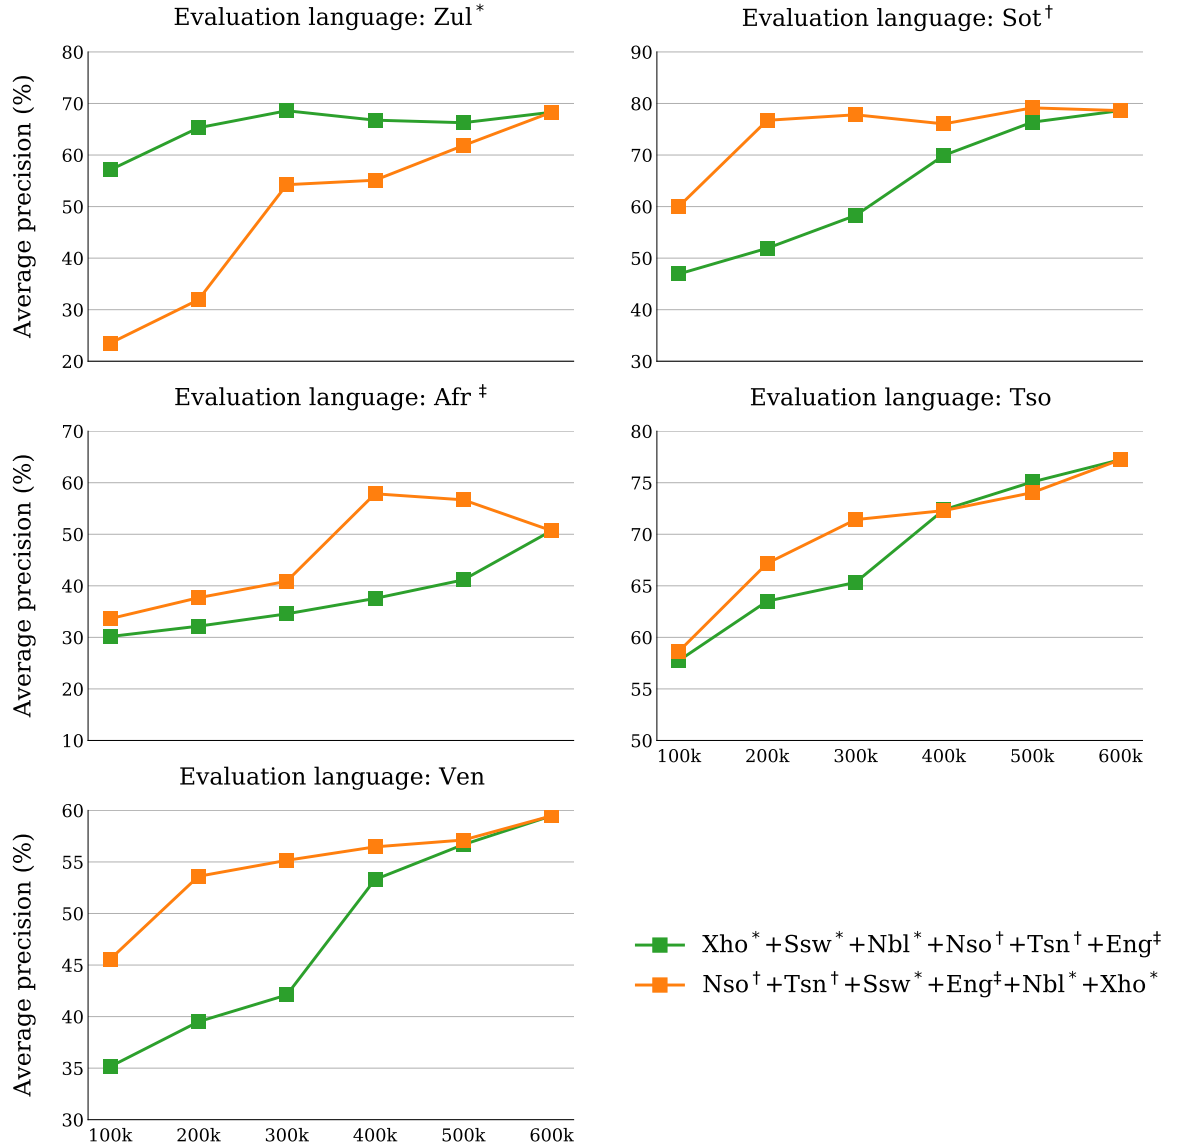


Figure 5.7: Same-different results from two sequences of multilingual models, trained by adding one language at a time. For each training language, 100k positive word pairs are used, which is indicated on the x -axis.

On Zul, the green sequence, which starts with a related language (Xho), initially achieves a higher score compared to the orange sequence in both Figures 5.7 and 5.8. Then, the score gradually increase by adding more related languages (Ssw, Nbl). Thereafter, adding additional unrelated languages (Nso, Tsn, Eng) show no performance increase. In fact, AP decreases slightly after adding the two Sotho-Tswana languages (Nso, Tsn), but not significantly. The orange sequence starts low on Zul until the first related language (Ssw) is added, causing a sudden increase. Adding the Germanic language (Eng) has little effect. Adding the last two related languages (Nbl, Xho) again causes the score to increase. A similar trend follows for Sot and Afr, where adding related languages causes a noticeable performance increase, especially when adding the first related language; after this, performance seems to plateau when adding more unrelated languages. (Afr is the

5.6. CHAPTER SUMMARY




Figure 5.8: QbE results on Zul, Sot and Afr, using the same sequences of multilingual models as in Figure 5.7. DTW result are shown for baseline performance.

one exception, with a drop when adding the last language in the orange sequence). On Tso, which does not have any languages from the same family in the training set, AP gradually increases in both sequences without any sudden jumps. Although Ven is not in the same family as Nso and Tsn at the lowest level of the tree in Figure 5.2, it belongs to the same family (Sotho-Tswana) at a higher level. This explains why it closely tracks the Sot results.

Summarising these results, we see that adding unrelated languages generally does not decrease scores, but also does not provide a big benefit (except if it is one of the earlier languages in the training sequence, where data is still limited). In contrast, it seems that training on languages from the same family is again beneficial; this is especially the case for the first related language, irrespective of where it is added in the sequence.

5.6 CHAPTER SUMMARY

We investigated the effect of training language choice when applying a multilingual AWE model to a zero-resource language. Using word discrimination and QbE search tasks on languages spoken in South Africa, we showed that training a multilingual model on languages related to the target is beneficial. We observed gains in absolute scores, but also in data efficiency: you can achieve similar performance with much less data when training on multiple languages from the same family as the target. We showed that even including

5.6. CHAPTER SUMMARY

just one related language already gives a large gain. From a practical perspective, these results indicate that one should prioritise collecting data from related languages (even in modest quantities) rather than collecting more extensive datasets from diverse unrelated families, when building multilingual acoustic word embedding models for a zero-resource language. In the following chapter, we present a framework for producing semantically organised AWEs using multilingual models developed in this chapter.

CHAPTER 6

HATE SPEECH DETECTION USING MULTILINGUAL ACOUSTIC WORD EMBEDDINGS

Leading up to this chapter, we mainly aimed at improving the quality of AWEs for low-resource languages: we developed a new AWE model (Chapter 3), explored multilingual adaptation techniques (Chapter 4), and considered language choice in multilingual AWE modelling (Chapter 5). In Chapter 5 we briefly measured AWE performance in a downstream query-by-example (QbE) speech search task. However, this was performed with limited development data in a controlled environment. In this chapter, we evaluate the practicality of using multilingual AWEs for real-world applications, beyond the experimental environments considered up to now.

The content of this chapter originates from a collaborative project with VoxCroft, a company specialising in developing speech applications in low-resource settings. We address a specific case study involving the development of keyword spotting (KWS) systems to monitor radio broadcast audio for low-resource languages. Here, we confront the challenge of lacking transcribed data in the target language and dealing with radio recordings that differs significantly from experimental audio data. In this challenging environment, we develop an ASR-free approach, utilising a pre-trained multilingual AWE model in a KWS system. Our main objective is to compare the performance of this AWE-based KWS system to the conventional ASR-based KWS systems, which mostly rely on transcribed data from the target language. We are particularly concerned about the trade-off between accuracy and feasibility for rapid deployment. To measure performance in this in-the-wild environment, we perform our final experiments on real Swahili radio broadcast audio, scraped from local radio stations in Kenya. Moreover, our KWS systems are specifically developed to detect instances of hate speech in radio broadcasts. Therefore, to search the audio, we use a set of keywords carefully selected by domain experts that are potential markers of hate speech.

6.1. BACKGROUND

The content of this chapter was presented at Interspeech:

C. Jacobs, N. C. Rakotonirina, E. A. Chimoto, B. A. Bassett, and H. Kamper, “Towards hate speech detection in low-resource languages: Comparing ASR to acoustic word embeddings on Wolof and Swahili,” in *Proceedings of Interspeech*, 2023.

Contributions per author:

C. Jacobs developed the AWE-based KWS systems (Section 6.4.2).

N. C. Rakotonirina performed the ASR experiments (Section 6.4.1).

E. A. Chimoto evaluated radio recordings flagged by our KWS systems (Section 6.5.3.2).

B. A. Bassett and H. Kamper supervised the project and helped with editing the paper.

6.1 BACKGROUND

The rise of the internet and communication platforms are increasing the spread of content, some of which include the use of derogatory or discriminatory language targeting individuals or groups based on their identity or characteristics. This form of expression is commonly referred to as *hate speech*. Given the overwhelming volume of content being generated on online platforms, several automatic monitoring systems have been explored to curb the spread of hate speech (usually in the form of written text), especially for resource-rich languages.

However, in many developing countries where there is a lack of technological advancement, radio remains the primary medium for communicating news and information to the public [111]. These radio stations play a vital role as platforms for public discourse, facilitating freedom of expression [112]. Therefore, it is essential to establish appropriate mechanisms for monitoring these radio stations to ensure responsible broadcasting practices.

A striking example of the consequences of hate speech over radio is the 1994 genocide in Rwanda where propaganda was disseminated through local radio stations targeting individuals from an ethnic minority. The dehumanising language used on air actively fuelled violence, resulting in devastating consequences. Similarly, in the 2007-2008 post-election violence in Kenya, vernacular radio stations played a significant role in exacerbating the conflict [113].

In response to this ongoing issue, the United Nations has taken proactive measures to address hate speech on radio [59]. By actively monitoring hate speech claims, authorities can promptly respond and mitigate instances, preventing tension escalation and the dissemination of harmful ideologies. The efforts of this chapter are in line with the objective of combating hate speech on radio broadcasts, specifically for low-resource languages.

6.2. RELATED WORK

Hate speech detection can be supported by using a *keyword spotting* (KWS) system that can effectively search an audio corpus for a predetermined set of keywords or phrases that are indicative of hate speech. Such a KWS system has the potential ability to quickly scan through large amounts of audio content and flag parts of the audio that have been identified as potential markers of hate speech. It is important to note that a KWS system solely relying on keyword or phrase identification may not be sufficient for comprehensive hate speech detection. For instance, given the keyword “blood”, a KWS system may retrieve utterances “we must spill their blood and eradicate them from our society” and “she donated blood to help save lives during the blood drive.” The former contains hate speech whereas the latter highlights a positive action. Therefore, the context in which these potential hate speech markers occur should be carefully considered and analysed by domain experts.

6.2 RELATED WORK

The conventional KWS approach involves transcribing the target audio with an automatic speech recognition (ASR) system and then searching for keywords in the output text [60, 61]. For low-resource languages, the absence of large quantities of transcribed speech in the target language can make it difficult to train high-quality ASR models. Several studies have looked into this problem, proposing methods to train ASR models on limited amounts of labelled data [55, 57, 114–116].

Menon et al. [56] were the first to evaluate and report results of a KWS system monitoring radio broadcasts in low-resource languages spoken in Uganda. Their system was tailored to identify specific keywords, contributing to the progress of development and relief initiatives in rural areas. Their system utilises an ASR model trained on around nine hours of transcribed data from the target language. However, relying on this amount of data poses a challenge as it will require a significant amount of time to collect and transcribe the audio.

Saeb et al. [55] extended the work of Menon et al. [56] and implemented an ASR-based KWS system that combines multilingual acoustic modelling and semi-supervised training relying on much less transcribed data from the target language. Here, they only use twelve minutes of annotated data from the target language to train an ASR model. However, they use several hours of untranscribed audio from the target language, coupled with a self-supervised learning strategy to improve their system.

Rapidly deploying ASR KWS systems for low-resource languages remains challenging due to the time-consuming process of collecting even a few minutes of transcribed data, which can take several weeks. Moreover, multiple hours of untranscribed speech are not necessarily available for all target languages for self-supervised training.

When we want to rapidly develop and deploy a KWS system for hate speech detection in

6.2. RELATED WORK

a new setting, it might be better to opt for an ASR-free approach. In a truly zero-resource setting, ASR-free KWS systems are favourable, where KWS systems are developed without the need for transcribed data. The most popular methodology extends QbE (Section 2.2.3), which uses a *spoken* instance of a keyword to search through an audio corpus, without the need for transcribing the audio into text. To use QbE for KWS, we need a small number of spoken templates to serve as queries for the keywords of interest; this is the approach followed in [57, 62] and the methodology we also use for our AWE-based KWS system in this chapter.

DTW (Section 2.2.2) is typically used to match the speech features of a spoken query to a search utterance [9, 48–50]. However, DTW has drawbacks in terms of its high time complexity and limitations, such as poor speaker independence, often requiring a large template set.

AWEs have been proposed as an alternative approach to finding matches in QbE where search segments and query segments are jointly mapped to the same fixed vector space [6, 52, 53, 87]. This allows for faster comparisons between query and search segments in linear time complexity. Various neural networks have been considered to obtain these AWEs as described in Section 2.3.

As demonstrated in Chapter 3 and Chapter 5, one approach to quickly obtain robust AWE for a new language is to train a multilingual AWE model on labelled data from multiple well-resourced languages, and then apply it to the target low-resource language. Although these multilingual AWE models have proven successful in controlled experiments, there has been limited work investigating the effectiveness of these systems beyond the experimental environment.

In this chapter, for the first time, we develop and evaluate a KWS system utilising multilingual AWEs in a real-world scenario. For this experiment, we apply the system to real radio broadcast audio, specifically targeting the detection of instances of hate speech. The radio data was collected from stations in Kenya, a Sub-Saharan country in Africa, where the predominant language is Swahili. Additionally, we conducted developmental experiments where we included Wolof, another low-resource language spoken mainly in Senegal, Sub-Saharan Africa.

We compare the AWE-based KWS system to ASR-based KWS systems. We train our ASR models, leveraging advancements in self-supervised speech representation learning (SSRL) (Section 2.2.1). These models enable fine-tuning of ASR models using limited transcribed data from the target language. Our main research question is to compare the AWE-based KWS system with ASR-based KWS systems, focusing on the trade-off between accuracy and rapid deployment.

Before we describe the different KWS systems, we first briefly give the overall methodology that we follow towards our end goal of hate speech detection on Swahili radio broadcasts.

6.3 OVERVIEW OF METHODOLOGY

We start by developing Wolof and Swahili KWS systems in a controlled environment where we use parallel audio-text data from experimental datasets to train and test our models. These datasets [117, 118] contain high-quality, noise-free read speech. We then test our models on held-out test data from these corpora. We refer to these experiments as *in-domain* because the test data comes from the same domain as the training data. In this controlled setup, we perform KWS using a set of keywords drawn from development data. These are not necessarily hate speech keywords. The benefit of this controlled setting is that we have transcriptions for the test data, and can therefore quantify performance exactly.

We then apply the models to the *out-of-domain* Swahili radio broadcasts. These recordings were obtained from three radio stations in Kenya. They consist of a mix of read and spontaneous speech, and also include non-speech audio like music and advertisements. Additionally, there may be instances of different accents and languages present within the recordings. This is the type of real-world audio that we would like to monitor for hate speech using KWS. As our keywords in this context, we utilise actual hate speech keywords. These were labelled as inflammatory words by native Swahili experts familiar with the media environment. We should emphasise that, when we detect occurrences of these words, they might not necessarily correspond to hate speech. For example, the English sentence “do not kill your neighbour” is not hate speech despite containing the word “kill”. In a deployed system the utterances flagged by our KWS approach would be passed on to human moderators for further review and hate speech assessment. In this work we also use a human moderator, but only to mark whether a detected keyword actually occurred in an utterance. We do this for our in-the-wild test since we do not have transcriptions for the radio broadcasts.

In the following section, we provide a detailed overview of the KWS systems that we developed. This will be followed by the implementation details of the KWS systems (Section 6.5).

6.4 LOW-RESOURCE KEYWORD SPOTTING

The ASR-based KWS systems of this chapter were developed by a colleague. Therefore, in Section 6.4.1, we provide a high-level overview of this approach. The AWE-based KWS system we developed as an ASR-free alternative, is discussed in detail in the follow-up section.

6.4. LOW-RESOURCE KEYWORD SPOTTING

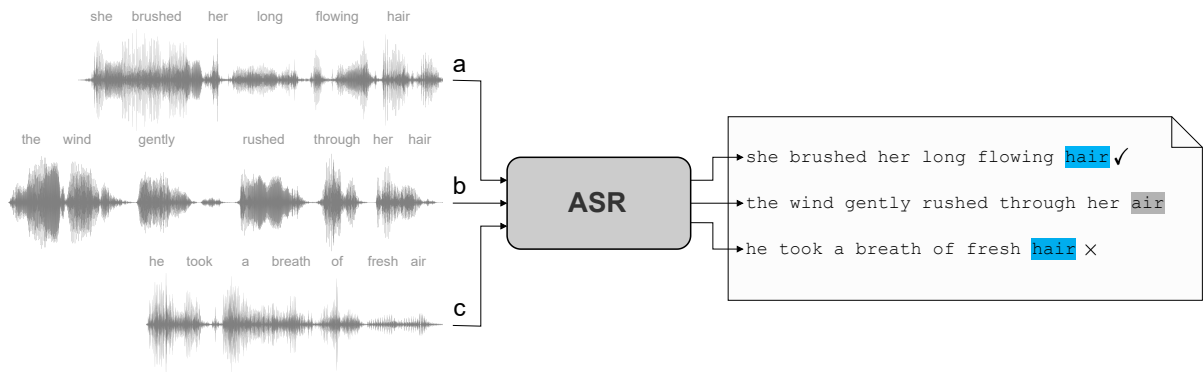


Figure 6.1: An illustration of ASR-based keyword spotting.

6.4.1 KEYWORD SPOTTING WITH ASR

In an ASR-based KWS system, the target corpus is first transcribed to text. Occurrences of a keyword in the corpus can be efficiently identified by simply matching the keyword tokens with the output transcriptions. An example of this ASR-based KWS is illustrated in Figure 6.1. In this example, when using the keyword “hair” to search the output transcriptions, the KWS system will correctly retrieve the first utterance (a). In the third utterance (c), “air” was incorrectly transcribed as “hair” and will therefore be an incorrect retrieval (although the keyword matches the transcription). In the second utterance (b), the word “hair” in the audio was incorrectly transcribed as “air” and will not be retrieved by the KWS system. Also, note that in this example the ASR model correctly transcribed all the other words and would therefore have a low word error rate (WER), but WER does not always correspond to KWS performance. In Section 6.5.3 we give a more detailed explanation of how we evaluate KWS performance.

Typically, to train an ASR model we need audio with transcriptions from the target language. In a low-resource setting, a lack of large quantities of transcribed speech data from the target languages prohibits the rapid development of ASR-based KWS systems.

For our ASR systems, we build on recent advancements in ASR models leveraging SSRL (Section 2.2.1). These SSRL models are trained in a self-supervised fashion using large amounts of unlabelled audio recordings to produce high-quality latent frame-level speech representations. It has been shown that these pre-trained SSRL models can be fine-tuned in an end-to-end fashion for ASR by mapping the latent frame representations to target text tokens. Fine-tuning SSRL models for ASR requires using much less training data than training an ASR model from scratch.

In our experiments, we use the pre-trained Wav2Vec2.0 XLSR model [76, 119]. This SSRL model is a large-scale cross-lingual speech model pre-trained on half a million hours of unlabelled speech data in 128 languages (including Swahili but not Wolof). The speech representation learned during cross-lingual pre-training improves performance for low-resource languages by leveraging data from the high-resource languages on which it is

6.4. LOW-RESOURCE KEYWORD SPOTTING




Figure 6.2: Search utterances (u_1, u_2, \dots, u_N) are segmented and mapped to a fixed-dimensional space, along with a query segment q , to produce embeddings $\{\mathbf{z}_n^{(i)}\}$ and \mathbf{z}_q , respectively. The subscript n represents a unique utterance within the search collection, and superscript i represents a segment extracted from utterance \mathbf{z}_n . Subsequently, search utterances are ranked based on the shortest distance between a search segment $\mathbf{z}_n^{(i)}$ and the query vector \mathbf{z}_q .

pre-trained. Using the pre-trained model as a starting point, the model is fine-tuned on labelled data from the language of interest.

Previous work has shown that this fine-tuning approach leads to competitive ASR performance even with limited amounts of labelled data [119]. In our experiments (Table 6.2 specifically) we show that the sample efficiency of the XLSR model also extends to KWS when training on just five minutes of labelled data.

6.4.2 KEYWORD SPOTTING WITH MULTILINGUAL ACOUSTIC WORD EMBEDDINGS

For a severely under-resourced language, there might not even be limited amounts of labelled data available for ASR fine-tuning. In this scenario, collecting only a few spoken instances of a keyword of interest would allow for KWS through QbE search (Section 2.2.3). QbE is the task of retrieving utterances from a search corpus using a spoken instance of a keyword, instead of a written keyword.

In the previous chapter, our results (on development data) showed improved performance in QbE search by using multilingual AWEs instead of DTW. We use a similar setup to the QbE experiments described in Section 5.3 for experiments in this chapter. The workings of a QbE system utilising AWEs are displayed in Figure 6.2. First, the search corpus is segmented into word-like segments. Search utterances are segmented by splitting into overlapping segments of some minimum to some maximum duration (see Figure 5.3). Concretely, utterances from the search corpus (u_1, u_2, \dots, u_N) are segmented to produce a set of variable-length speech segments $S = \{\mathbf{s}_n^{(i)}\}_{n=1}^N$, where i is the i th search segment extracted from utterance n .

6.5. EXPERIMENTAL SETUP

A multilingual AWE model (trained on labelled data from multiple well-resourced languages but not the target language) is then applied to the segmented search corpus S and query segment q , mapping them to the same fixed-dimensional vector space. In this experiment, we use the CAE-RNN network (Section 2.1.3) and train a multilingual CAE-RNN AWE model. This model is trained on word pairs using labelled data from multiple languages (but does not include the target language). This multilingual AWE model can be transferred and applied to any unseen language to embed word segments of arbitrary length. Utilising such a universal, language-independent model would enable fast deployment of practical QbE-based KWS systems.

Applying the AWE model to search segments S and query segment q produces search embeddings $\{\mathbf{z}_n^{(i)}\}$ and query embedding \mathbf{z}_q , respectively. The utterances in the search corpus are ranked based on comparing the query-segment distances $d(\mathbf{z}_q, \mathbf{z}_n^{(i)})$. For each query, the lowest distance for each utterance is used as the final utterance score. We can then predict that all query-utterance pairs with a distance below some threshold are positive matches (this requires a threshold to be tuned on development data). With these predictions, we can calculate common KWS metrics and compare them to those achieved by the ASR-based KWS systems.

6.5 EXPERIMENTAL SETUP

In this section, we present the implementation details of both the ASR-based (Section 6.4.1) and AWE-based (Section 6.4.2) KWS systems. It is important to note that the overview provided in Section 6.3 serves as a reminder for the reader. By the end of this section, the difference between the developmental experiments conducted in a controlled environment and our in-the-wild KWS testing, particularly concerning data and evaluation methods, should be clear.

6.5.1 DATA

Here we emphasise the difference between the training and evaluation data we use during development and the data we use for our in-the-wild testing.

6.5.1.1 DEVELOPMENT

For our controlled experiments, we train and evaluate our models on Swahili and Wolof data, collected from the Common Voice [118] and ALFFA [117] datasets, respectively. Common Voice is an open-source data collection initiative by Mozilla Foundation where volunteers can contribute by recording their voices while reading prompts collected from a variety of sources. These recordings are then verified for quality and accuracy. For our experiments we use the Swahili data from the Common Voice Corpus 8.0 dataset

6.5. EXPERIMENTAL SETUP

which gives train, test, and developing sets totalling 150 hours. In order to match the sampling rate used during the training of the XLSR model, the original recordings, which are sampled at 48 kHz, are downsampled to 16 kHz.

The current data repository of Common Voice does not include Wolof. We therefore use a separate dataset from the ALFFA project, which contains Wolof recordings and their corresponding transcriptions. Wolof is primarily spoken in Senegal, another country in Sub-Saharan Africa. However, the official language in Senegal is French, which serves as the lingua franca in all governmental institutions. Due to the prevalence of French as the official language, there are limited written documents available for Wolof to construct parallel training data. For the Wolof language, we have designated training, testing, and development sets, consisting only of 16 hours, 2 hours, and 2 hours, respectively.

For both datasets, we use the default train, development, and test splits. In our controlled experiments, we evaluate in-domain KWS performance on the Swahili and Wolof test splits using available transcriptions. The test split of both languages serves as the search collection and is not used during model development. The Swahili search collection contains 8 941 utterances and has a duration of approximately 14 hours, while the Wolof set contains 2 000 utterances and has a duration of approximately 2 hours.

6.5.1.2 IN-THE-WILD

We perform our in-the-wild test on recorded radio broadcasts (collected by VoxCroft) from three different radio stations in Kenya: Ghetto Radio, Radio Salaam, and Radio Taifa. We do not have any transcriptions, speaker information or content information on these audio clips. The nature of these radio recordings has distinct characteristics that are different from the prompted read speech we use during development. These broadcasts can contain a variety of elements, including music in the background, fast-paced speech by radio hosts, advertisements, background music, and dialogues involving multiple speakers.

Our systems are only applied to these out-of-domain recordings and are nowhere used during the development of our KWS systems. Before applying our systems to this data we first apply a preprocessing step where we segment the recordings using an unsupervised diarisation system that also attempts to remove music. Apart from this preprocessing step, we simply apply our systems to this out-of-domain data without further changes or calibration. After this preprocessing step, 19 716 utterances are obtained, which is our out-of-domain search collection. The duration of all utterances is between three and 30 seconds. Given that we do not have transcriptions of these recordings, we cannot evaluate system performance in the same way we do in the controlled experiments (where we have audio with corresponding transcriptions). As we explain below (Section 6.5.3.2), we use a domain expert to evaluate our system’s performance on this set.

6.5.2 MODELS

We now provide the implementation details of the ASR model and the multilingual AWE model for ASR-based (Section 6.4.1) and AWE-based KWS (Section 6.4.2), respectively.

6.5.2.1 ASR MODEL

To train our ASR models, we fine-tune the pre-trained XLSR model by adding a projection head to the final transformer layer. Subsequently, the model is trained using transcribed data from the target language by mapping the frame-level representations to character tokens. The model is optimised using the connectionist temporal classification (CTC) loss [120]. The CTC loss is effective for aligning input sequences and producing output sequences that may not be temporally aligned, as is the case with speech and text. After the fine-tuning step, the model produces a sequence of characters given an arbitrary audio recording, that is then decoded to give the predicted transcription.

For the ASR-based KWS systems, we control the amount of training data and train three ASR models for Wolof and Swahili. Specifically, we fine-tune the XLSR model [119] for each language on 30 hours, one hour, and five minutes of labelled training data. These models are used for our development experiments before applying them to the radio broadcast recordings.

6.5.2.2 MULTILINGUAL ACOUSTIC WORD EMBEDDING MODEL

For the AWE-based KWS system (Section 6.4.2), we train a multilingual AWE model on five different Common Voice languages: Abkhazian (ab), Czech (cs), Basque (eu), Swedish (sv), Tamil (ta). Word boundaries are obtained using the Montreal forced-aligner [121]. We pool the data from all languages and extract 300k positive word pairs which we use to train a multilingual CAE-RNN AWE model similar to the CAR-RNN model of Chapter 3. We use another well-resourced language, Bulgarian, for validation during training using the same-different evaluation (Section 2.5). The encoder and decoder each consist of three unidirectional RNNs with 400-dimensional hidden vectors and an embedding size of 100 dimensions. As inputs to the encoder, we use the XLSR model to extract speech features (Section 2.2.1) instead of MFCCs used for AWE modelling in previous chapters. We use the output of the 12th transformer layer, producing input features with 1024 dimensions. In the work of [122], they show that the lower layers of the XLSR transformer model capture more discriminative phonetic properties than the final layers when simply using the model to extract frame-level speech representations. Preliminary experiments showed using these features outperforms MFCCs in AWE modelling. We provide a short analysis in Section 6.7.1 comparing multilingual AWEs trained on MFCC and XLSR. All models are optimised using Adam [103] with a learning rate of 0.001.

6.5.3 EVALUATION

We outline our KWS evaluation methods for the ASR-based and AWE-based KWS systems for the two scenarios: controlled in-domain experiments with available transcriptions and out-of-domain in-the-wild tests where transcriptions are unavailable for evaluation.

6.5.3.1 CONTROLLED EXPERIMENTS

To measure in-domain KWS performance, we apply the ASR-based and AWE-based KWS systems to the same search corpora. Each utterance in the search corpus is assigned a label of 1 if the keyword is detected by the respective KWS system, or 0 if the keyword is not present. We evaluate our KWS systems by using the predicted labels and compare them against the true labels obtained from available transcriptions. Following this approach, we can calculate precision and recall scores.

Precision measures the accuracy of a KWS system in predicting whether an utterance contains a keyword. It is calculated as the number of utterances correctly predicted as keyword-containing utterances (# hits) divided by the total number of utterances predicted to contain a keyword (# keywords detected):

$$P = \frac{\# \text{ hits}}{\# \text{ keywords detected}} \quad (6.1)$$

Recall quantifies the KWS system's ability to retrieve keyword-containing utterances from the search corpus. It is calculated as the number of correctly predicted keyword-containing utterances (# hits) divided by the total number of keyword-containing utterances in the search corpus (# keywords):

$$R = \frac{\# \text{ hits}}{\# \text{ keywords}} \quad (6.2)$$

In addition to precision and recall, we report the F_1 -score to quantify the trade-off between these two metrics. The F_1 -score represents the balanced combination of precision and recall, calculated as the harmonic mean of the two metrics:

$$F_1 = 2 \times \frac{P \times R}{P + R} \quad (6.3)$$

To evaluate our ASR KWS systems, we simply check whether a keyword occurs in the predicted transcript after applying it to the search corpus and label each utterance accordingly (1 or 0) to calculate the metrics above.

The AWE-based KWS system does not output transcribed text but a query-utterance score for each query segment (bottom of Section 6.4.2). First, sub-segments for utterances in the search collection are obtained by extracting windows ranging from 20 to 35 frames with a 5-frame overlap (we experimented with different window and overlap sizes on development data where this configuration worked best). Then, an AWE model is applied

6.5. EXPERIMENTAL SETUP

Table 6.1: Swahili hate speech keywords used for in-the-wild KWS on Swahili radio broadcasts, with their English translations.

Swahili	English	Swahili	English
vita	war	wezi	thieves
damu	blood	majimbo	states
hama	move	wakora	conmen
kabila	tribe	panga	machete
utapeli	fraud	takataka	garbage
mende	cockroaches	mjinga	stupid
kitendawili	riddle	fala	stupid

to all the sub-segments and query segment. Finally, the similarity between a query and search segment is calculated using cosine distance.

We use a set of ten templates of 36 unique keywords with a minimum character length of five for Swahili, and 15 unique keywords with a minimum character length of four, for Wolof. These keywords were randomly sampled from a larger set of words that appear at least ten times in both the development and test sets. These keywords do not represent hate speech keywords. Query templates are drawn from development data, embedded, and averaged to obtain a single AWE embedding representing a keyword. Preliminary experiments showed that averaging all templates from the same keyword improves KWS performance compared to using a single template for searching. The threshold value in the AWE-based KWS is tuned for the highest F_1 -score across all keywords on the controlled test data (more in Section 6.7.2).

6.5.3.2 IN-THE-WILD

For the in-the-wild KWS, we apply the systems directly to the out-of-domain Swahili radio broadcasts. We use a set of keywords labelled as inflammatory by expert analysts from VoxCroft who are familiar with the media environment for the purpose of hate speech detection. These keywords with their English translations are given in Table 6.1. For the AWE-based KWS system, ten query templates per keyword are extracted from the in-domain Swahili data. We evaluate KWS performance by asking each approach to give 100 utterances out of the search collection that are most likely to contain any of the hate speech keywords. In the absence of per-word confidence scores for the ASR models, here we simply take 100 random utterances that were predicted to contain a keyword (for most models, the total number of utterances was around 150).

For the AWE-based KWS, we use the 100 highest-ranked utterances. Because we do not have transcriptions for this data, we provide a native Swahili speaker with untranscribed recordings of the utterances that were predicted to contain a keyword. They then mark

6.6. EXPERIMENTS

Table 6.2: ASR and AWE KWS results (%) on in-domain test data. For the ASR systems, the XLSR model is fine-tuned on each target language, controlling the amount of training data. For the AWE KWS system, a supervised multilingual AWE is trained on multiple well-resourced languages and applied to the two low-resource target languages.

Model	Swahili			Wolof		
	Prec.	Rec.	F_1	Prec.	Rec.	F_1
<u>ASR:</u>						
XLS-R (30-h)	97.6	98.6	98.0	93.1	87.1	90.0
XLS-R (1-h)	96.4	76.1	85.1	93.1	74.3	82.7
XLS-R (5-min)	95.4	42.6	58.9	88.0	45.6	60.0
<u>Multilingual AWE:</u>						
CAE-RNN (ab+cs+eu+sv+ta)	57.1	56.1	56.6	44.0	58.4	50.2

whether the keyword was indeed present. Note that due to the absence of transcriptions for the radio recordings, we are unable to calculate recall or the F_1 -score. Therefore, we can only measure performance by calculating precision, based on 100 retrieved utterances per system (Table 6.4).

6.6 EXPERIMENTS

We first look at the results from our controlled test on in-domain Swahili and Wolof data. We then look at the results on the out-of-domain data, which is of particular interest to us as it would indicate the effectiveness of both these systems for real-world application.

6.6.1 KWS RESULTS IN A CONTROLLED ENVIRONMENT

The first three lines in Table 6.2 report the ASR KWS results on in-domain Swahili and Wolof data. The results show that fine-tuning XLSR on only five minutes of labelled data achieves decent KWS performance, with F_1 -scores of 59% and 60% for Swahili and Wolof, respectively. The multilingual AWE KWS system using roughly one minute of template data, performs worse on most metrics. The only metric on which the AWE system is better is recall, where it achieves higher scores than the 5-minute ASR system on both languages. In practice, the AWE system could therefore be a better option if recall is important.

We briefly consider how ASR performance is affected by the amount of training data. For reference, WERs for the 30-hour, 1-hour, and 5-minute models are respectively 9%, 36%, and 62% on Swahili. For Wolof, the WERs are 27%, 44%, and 68%. As expected, more training data gives better WER and KWS scores. But it is noteworthy that with just five minutes of training data, we can already spot keywords with high precision (95% and

6.6. EXPERIMENTS

Table 6.3: Supervised AWE KWS results (%) on in-domain Swahili test data. The supervised monolingual CAE-RNN AWE model is trained using labelled Swahili data. It is applied with two segmentation configurations: using ground-truth word boundaries (true segm.), and using a variable-length window which is swept across the search collection (random segm.). Here we also report the standard QbE metrics $P@10$ and $P@N$ (Section 5.4.3.2).

Model	Prec.	Rec.	F_1	$P@10$	$P@N$
<i>Supervised monolingual:</i>					
CAE-RNN (true segm.)	95.5	90.8	93.1	98.6	94.3
CAE-RNN (random segm.)	79.2	76.1	77.6	92.2	90.9
<i>Supervised multilingual:</i>					
CAE-RNN ab+cs+eu+sv+ta (random segm.)	57.1	56.1	56.6	87.8	64.7

88% for Swahili and Wolof, respectively). This is especially useful for our use case, where we want to rapidly develop KWS applications in severely low-resourced settings. The other noteworthy finding from the ASR results is that scores for Swahili are not notably higher than Wolof, although the former is one of XLSR’s pretraining languages.

We now turn to the AWE-based approach. Specifically, we ask what the upper bound on performance would be if we had more training data or a more idealised search setting. Therefore, for a moment, we assume we have labelled data available and perform top-line experiments on the in-domain Swahili data, shown in Table 6.3. A supervised Swahili AWE model is trained on 30 hours of labelled data and applied to the search collection that is segmented using true word boundaries. Results are shown in the first row of Table 6.3, serving as the top-line performance for AWE KWS, with a high F_1 -score of 93%. Compared to the 30-hour ASR system (Table 6.2), the idealised AWE system comes closer to the F_1 of 98%. The second row in Table 6.3 shows KWS performance using a supervised Swahili AWE model but applied to search segments extracted by sliding a variable-length window across the search collection—the way we apply the AWE approach in practice and in Table 6.2. Segmentation without true word boundaries incurs a significant penalty, resulting in an F_1 -score drop to 77%. It is clear that the sliding window approach has a large effect on downstream performance, so future work should consider more sophisticated unsupervised word segmentation techniques.

6.6.2 KWS RESULTS IN THE WILD

We now turn to our main research question: comparing ASR to multilingual AWE-based KWS on real-life, out-of-domain audio in a low-resource setting. As mentioned in Section 6.5.3.2, systems are applied to out-of-domain Swahili radio broadcasts, after which, for each system, the top 100 utterances predicted to contain a hate speech keyword

6.6. EXPERIMENTS

Table 6.4: In-the-wild KWS results (%) on Swahili radio broadcasts. For each system, we report the percentage of utterances retrieved containing a keyword (precision) and the percentage of utterances retrieved containing music.

Model	Precision	Music
<u>ASR:</u>		
XLS-R (30-h)	52	30
XLS-R (1-h)	42	29
XLS-R (5-min)	36	17
<u>Multilingual AWE:</u>		
CAE-RNN ab+cs+eu+sv+ta	45	2

(Table 6.1) are manually reviewed. The results are given in Table 6.4.

The table reports precision: the proportion of retrieved top-100 utterances that correctly contain a hate speech keyword. Surprisingly, we see that the multilingual AWE KWS system achieves a precision of 45%—better than the 5-minute and even the 1-hour ASR system in this in-the-wild test. This is in contrast to the in-domain KWS results, where the 1-hour ASR model outperformed the AWE KWS system (Table 6.2). Further investigation is required to understand exactly why the relative performance of the ASR and AWE systems are affected differently when applied to out-of-domain data. However, it is worth noting that several studies have shown that ASR system performance can drop dramatically when it is applied to data outside of its training domain [123–125]. One example where this can be seen in this case is how many of the ASR retrievals contain music (which we asked the human annotator to mark).

The in-the-wild search collection has been diarised automatically, which included a step to remove music segments. But this preprocessing step is not perfect: the search collection still ends up with some music (which neither the ASR nor AWE systems have seen in training). Table 6.4 shows that, out of the top-100 utterances for the 30-hour ASR system, 30% contained music. This decreases for the 1-hour (29%) and 5-minute (17%) ASR systems. But the AWE KWS system only retrieves two utterances containing music. And of these two, one utterance actually did contain a hate speech keyword in the music lyrics (there are also examples of such correct matches in the ASR music retrievals). Nevertheless, this shows that the AWE approach seems to be more robust to domain mismatch compared to training an ASR system on one hour or five minutes of labelled data. We therefore conclude that implementing an AWE-based KWS system as an ASR-free alternative for rapid deployment in an unseen language proves to be fruitful.

6.7. FURTHER ANALYSIS

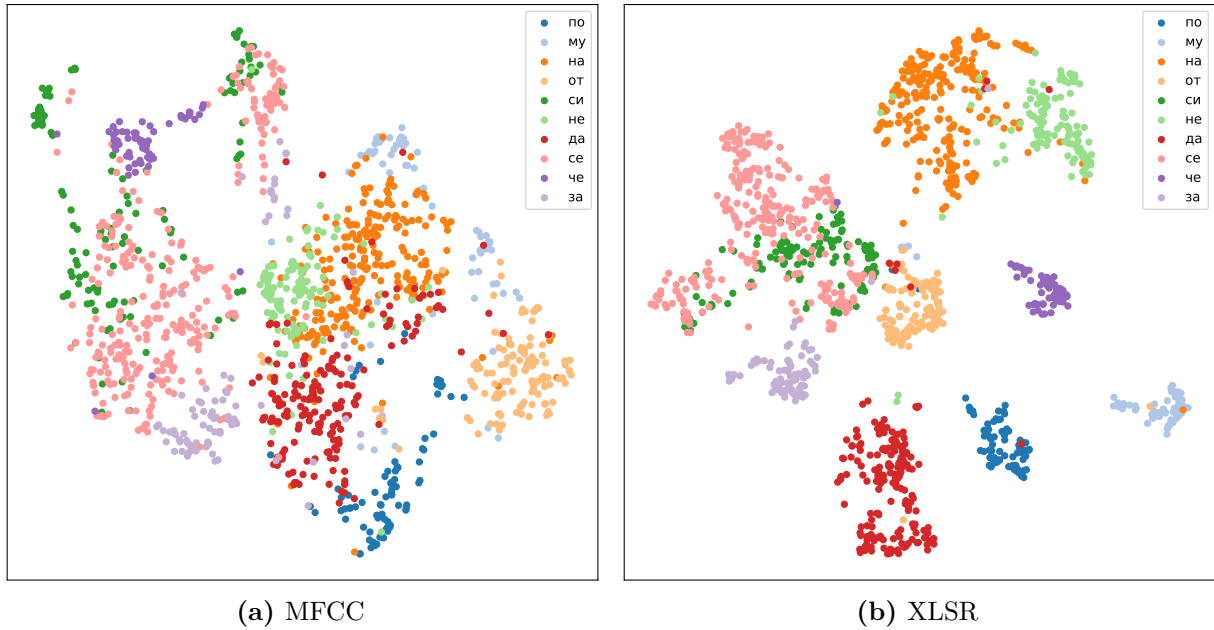


Figure 6.3: UMAP [126] visualisation of the ten most frequent two-character word segments in Bulgarian development data. These plots show the difference between AWEs produced by a multilingual AWE model trained on (a) MFCC features and (b) XLSR features.

6.7 FURTHER ANALYSIS

We briefly examine the transition from MFCCs to XLSR features as input to our multilingual AWE model. We also investigate the potential of selecting a more optimal threshold value for the AWE-based KWS system (Section 6.4.2).

6.7.1 XLSR vs MFCCs

For our experiments in this chapter we substituted MFCCs for XLSR speech features as input to our AWE models. Here we briefly look at the difference in embedding quality produced between a multilingual AWE model trained on XLSR features compared to MFCCs. We use the same training setup in Section 6.5.2.2 and train a multilingual AWE using MFCCs. For evaluation, we use the same developmental language, Bulgarian, and perform the same-different word task. On the same set of evaluation pairs, we obtain an AP of 29% and 58% for the MFCC and XLSR multilingual AWE models, respectively. With further investigation, we found that a model trained on XLSR features is much better at learning higher-quality embeddings for shorter word segments. Figure 6.3 shows the difference between the quality of embeddings produced for two character words when using (a) MFCCs and (b) XLSR as input features. Overall, we see word segments from the same class to be much better clustered together and separated from word segments from different classes. Using XLSR features as input clearly improves the quality of AWEs, leading to improved performance in the AWE-based KWS approach.

6.7. FURTHER ANALYSIS




Figure 6.4: Precision-recall curves from in-domain AWE-based KWS on in-domain Swahili test data using the same 36 keywords as in Section 6.6.1.

6.7.2 THRESHOLD

As mentioned in Section 6.4.2, the threshold value selected for the AWE-based KWS system was tuned on the F_1 -score by varying the threshold value over all keywords. For interest, we tune the threshold value per keyword and plot a precision-recall curve for each of the 36 keywords, shown in Figure 6.4. We also plot the precision-recall curve (in black) using the same threshold value for all keywords. From the figure we see that most per-keyword scores are higher than the average (blue lines above black line). When we average the F_1 -scores using the optimal threshold value per keyword, we obtain an F_1 score of 69% compared to 57% (bottom of Table 6.2) when using the same threshold for all keywords. This raises the question of the selected threshold value and whether simply treating keywords as equal might be a suboptimal approach. Practically, however, a per-keyword threshold can not be calculated given that for a new arbitrary keyword in a new unseen language we do not have development data to tune this.

But there is a possibility that a threshold can be selected based on other properties of a new keyword. These factors include the length of the speech segment, the number of phones present, and the relative distances between other unknown word segments in the acoustic space. For example, we may find that, in general, the threshold value for shorter word segments is lower than for longer speech segments. Future work should investigate the correlation between the threshold value of keywords and the potential factors mentioned above.

6.8 CHAPTER SUMMARY

This chapter addressed the problem of keyword spotting (KWS) for the purpose of hate speech detection in low-resource languages. We compared two KWS systems for data from Swahili and Wolof: a fine-tuned automatic speech recognition (ASR) model using different amounts of training data, and an ASR-free KWS system that utilises multilingual AWEs. The results show that fine-tuning a pre-trained multilingual ASR model using even a small amount of labelled data can outperform an AWE-based KWS system in controlled environments where training and test data come from the same domain. However, the AWE-based KWS system is more robust on out-of-domain radio broadcast data and achieves comparable results to an ASR model fine-tuned on 30 hours of labelled data. In the end, the merits of ASR vs AWE KWS will come down to the practical setting: it will depend on whether labelled training data can be collected from the target domain or not, and whether precision or recall is more important.

CHAPTER 7

LEVERAGING MULTILINGUAL TRANSFER FOR UNSUPERVISED SEMANTIC ACOUSTIC WORD EMBEDDINGS

In the previous chapters, we focussed on producing vector representations for spoken word segments where different acoustic realisations from the same word type should have similar embeddings. This can clearly be seen in the way we evaluated the embeddings: we use a word discrimination task (Section 2.5), a QbE speech search task (Section 5.3), and a keyword spotting task (Section 6.4.2), all of which require exact matching of spoken word segments. The performance of both these tasks relies on the discriminative ability between embeddings from different word types. Specifically, we aimed at producing these embeddings in low-resource settings, where annotated data in the target language is not available; here a multilingual transfer approach demonstrated superior performance.

In this chapter, we shift our focus to a different type of representation learning: *semantic acoustic word embeddings*. In a semantic-acoustic space, embeddings should not only be similar for different realisations of the same word type. Additionally, embeddings from distinct word types should also be close to each other if the respective words share meaning. These embeddings, mapping to spoken words, should reflect the same semantic properties as word embeddings trained on text data (Figure 1.1a). In other words, proximity in this new acoustic-semantic space should reflect word-type similarity as well as semantic relatedness between embeddings from different word types.

Again, we consider a low-resource scenario where word labels in the target language are unavailable. In this setting, there have been limited efforts to address the challenging problem of producing semantic AWEs. As a contribution, we propose a novel semantic AWE modelling approach: leveraging a pre-trained multilingual AWE model to assist semantic AWE modelling. This marks the first instance of using a multilingual transfer strategy in this context. Furthermore, existing research on semantic AWEs lacks comprehensive evaluation. We therefore perform our experiments in a constrained setup and evaluate our

7.1. BACKGROUND

embeddings in an intrinsic word similarity task (measuring semantic relatedness between words), and an extrinsic semantic QbE search. In the latter, a system should retrieve utterances containing different realisations of the query as well as utterances containing word segments that are semantically related to the query.

Parts of this chapter are presented in the submitted journal paper:

C. Jacobs and H. Kamper, “Leveraging multilingual transfer for unsupervised semantic acoustic word embeddings,” submitted to *IEEE Signal Processing Letters*, 2023.

7.1 BACKGROUND

We remind the reader of the properties and differences between AWEs and textual word embeddings described in Section 1.2.1. In the zero-resource setting, unsupervised representation learning of whole word speech segments primarily focuses on producing embeddings that allow discrimination between word segments of different word types. In this context, we train our AWE models using isolated word pairs. For example, the CAE-RNN (Section 3.2.1) is trained to reconstruct a different instance of the same word type as the input, and the CONTRASTIVERNN (Section 3.3) is explicitly trained to minimise the distance between instances of the same word type and maximise the distance between all instances of different word types. Consequently, different instances of the same word type should have similar embeddings and will be different from instances from different word types. These representations are learned without considering the context in which they appear within a spoken corpus. From this point onwards, we refer to these representations, which only encapsulate the acoustic properties linked to word type, as phonetic AWEs. An example vector space containing these phonetic AWEs is shown in Figure 7.1a. Here, different spoken instances of the same word type, for example “bank”, end up close to each other. Different spoken instances of the word “tank” will also end up close to each other. Since “bank” and “tank” are phonetically similar, instances of both word types might also end up close to each other in vector space. These are the type of embeddings generated in all our AWE experiments from previous chapters.

In contrast, in the text domain, word embedding models map written words to fixed vectors, capturing the contextual meaning of words within a text corpus. These models learn from the co-occurrence information in large unlabelled text corpora. For example, the Skip-gram model of Mikolov et al. [127] employs a neural network model that, given a target word, aims to predict the neighbouring words in that sentence (more in Section 7.2). As a result, words that are related in meaning end up having similar embeddings. Moreover, these textual word embedding (TWE) models assign a unique embedding to each word type within a fixed vocabulary. Figure 7.1b illustrates an example semantic space, where the written words “finance”, “bank”, and “currency” are mapped nearby, and “tank” and

7.1. BACKGROUND

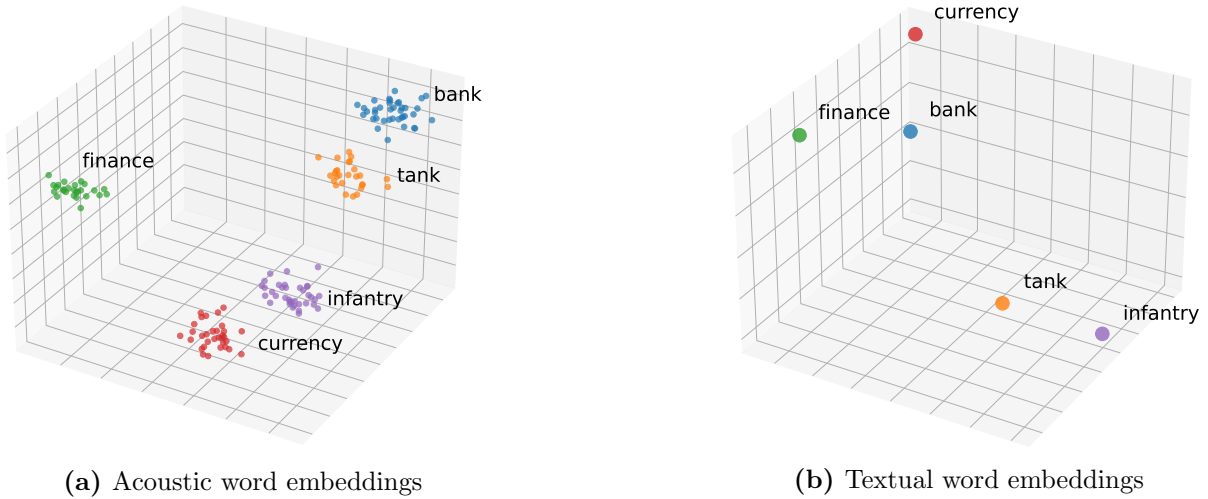


Figure 7.1: A visualisation of the fundamental difference between (a) AWEs and (b) TWEs. (a) Preserves word-type similarity among different acoustic realisations of five different word types. (b) Preserves semantic relatedness among the same word types mapping to written words.

“infantry”, reflecting their shared contextual usage.

Instead of written words, our goal in this chapter is now to map spoken word segments to vector representations that preserve this contextual relationship among words described by Mikolov et al. [127] (Figure 7.1b). A unique aspect of our problem lies in the continuous nature of speech, where each spoken word segment will be different from another. Consequently, each spoken segment will have a unique vector representation, which is different to TWEs where a single vector is assigned to each word type.

For instance, we would now like to see spoken instances of “finance”, “currency”, and “bank” mapped close to each other and, “tank” and “infantry”, to be in close proximity in the resulting acoustic-semantic space, as illustrated in Figure 7.2.

As part of our approach to create semantic AWEs, we will extend the textual word embedding approach called skip-gram. This approach has also been used as the basis for




Figure 7.2: A visualisation of a semantic acoustic word embedding space.

7.2. SKIP-GRAM

Table 7.1: Training samples extracted from a sentence using a context size of two. Example sentence: “two boys are kicking a ball to each other in the park”.

Context size = 2	Target word	Context words
[two boys are]	two	boys, are
[two boys are kicking]	boys	two, are, kicking
[two boys are kicking a]	are	two, boys, kicking, a
[boys are kicking a ball]	kicking	boys, are, a, ball
...
[each other in the park]	in	each, other, the, park
[other in the park]	the	other, in, park
[in the park]	park	in, the

other existing semantic AWE approaches. We therefore present the skip-gram approach in detail below, before going into further details on previous work on semantic AWEs.

7.2 SKIP-GRAM

The Skip-gram model of Mikolov et al. [127] is trained on large amounts of text corpora. The objective of this model is to learn word representations that are useful for predicting neighbouring words in a sentence. Formally, given a sequence of training words, $w_1, w_2, w_3, \dots, w_T$, the model aims to maximise the average log probability [28],

$$\frac{1}{T} \sum_{t=1}^T \sum_{-c \leq j \leq c, j \neq 0} \log p(w_{t+j} | w_t), \quad (7.1)$$

where c is the number of words before and after the target centre word w_t . Consider an example of multiple pairs of target and context words as training samples for a single sentence in a text corpus, shown in Table 7.1. For the training sample with input (the target word) “kicking” the model will learn to predict the words “boys”, “are”, “a”, “ball”.

Figure 7.3 shows the Skip-gram model architecture. The model consists of two fully connected feedforward layers with weight parameters $\mathbf{W}_1 \in \mathbb{R}^{(V \times D)}$ between input and hidden layer, and $\mathbf{W}_2 \in \mathbb{R}^{(D \times V)}$ between hidden layer and output, where V is the vocabulary size (number of unique words in text corpus) and D is the dimension of the hidden layer. For each training sample, the target and context words are encoded into one-hot vectors where the non-zero index corresponds to a word’s index in the vocabulary set. Given a one-hot vector as input, the model produces a latent embedding of size D . The embedding is then given as input to the second layer to produce an output vector of size V . A softmax function is applied to the output vector to normalise the output to a probability distribution over V classes, where the value at each index shows the prediction of the model given the input word. The model then updates weight parameters \mathbf{W}_1 and

7.3. RELATED WORK

\mathbf{W}_2 to maximise the objective function in Equation 7.1. Note that the one-hot input vector selects the i -th row from weight parameter \mathbf{W}_1 to produce the latent embedding. The one-hot vector essentially acts as a lookup table for each word in the vocabulary. Therefore, after training, \mathbf{W}_1 represents the learned word embeddings where the i th row of size D represents the i th word in the vocabulary set.

Keep in mind that the meaning of the words itself is not important, but rather the context in which they appear with other words, such that the words can be replaced with any unique character sequence (i.e. label) as long as the sequence in which they appear are retained. This allows a model to be trained on a text corpus with unknown word labels. In one of our speech models, we make use of this property by inferring unknown word labels for spoken word segments.

The training strategy of the Skip-gram model, which involves predicting surrounding words from a target centre word, served as an inspiration to the speech community as a means to obtain semantic embeddings from unlabelled spoken word segments. In the following section, we explore the existing body of work addressing this problem.

7.3 RELATED WORK

Learning semantic AWEs from speech is extremely difficult due to channel variability, noise, and speaker-specific information that are not present in written text. Moreover, the presence of inherent phonetic interference poses a significant obstacle to semantic AWE modelling. For example, distinguishing between spoken instances of “steal” and “steel” is challenging, even with the context provided, since from a model’s perspective they would appear identical. Some studies, therefore, consider the use of another modality as a grounding signal, for example, images [91, 128, 129] or text labels [130] as a weak form of supervision.

Only a handful of studies have looked at learning semantic AWEs from unlabelled speech alone. Chung and Glass [63] were the first to present and gain widespread recognition




Figure 7.3: The Skip-gram model architecture.

7.3. RELATED WORK

for producing semantic AWEs without word labels. They introduced a model named `SPEECH2VEC`, borrowing the training methodology of the Skip-gram model, by training on pairs of context words. `SPEECH2VEC` uses an RNN decoder-encoder structure to accommodate variable-length speech segments as input and target output (similar to the CAE-RNN), while the Skip-gram model uses a two-layer fully connected feedforward network with one-hot encoded vectors as input and output.

`SPEECH2VEC` is evaluated as follows: After training, all word segments from the training data are embedded using the encoder RNN. Each word type is then represented by a single embedding by averaging all instances of that type. The averaged embeddings are then used in multiple intrinsic word similarity tasks measuring semantic relatedness. With this setup, `SPEECH2VEC` achieved surprisingly good results, even outperforming textual word embeddings (trained on the corresponding transcripts) in some tasks. This certainly raised our suspicion, as well as others [131], concerning the validity of the embeddings produced by this model. In Section 7.4.1 we describe the model in more detail and address some of the concerns with `SPEECH2VEC`. Nevertheless, `SPEECH2VEC` has not been evaluated for single sample performance nor word segments that were not included in the model’s training data. The quality of embeddings produced for unseen segments is essential to determine the usefulness of embeddings for downstream applications, such as semantic QbE.

Chen et al. [64] proposed a two-step framework towards semantic AWE modelling. Speech segments are first given to a speaker disentanglement network (trained using the unlabelled word segments), which outputs two embeddings, separating speaker and phonetic information. The idea is that the phonetic embedding should only contain information related to word type (the same property of our multilingual AWEs). To produce semantic embeddings, they construct context training pairs using the phonetic embeddings and train a network optimising a contrastive loss similar to the `CONTRASTIVERNN` (Equation 3.6). Their results are not compared to `SPEECH2VEC` and also use averaged embeddings from the training data for evaluation.

`SPEECH2VEC` [63] and the work of Chen et al. [64] only consider the unlabelled speech from the target language during model development. In this chapter, we propose strategies for semantic AWE modelling, utilising a pre-trained multilingual AWE model. We evaluate our models on an intrinsic word similarity task, measuring the quality of embeddings for single word segments not used during training. To demonstrate the practical applicability of our embeddings, we also apply them to a downstream semantic QbE task.

In both cases [63, 64], the problem is simplified by assuming that we know where words start and end in the target language (but the word classes are still unknown). However, in a true zero-resource setting, word boundaries will be unknown. In the experiments presented in this chapter, we also assume word boundaries are known. This problem needs to be addressed in the future to consider complete unsupervised semantic AWE modelling

(Section 8.2).

7.4 SEMANTIC AWEs TRAINED FROM SCRATCH

Until now, our AWE models mapped speech segments from the same word type nearby in vector space (Figure 7.1a). In our notation, a speech segment $X = (\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_T)$ is projected to a fixed vector \mathbf{z} . From this point onwards we add the subscript p to the AWEs from previous chapters. In other words, \mathbf{z}_p indicates that the embedding preserves phonetic information related to word type only.

The goal of this chapter is to produce semantic embeddings \mathbf{z}_s . These embeddings should not only reflect phonetic similarity but also capture word meaning (Figure 7.2). SPEECH2VEC [63] and the work of Chen et al. [64] achieve this by only training on context pairs extracted from unlabelled speech data (from the target language). We use the training strategies presented in these two methods as a baseline for our multilingual enhanced semantic AWE models (Section 7.5).

7.4.1 SPEECH2VEC

SPEECH2VEC is a variant of the CAE-RNN AWE model (Section 3.2.1). The CAE-RNN is an encoder-decoder model that takes a speech segment as input and aims to reconstruct a different instance of the same word type. SPEECH2VEC uses the same encoder-decoder framework, but instead of training on pairs of instances from the same word type (as for phonetic AWE modelling), it uses context word pairs $(X_{\text{trg}}, X_{\text{ctx}})$. X_{trg} is a target centre word segment while X_{ctx} is a context word appearing somewhere in a window around the centre. These context pairs are constructed similarly to those used for training the Skip-gram model on text data (see Table 7.1). However, in this case, X_{trg} and X_{ctx} represent spoken words. Importantly, these context pairs are constructed without word labels by only considering the relative position of words within an utterance (requiring word boundaries). Figure 7.4 shows the SPEECH2VEC model for a single training pair. SPEECH2VEC optimises the reconstruction loss of Equation 3.1, where X' in that equation is replaced with X_{ctx} . Ideally the resulting embeddings \mathbf{z}_s should therefore be similar for words that co-occur in the speech data. Although this approach seems theoretically sound, concerns have been raised regarding the performance of this model. We address this now.

Criticism. The authors of SPEECH2VEC [63] did not make their code available, despite multiple requests from the public. They did however release a set of embeddings produced by the model.¹ With numerous attempts to replicate the model’s embeddings exactly as described in the paper, and reaching out to the authors without receiving a response, we have not been able to achieve results close to theirs. We are also not aware of any credible

¹<https://github.com/iamyuanchung/speech2vec-pretrained-vectors>

7.4. SEMANTIC AWEs TRAINED FROM SCRATCH

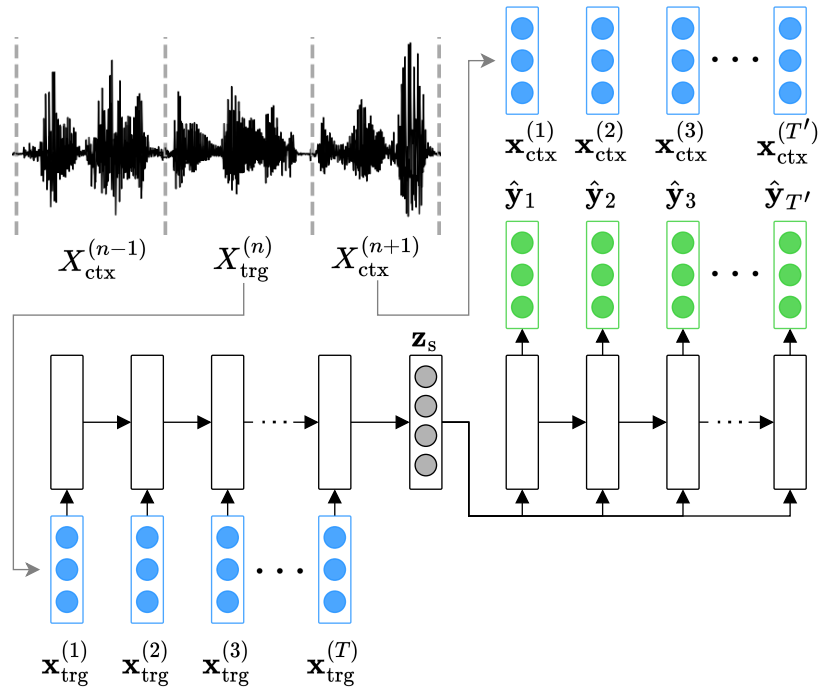


Figure 7.4: The SPEECH2VEC model architecture. This is similar to the CAE-RNN model architecture (Figure 3.1) but adapted for context word pairs $(X_{\text{trg}}, X_{\text{ctx}})$ instead of positive pairs (X, X') . Dashed lines in the audio indicate true word boundaries.

publications that successfully replicated SPEECH2VEC to obtain similar quality semantic embeddings on the same evaluation tasks. Other failed attempts to replicate the results of SPEECH2VEC include the works of Jang², Wang³ and EarthSpecies⁴.

Chen and Cao [131] recently performed an in-depth analysis questioning the authenticity of the original SPEECH2VEC model. On a homophone inspection analysis, they show that SPEECH2VEC is immune to phonetic interference. For example, the model will be able to perfectly discriminate between spoken instances of “steal” and “steel” despite being phonetically identical, which should inevitably affect model performance. By inspecting the vocabulary of the released embeddings, they found it exactly matches the vocabulary of a word embedding model trained on the corresponding transcripts. They show this should not be true given that some information (words) are lost when applying a forced-aligner to segment the audio. Based on these findings, there is compelling evidence that the embeddings released by SPEECH2VEC have been influenced by word labels in some way.

Moreover, the CAE-RNN AWE was originally designed to learn embeddings that are invariant to speaker identity, channel characteristics, etc. (Section 3.2.1). In the SPEECH2VEC setup, training pairs are from the same speaker and utterance (the way context pairs are constructed), so that the model is not explicitly able to learn embeddings invariant to those properties. Despite the scepticism of SPEECH2VEC we include the model

²<https://github.com/yjang43/Speech2Vec>

³<https://github.com/ZhanpengWang96/pytorch-speech2vec>

⁴<https://github.com/earthspecies/audio-embeddings/issues/6>

7.5. OUR APPROACH: SEMANTIC AWEs USING MULTILINGUAL TRANSFER

as a baseline in our experiments. We use our own version of the model and perform experiments in a different setup (but still refer to it as SPEECH2VEC).

7.4.2 SEMANTIC CONTRASTIVERNN

Similar to how the CAE-RNN AWE model is adapted for semantic AWE modelling (SPEECH2VEC), we modify the CONTRASTIVERNN AWE model (Section 3.3) to learn semantic AWEs. We substitute X_a , X_p and X_{n_1}, \dots, X_{n_K} in Equation 3.6, with X_{trg} (target centre), X_{ctx} (context), and out-of-context words. Now, by minimising the contrastive loss of Equation 3.6, the model should optimise relative distances between embeddings such that it reflect semantic similarity. This is similar to the contrastive optimisation for semantic AWEs by Chen et al. [64], without speaker disentanglement (requiring speaker labels).

For both SPEECH2VEC and semantic CONTRASTIVERNN, a semantic AWE model is trained from scratch, therefore requiring the models to learn phonetic and semantic similarity simultaneously. We now look at new semantic AWE modelling strategies that we propose, utilising a pre-trained multilingual AWE model.

7.5 OUR APPROACH: SEMANTIC AWEs USING MULTILINGUAL TRANSFER

This section introduces a novel semantic AWE modelling approach. We propose incorporating phonetic word knowledge through a pre-trained multilingual AWE model during semantic modelling. Incorporating phonetic details may assist the model to prioritise learning word semantics by filtering out unrelated characteristics, for example, speaker information. Specifically, we propose three strategies utilising a pre-trained multilingual AWE model in various ways.

7.5.1 CONTRASTIVERNN WITH MULTILINGUAL INITIALISATION

Instead of training semantic models from scratch (Section 7.4) we warm-start them using the learned weights of a pre-trained multilingual AWE model. In our experiments, we use the learned weights of a multilingual AWE model’s encoder to initialise the encoder RNN of the CONTRASTIVERNN and subsequently train a semantic CONTRASTIVERNN (Section 7.4.2). This learned parameter initialisation method is similar to the multilingual adaptation strategy explored in Chapter 4: a pre-trained multilingual AWE model is adapted to a target low-resource language using discovered word pairs. Here, we adapt the multilingual weights to produce semantic AWEs by training on context pairs from the target language.

7.5. OUR APPROACH: SEMANTIC AWEs USING MULTILINGUAL TRANSFER

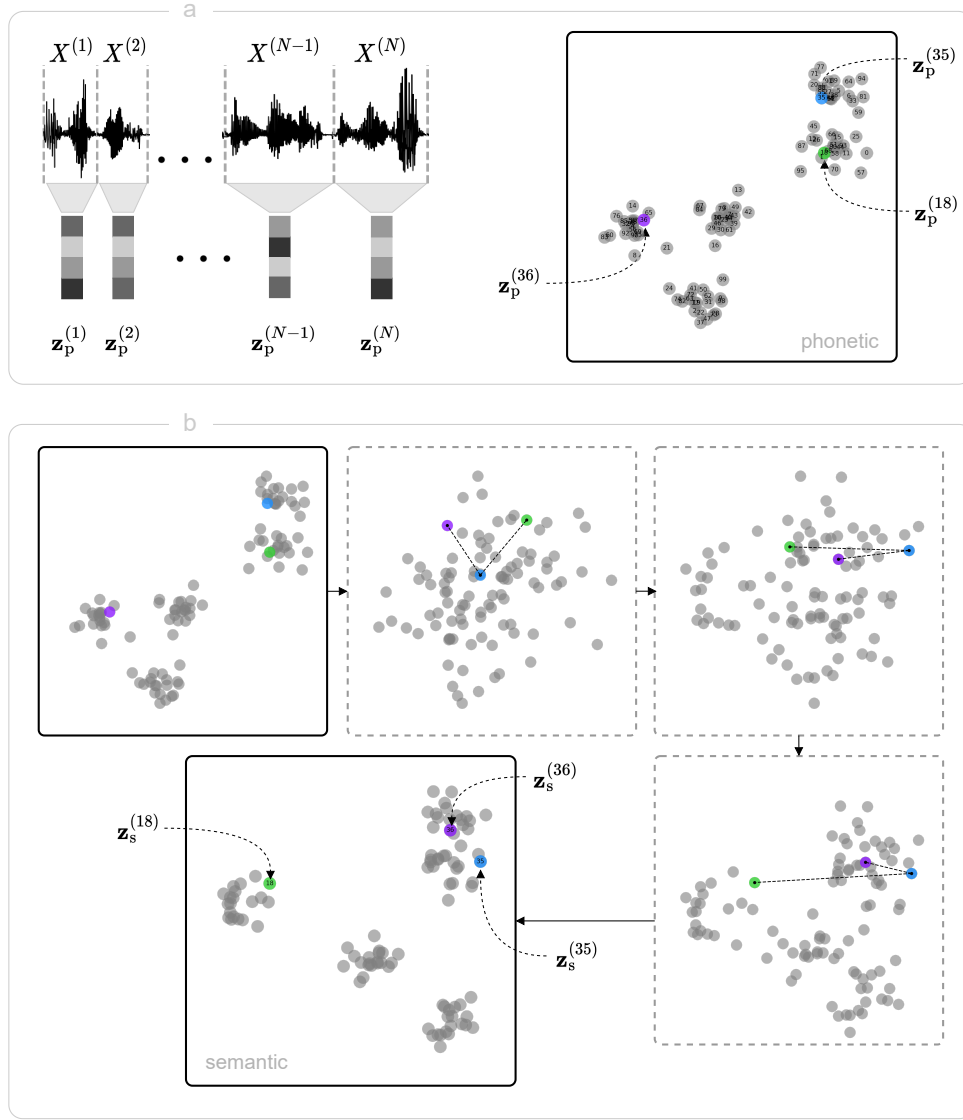


Figure 7.5: An illustration of the PROJECT AWE semantic strategy, as detailed in Section 7.5.2.

7.5.2 PROJECTING MULTILINGUAL AWEs

An alternative approach would be to first apply the multilingual AWE model to the segmented speech segments $\{X^{(n)}\}_{n=1}^N$ where N is the number of word tokens in the training corpus. Then, train a projection network, mapping the resulting phonetic AWEs to a new semantic AWE space. Figure 7.5 illustrates this mapping process. Applying the multilingual model to the unlabelled speech segments $\{X^{(n)}\}$ produces a set of phonetic AWEs $\{z_p^{(n)}\}$ (upper part in figure). In the phonetic space, we note that word segments are mapped to unknown word classes (five in this simplified example). This implies $z_p^{(35)}$ (blue) and $z_p^{(71)}$ most likely belong to the same unknown word class, as well as, $z_p^{(18)}$ (green) and $z_p^{(87)}$. Groups containing $z_p^{(35)}$ and $z_p^{(18)}$ are in close proximity, indicating they might also share phonetic similarities.

In the semantic space, we would like to see embeddings of segment $X^{(35)}$ and $X^{(36)}$ to

7.5. OUR APPROACH: SEMANTIC AWEs USING MULTILINGUAL TRANSFER

be positioned close to each other (assuming the unknown classes associated with these segments frequently co-occur). We do this by training a projection network, mapping the phonetic AWEs to a new set of semantic embeddings $\{\mathbf{z}_s^{(n)}\}$. We train this network using context pairs and the contrastive loss of Equation 3.6, optimising the distances between the output embeddings \mathbf{z}_s . The lower part of Figure 7.5 illustrates the transformation of the phonetic space into a new semantic space, through optimising the contrastive loss. We now observe neighbouring word segments $X^{(35)}$ and $X^{(36)}$ end up having similar semantic representations $\mathbf{z}_s^{(35)}$ and $\mathbf{z}_s^{(36)}$.

This approach is similar to the semantic AWE method presented by Chen et al. [64]. Here they used data from the target language to train a network separating speaker and phonetic information (which requires speaker labels). Instead, we use a pre-trained multilingual AWE model designed for generating phonetic representations for unlabelled word segments.

7.5.3 CLUSTER+SKIP-GRAM

This new approach is based on the Skip-gram model (Section 7.2). The Skip-gram model uses a fixed dictionary of discrete word class labels to create input and output vectors for training on text data. In the low-resource setting we do not have word labels for speech segments.

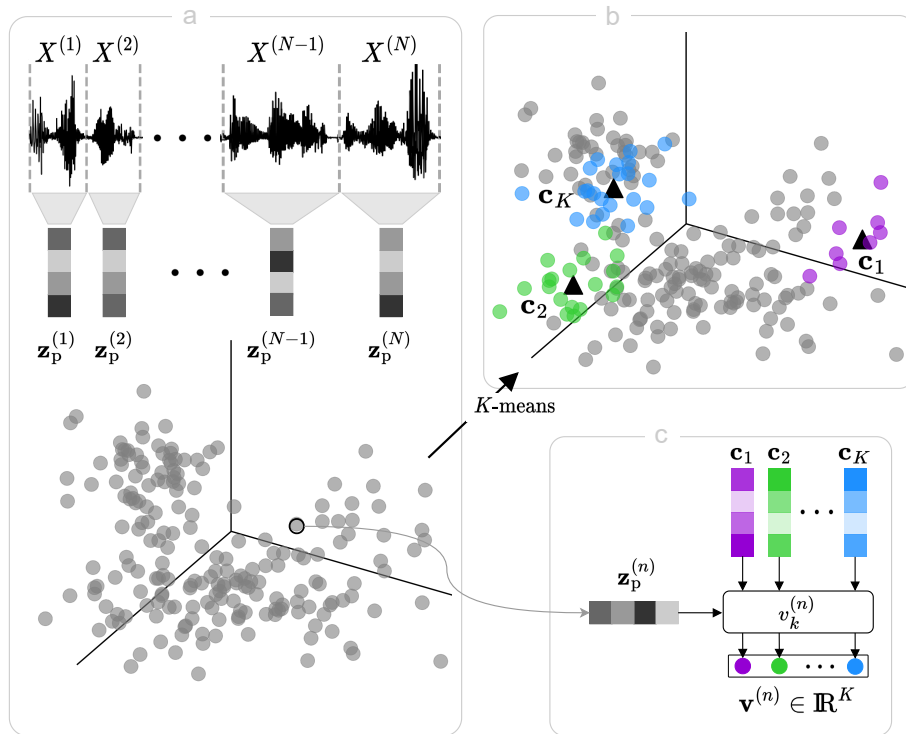


Figure 7.6: Our CLUSTER+SKIP-GRAM semantic AWE approach. Speech segments $X^{(n)}$ are represented by soft pseudo-word label vectors $\mathbf{v}^{(n)}$ which are then used to train a Skip-gram model.

7.5. OUR APPROACH: SEMANTIC AWEs USING MULTILINGUAL TRANSFER

Instead, we use the phonetic similarities in the original AWE space to derive a soft pseudo-word label for each speech segment. This process is illustrated in Figure 7.6. First, a pre-trained multilingual AWE model is applied to the segmented speech corpus $\{X^{(n)}\}$, producing a set of phonetic AWEs $\{z_p^{(n)}\}$. The colourless embedding space represents the unsupervised setting where embeddings have no associated word types or labels. We then apply K -means clustering to the phonetic embedding space, producing a set of centroids $\{c_k\}_{k=1}^K$, with K representing the predetermined number of clusters for partitioning the vector space.

Ideally, these clusters should resemble distinct word classes. For example, all spoken instances of “apple” should be assigned to cluster one (purple) and all instances of “water” should be assigned to cluster two (green). Note, with a perfect clustering, where all embeddings from an unknown word label are exclusively assigned to a unique cluster, training a Skip-gram model using the cluster labels would be equivalent to training on the corresponding text labels (see bottom of Section 7.2). However, due to clustering inaccuracies, not all embeddings belonging to the same unknown word class will be assigned to the same cluster. The main challenge arises from embeddings of the same word class being assigned to multiple clusters; a consequence of the uncertainty regarding vocabulary size and word class frequency range. For example, embeddings assigned to cluster 2 (green) and cluster K (blue) in Figure 7.7 might represent the same word class given that there is no clear separation between them. We experimented with cluster labels as training labels for Skip-gram modelling, but this gave very poor performance on development data.




Figure 7.7: Given an acoustic embedding $z_p^{(n)}$ (assigned to cluster 2), instead of using a one-hot encoded input $\mathbf{x} = \{0, 1, 0, \dots, 0\}$, the value at each index is set using Equation 7.2, for example, $\mathbf{v} = \{0.03, 0.45, 0.01, \dots, 0.43\}$. Note that the values at index 2 and K are the highest given that $z_p^{(n)}$ is the closest to cluster means c_2 and c_K . Black triangles indicate cluster means.

7.5. OUR APPROACH: SEMANTIC AWEs USING MULTILINGUAL TRANSFER

Rather than assigning a hard label to each segment based on its associated cluster number, we calculate a soft vector label of an AWE belonging to each cluster:

$$v_k^{(n)} = \frac{\exp\left(-\text{sim}(\mathbf{z}_p^{(n)}, \mathbf{c}_k)/\sigma^2\right)}{\sum_{j=1}^K \exp\left(-\text{sim}(\mathbf{z}_p^{(n)}, \mathbf{c}_j)/\sigma^2\right)} \quad (7.2)$$

where $\text{sim}(\cdot)$ denotes cosine similarity and σ is a hyperparameter controlling the influence of distant centroids. This non-linear scaling approach assigns more weight to embeddings closer to the target embedding $\mathbf{z}_p^{(n)}$. Figure 7.7 extends the soft labelling calculation illustration in the bottom right in Figure 7.6. Now, each word segment $X^{(n)}$ is associated with a new unique vector $\mathbf{v}^{(n)} \in \mathbb{R}^K$, where the value at index k indicates the likelihood of the segment being contained in cluster k . Even if segments from the same unknown word class are assigned to multiple clusters, their relative position in the embedding space is captured in \mathbf{v} .

Lastly, we use these soft vectors as inputs and targets to train a Skip-gram model, illustrated in Figure 7.8. Compared to the Skip-gram model trained on text data (Figure 7.3), there is a key difference: using these soft representations as input, compared to a one-hot vector. This means that \mathbf{W}_1 do not act as a look-up table, instead, each row contributes to the hidden embedding with corresponding weight (visually illustrated with shading in Figure 7.8). In other words, the hidden layer contains the weighted sum of all rows in \mathbf{W}_1 , which we take as the semantic embedding \mathbf{z}_s . The same applies to \mathbf{W}_2 when using soft targets.

During inference, a unique embedding \mathbf{z}_s is produced for a word segment by first obtaining a phonetic AWE using the multilingual model, then calculating representation \mathbf{v} using the trained cluster means, and finally applying it to the learned weight matrix \mathbf{W}_1 .




Figure 7.8: The skip-gram model using soft inputs and targets.

7.6 SEMANTIC QUERY-BY-EXAMPLE

In Chapter 5 and Chapter 6, we performed QbE experiments using AWEs to calculate the phonetic similarity between speech segments. This involved searching an unlabelled audio corpus with a query word segment, aiming to retrieve utterances containing instances of the query. We refer to this as exact QbE, where the query only retrieves utterances if it contains a word sharing the same class label. In this chapter, apart from intrinsic AWE evaluations, we consider the task of *semantic* QbE as a downstream application for semantic AWEs. In semantic QbE the goal is to retrieve utterances from an audio collection that not only match segments containing the exact query segment but also utterances which contain word segments that are semantically similar to the query. For example, using a query segment, “football”, the system should retrieve utterances containing exact query matches for instance, “the football player is injured”, as well as utterances containing word segments semantically related to the query for instance, “two kids play soccer in the field”, and “the player was offside”.

We perform our semantic QbE experiments using a similar setup as in [132]. They performed experiments to produce semantic embeddings on the same data corpus we use for our semantic AWE experiments. However, during training, they use images as a weak form of supervision to produce semantic embeddings. To the best of our knowledge, we are the first to perform semantic QbE using semantic AWEs trained on unlabelled speech segments.

7.7 SETUP

In this section we outline the data corpus we use for our semantic AWE modelling. We provide implementation details for the two semantic models, SPEECH2VEC and CONTRASTIVERNN, trained from scratch (Section 7.4). Subsequently, we present the specifics of the three semantic AWE models leveraging a pre-trained multilingual AWE model we introduced in Section 7.5. Finally, we discuss the evaluation process for our semantic AWEs.

7.7.1 DATA

For this experiment, we use the Flickr8k Audio Captions Corpus (FACC) [128]. The corpus contains 40k spoken captions in English describing the content of a Flickr image [133]. These images depict scenes for which humans have given verbal descriptions. Example images with corresponding spoken captions are shown in Figure 7.9. Note, the text in Figure 7.9 is only for the purpose of illustration and not used during training (the corresponding speech are used). The corpus is divided into train, development, and test

7.7. SETUP

splits containing 30 000, 5 000, 5 000 utterances respectively.

Although the amount of training data in the Flickr8k corpus is relatively low, the content of the data suits our semantic experiments well: images come from a fairly narrow domain, with five different spoken captions describing the content of each image. Therefore, semantic concepts reoccur in different utterances. This dataset has also been used for acoustic semantic modelling by others [91, 128, 132].

We train and evaluate our models using two sets of speech features. Firstly, we employ 13-dimensional MFCCs. Secondly, as in Chapter 6, we conduct experiments with self-supervised speech features: we use the 12th transformer layer of the multilingual XLSR model [119] to get 1024-dimensional features. Utterances are normalised per speaker and segmented using true word boundaries from forced alignments [121]. We use these word segments from the training data to construct context word pairs. For all the semantic models, we use a context window of three words before and after a centre word: this yields roughly two million context pairs in our training data.

7.7.2 SEMANTIC MODELS TRAINED FROM SCRATCH

The encoder and decoder of our SPEECH2VEC (Section 7.4.1) implementation each consists of three unidirectional RNNs with 400-dimensional hidden vectors and an embedding size of 100. The semantic CONTRASTIVERNN (Section 7.4.2) uses the same encoder structure. We train both these models using the context pairs extracted from the training set. For the CONTRASTIVERNN we randomly sample word segments outside the context window as



A golden-colored dog, with his eyes alert, holds a brightly colored tennis ball in his mouth.
A brown and white dog holds a tennis ball in his mouth.
A dog has a tennis ball in its mouth.
A tan dog is playing with the green ball.
Dog running with a tennis ball in its mouth.



Boys kicking soccer ball in the grass under a tree.
Two boys are kicking a ball to each other in the park.
Two boys kick around a ball in a meadow.
Two kids play soccer in a field.
Two boys have made a goal out of two jackets in order to play soccer.



A man in a red jacket looking out over snow-covered mountains.
A man in red snow gear looks at an expanse of snow and mountains.
A person in a red and black ski outfit is looking over a snowy, mountainous landscape.
A person in a red coat looking out at a snowy landscape.
View from the summit.

Figure 7.9: Examples of images with corresponding spoken captions from the FACC.

negatives; for each positive, we sample 20 negatives.

7.7.3 SEMANTIC MODELS USING MULTILINGUAL TRANSFER

All the multilingual transfer strategies presented in Section 7.5 use a pre-trained multilingual AWE model. The multilingual model is trained in the same way multilingual models were trained in previous chapters, with embedding dimension set to 100. We train a multilingual CAE-RNN AWE model on labelled data from four different well-resourced languages included in Common Voice [118]: Italian, Dutch, Russian, and Czech. Note, the entire speech corpus is in English and therefore not included as a training language. We pool the data from all four languages and extract 300k training pairs.

Multilingual initialisation. (Section 7.5.1) For the multilingual weight initialisation method, we use the trained parameters of the multilingual CAE-RNN encoder to initialise the CONTRASTIVERNN’s encoder. Furthermore, we freeze the weights of the first two encoder layers to preserve the phonetic structure of input words during further training. This initialisation approach is similar to how we initialised AWE models for adaptation in Chapter 4: a multilingual AWE model is adapted using discovered word pairs from the target language. Here, we adapt for semantic AWEs by training on context pairs from the target corpus.

Project AWE. (Section 7.5.2) For our projection network, we use a feed-forward network of two linear layers with an inner dimension of 1024 and input and output dimensions of 100. The input dimension is set to match the embedding dimension produced by the multilingual AWE model. The inner and output dimensions are chosen based on experimenting on development data. Again, we train on context pairs and sample 20 negatives for each positive. The network optimises by minimising the contrastive loss of Equation 3.6.

Cluster+Skip-gram. (Section 7.5.3) For K -means clustering, we use $K = 5\,000$ clusters and set $\sigma = 0.01$ in Equation 7.2. The number of unique word types in the target corpus is 6 307 where 2 150 appears once. We experimented with different values for K ranging from 1 000 to 10 000 which produced similar results. We initialise our cluster means with a random selection of acoustic embeddings, perform 50 permutations, and keep the best means. We reimplement the Skip-gram model [127] described in Section 7.2, setting the latent embedding dimension to $D = 100$. The model is trained using the soft-vectors as input and optimises a negative log-likelihood with class probabilities set to the context target word’s soft-vector. The model is optimised using Adam [103], with a learning rate of 0.001.

7.7. SETUP

Table 7.2: Word-pair similarity ratings extracted from the WordSim353 [134] evaluation set. Similarity scores indicate word relatedness based on human judgements.

money	bank	8.50
physics	proton	8.12
tennis	racket	7.56
journey	car	5.85
smart	student	4.62
problem	airport	2.38
professor	cucumber	0.31

7.7.4 EVALUATION

We evaluate the quality of resulting semantic embeddings in an intrinsic word similarity task. To measure the applicability for real-life applications, we also apply the semantic AWEs in a downstream semantic QbE task for extrinsic evaluation.

7.7.4.1 INTRINSIC SEMANTIC AWE EVALUATION: WORD SIMILARITY

In the text domain, word embeddings are evaluated for semantics by measuring the similarity between word pairs through vector distance calculations. These vector distances are then compared to standard reference sets containing similarity ratings assigned by humans. As an illustration, Table 7.2 shows some reference word-pair similarities from the WordSim353 [134] evaluation set. For instance, in a resulting semantic space, the distance between the embedding of “tennis” and “racket” should be smaller than the distance between “journey” and “car”, but greater than the distance between “physics” and “proton.” Furthermore, “money” and “bank” should have the smallest distance than all the shown word-pair distances. Spearman’s ρ is used to quantify the similarity between the set of word-pair distances (WP_{embed}) and the set of word-pair human judgement ratings (WP_{human}), calculated as follows:

$$\rho = \frac{\text{cov}(R(WP_{\text{embed}}), R(WP_{\text{human}}))}{\sigma_{R(WP_{\text{embed}})}\sigma_{R(WP_{\text{human}})}} \quad (7.3)$$

with R converting the scores to ranks, and $\text{cov}(\cdot, \cdot)$ and σ denotes, covariance and standard deviation, respectively.

In our controlled experiment, our vocabulary is limited, causing many words from the standard evaluation sets typically used to assess word embeddings trained on large amounts of text to be absent from our training and evaluation data. To address this, we compile our own list of word pairs that frequently appear in our dataset. Instead of using human judgement, we assign similarity scores to those pairs based on the distance between textual word embeddings generated by an off-the-shelf Skip-gram model trained

7.7. SETUP

For the sentence below, select all words that could be used to search for the described scene.

Sentence: a skateboarder in a light green shirt.

☐ dogs ☐ beach ☐ children ☐ white ☐ swimming
☐ wearing ☐ skateboard ☐ None of the above

Figure 7.10: An example annotation job presented to a human for manual semantic labelling, with keyword options. In this case, “skateboard” was selected by all five annotators and “wearing” by four. Reproduced from [132].

on the corresponding transcribed utterances. Given that our goal is to obtain semantic embeddings similar to Word2Vec, we believe this approach is reasonable.

To obtain a single semantic embedding for each word class, we calculate the average of all embeddings from the same class and report ρ_{avg} . Given that we are particularly interested in obtaining semantic embeddings for individual word segments, single-sample performance is also measured by randomly selecting one instance of each spoken word. This is repeated ten times and averaged to get a single score ρ_{single} .

Table 7.3: The number of annotators (out of five) that selected the keyword for that utterance is shown. The majority vote indicates whether the keyword is deemed semantically relevant to an utterance, according to the labelling. Adapted from [132].

Keyword	Example utterance	Count	Majority vote
jumps	biker jumps off of ramp	5 / 5	Y
ocean	man falling off a blue surfboard in the ocean	5 / 5	Y
race	a red and white race car racing on a dirt racetrack	5 / 5	Y
snowy	a skier catches air over the snow	5 / 5	Y
bike	a dirt biker rides through some trees	4 / 5	Y
children	a group of young boys playing soccer	4 / 5	Y
young	a little girl in a swing laughs	4 / 5	Y
field	two white dogs running in the grass together	3 / 5	Y
carrying	small dog running in the grass with a toy in its mouth	2 / 5	Y
face	a man in a gray shirt climbs a large rock wall	2 / 5	N
large	a group of people on a zig path through the mountains	1 / 5	N
sitting	a baby eats and has food on his face	1 / 5	N
hair	two women and a man smile for the camera	0 / 5	N

7.7.4.2 SEMANTIC QbE

We use the same setup as [91] to evaluate downstream semantic QbE performance. Semantic labels for 1 000 test utterances from FACC were collected from human annotators, using a set 67 keyword classes [132]. Specifically, each of the 1 000 utterances was labelled by five annotators, indicating whether a particular keyword is semantically relevant to that utterance (regardless of whether the word instance appears verbatim in the utterance). An example annotation job is shown in Figure 7.10.

We use the majority decision to assign a hard label for whether a keyword is relevant to an utterance as in [132]. Table 7.3 shows annotations for some of the utterances. For instance, a correct retrieval using the keyword “jump” would be “bike jumps of ramp”, while an incorrect retrieval using the keyword “face” would be “a man in a gray shirt climbs a large rock wall”. Using these hard labels, we calculate semantic $P@10$, $P@N$, EER (metric details in Section 5.3).

Given the ambiguity among annotators when determining if a keyword is semantically related to an utterance, we also use the counts as soft labels and calculate Spearman’s ρ (Equation 7.3). Here, ρ measures the correlation between a system’s ranking and the actual number of annotators who deemed a query keyword relevant to an utterance.

To simplify the QbE task, we still assume that ground truth word boundaries are known: a query AWE is therefore compared to AWEs for the word segments in an unlabelled search utterance.

7.8 EXPERIMENTAL RESULTS

We first compare embedding quality through measuring word similarity between the semantic AWE models trained from scratch (Section 7.4) and the semantic AWE models we proposed leveraging multilingual AWEs (Section 7.5). We then apply one of our proposed semantic AWE models in a semantic QbE search for extrinsic evaluation. This is followed by a qualitative analysis of our semantic AWEs.

7.8.1 INTRINSIC EVALUATION: WORD SIMILARITY

Table 7.4 presents the intrinsic scores of embeddings from the semantic AWE models, trained either from scratch (top section) or using multilingual transfer (bottom). The benefit of multilingual transfer is evident in the scores of the projection and CLUSTER+SKIP-GRAM approaches, with the latter outperforming all other models regardless of the input features used or whether single or averaged embeddings are evaluated. The single-sample performance ρ_{single} is particularly significant as it shows that individual representations can be compared accurately—a useful property for downstream applications such as semantic QbE (Section 7.6).

7.8. EXPERIMENTAL RESULTS

Table 7.4: Spearman’s ρ (%) for semantic word similarity of AWEs on development data. Embeddings are either averaged over word segments of the same class, or a single segment is sampled per class. AWE models take either MFCC or XLSR features as input.

Model	MFCC		XLSR	
	ρ_{single}	ρ_{avg}	ρ_{single}	ρ_{avg}
<i>Trained from scratch:</i> (Section 7.4)				
SPEECH2VEC	1.0	5.3	7.2	21.3
CONTRASTIVERNN	0.7	5.1	4.6	25.9
<i>Using multilingual transfer:</i> (Section 7.5)				
CONTRASTIVERNN multilingual init.	0.6	5.5	6.1	24.5
PROJECT AWE	3.6	18.2	18.3	33.6
CLUSTER+SKIP-GRAM	18.0	35.6	35.9	41.7

The CONTRASTIVERNN is the one exception that does not show a clear gain from initialising with multilingual weights compared to training from scratch. As a sanity check, we evaluate the phonetic multilingual AWEs before semantic training, in other words, embeddings produced by the foundation multilingual AWE model (top of Section 7.7.3) used for transfer in the bottom section. We obtain results for $\rho_{\text{single}} = 0.59\%$ and $\rho_{\text{avg}} = -0.13\%$. As expected, this indicates that phonetic multilingual AWEs do not capture semantic information.

Again, we see the benefit of using self-supervised speech representations as input to AWEs instead of conventional features, as also in previous work [122, 135]. We therefore use XLSR features for all further experiments. We apply the semantic AWE model achieving the best results on the word similarity task, CLUSTER+SKIP-GRAM, in our semantic QbE search.

7.8.2 EXTRINSIC EVALUATION: SEMANTIC QbE

Table 7.5 compares the CLUSTER+SKIP-GRAM (semantic) and multilingual AWE (phonetic) models when used in a downstream QbE system. We evaluate both exact and semantic QbE, where the latter gets awarded for retrieving exact query matches as well as utterances labelled as semantically related to the search query. To situate results, we use a random baseline model that assigns a random relevance score to each utterance. (The relatively high scores of the random approach are due to the narrow domain of the evaluation data, Section 7.7.1.)

Looking at the EER and Spearman’s ρ for semantic QbE, we see that the CLUSTER+SKIP-GRAM model achieves the highest score, outperforming the purely phonetic AWEs from the multilingual AWE model. The reason why the phonetic multilingual AWE model outperforms the semantic model in $P@10$ and $P@N$ is due to its proficiency in detecting

7.8. EXPERIMENTAL RESULTS

Table 7.5: Exact and semantic QbE results (%) on test data.

Model	Exact QbE			Semantic QbE			
	$P@10$	$P@N$	EER	$P@10$	$P@N$	EER	ρ
<i>Baselines:</i>							
Random	5.0	5.0	50.0	9.1	9.1	50.0	5.7
Multilingual AWE	90.7	82.8	4.7	92.8	59.4	24.1	17.0
<i>Semantic AWE model:</i>							
CLUSTER+SKIP-GRAM	85.8	71.3	9.3	88.2	52.1	21.6	28.2

exact matches (which are also correct semantic matches).

In this chapter, we are mainly interested in observing information captured by our embeddings related to word meaning. However, in the semantic QbE task, a system also gets rewarded for retrieving utterances containing exact instances of the query. To get a better sense of the ability of a model to retrieve non-verbatim semantic matches, we mask out all exact occurrences of the query before we segment the search collection. The system will now only be able to retrieve utterances based on word meaning. The results are shown in Table 7.6. Now we see a clear benefit in using the CLUSTER+SKIP-GRAM model, with the phonetic multilingual AWE model becoming close to a random search given that there are no exact matches.

Our core goal is semantic QbE, but it is worth briefly touching on exact QbE performance of the CLUSTER+SKIP-GRAM model in Table 7.5. In exact QbE, the system should only retrieve utterances containing instances of the query word type. The CLUSTER+SKIP-GRAM model learns to predict frequently co-occurring words, without explicit training for preserving word-type information. Although only trained for semantics, the CLUSTER+SKIP-GRAM still achieves reasonable exact retrieval performance, with only a drop of between 5% and 10% in scores compared to the multilingual AWE model. It is therefore clear that this semantic model is able to retain phonetic properties for word-type while also capturing semantic information related to context.

Table 7.6: Semantic QbE results (%) on test data, where any instance of a query is masked from the search utterances.

Model	$P@10$	$P@N$	EER	Spearman's ρ
<i>Baselines:</i>				
Random	9.2	9.0	50.0	5.4
Multilingual AWE	21.5	15.6	46.1	6.4
<i>Semantic AWE model:</i>				
CLUSTER+SKIP-GRAM	29.9	23.1	32.3	22.9

7.8. EXPERIMENTAL RESULTS




Figure 7.11: PCA projection of semantic AWEs (averaged), produced by our CLUSTER+SKIP-GRAM model on development data. We highlight the five nearest neighbours for “several” (pink), “bike” (green), and “orange” (blue).

7.8.3 QUALITATIVE ANALYSIS

Figure 7.11 visualises the semantic embedding space produced by the CLUSTER+SKIP-GRAM model. It is clear that the acoustic realisations of semantically related words end up in similar areas in the space. For example, the model learned that spoken instances of “orange”, “red”, “blue”, “yellow”, and “green” should be close to each other. For additional visualisation, we direct the reader back to Figure 1.1 and Figure 1.2. Figure 1.1a displays a selected number of TWEs produced by a Skip-gram model trained on the transcribed speech we use for acoustic semantic modelling in this chapter (Section 7.7.1). Figure 1.1b displays the AWE produced by the multilingual AWE model we use in our semantic AWE approaches (Section 7.5). Finally, Figure 1.2 displays the embeddings (without averaging) produced by our CLUSTER+SKIP-GRAM model. From these visualisations, it is evident that the semantic AWEs satisfactory capture of word-type and semantic information.

The intrinsic results in Table 7.4 and qualitative analysis in Figure 1.2 show that our

Table 7.7: The top ten closest embeddings for keywords, derived from text (left) and speech segments (right), respectively. Speech embeddings were produced using the CLUSTER+SKIP-GRAM model and averaged per word-type before nearest neighbour calculation.

american		swimming		boys		football	
flags	african	pool	swimmers	children	girls	opponent	baseball
spiderman	waving	swim	swims	kids	children	soccer	hockey
tin	older	water	swim	men	kids	hockey	soccer
camouflage	costume	wet	into	girls	women	player	uniform
muzzles	elderly	diving	zooms	boy	teens	players	miami
participating	flags	headfirst	swings	players	men	volleyball	player
elderly	asian	wading	pool	people	are	hitting	approaches
tiny	leather	paddling	scratching	dolphins	traveller	uniform	apportioned
instrument	establishment	swims	scratches	gateway	cans	field	saxophone
asian	instruments	patiently	splashes	girl	adults	swim	players

7.9. CHAPTER SUMMARY

semantic embeddings, to some extent, resemble the semantic information captured in the corresponding text embeddings. To get a better understanding of the difference between the embeddings produced from speech and the corresponding transcripts, we look at the variation in the nearest embeddings of some words. Table 7.7 displays chosen keywords, comparing semantic AWEs (on the right) with textual word embeddings (on the left). Firstly, we notice that across all keywords some words are shared (highlighted in bold). Secondly, we observe a strong phonetic influence on the semantic embeddings trained on speech. They are more likely to have words higher ranked that are phonetically similar to the keyword. For example, “american” and “football”, have “african” and “baseball”, as their closest neighbour, respectively. In this instance, we could say they are semantically similar from our own judgement, but they do not appear among the nearest words for the corresponding text embeddings. Looking at the keyword “swimming” we can see where this phonetic disturbance has a negative impact. For “swimming”, the word “swims” appear in the top neighbours for both speech and text. Because “swims” are phonetically similar to “zooms” and “swings”, the model fails to separate them in vector space. As a result, “swimming” also ends up close to “zooms” and “swings”. In this instance, we would not regard them as semantically relevant to each other.

7.9 CHAPTER SUMMARY

In this chapter, we addressed a new form of acoustic representation learning, where word embeddings should capture the meaning of words, rather than phonetic information. To achieve this, we presented several semantic AWE modelling strategies. We specifically promoted transferring knowledge from a pre-trained multilingual AWE model trained for word-class discrimination. Through intrinsic and extrinsic evaluations, we demonstrated the effectiveness of our strategies in learning semantic representations from unlabelled speech data. Our best semantic AWE approach involves a soft clustering on the original multilingual AWEs, serving as input to a Skip-gram model.

The main shortcoming of our work (as also in others [63, 64]), is that the word segmentation is assumed to be known. This was reasonable given our goal of comparing different semantic AWE approaches on a sensible benchmark, but future work should look into incorporating unsupervised word segmentation methods [136–138] in order to do fully unsupervised semantic AWE modelling.

CHAPTER 8

SUMMARY AND CONCLUSIONS

Developing robust speech systems for zero-resource languages—where no transcribed speech resources are available for model training—remains a challenge. Most of these systems, at their core, require speech segments of different duration to be compared. Previous studies showed the benefit of AWEs over alignment methods for this task. In a zero-resource setting, the two existing training strategies are unsupervised monolingual learning (Section 2.4.1) and supervised multilingual transfer (Section 2.4.2). This thesis mainly contributes to the development of AWEs following the multilingual transfer strategy. We presented multiple approaches toward improving the intrinsic quality of AWEs as well as the implementation of AWEs in a downstream speech retrieval task. Furthermore, we extended the multilingual transfer strategy to semantic representation learning, demonstrating the versatility of multilingual AWEs in upstream tasks. We first summarise our main findings. We then provide recommendations on how the work can be extended and applied in future research.

8.1 MAIN FINDINGS

In Chapter 3, we presented the CONTRASTIVERNN AWE model. This model optimises a distance metric function not considered for AWE modelling in previous work. We compared the CONTRASTIVERNN to two existing AWE embedding models, CAE-RNN [21] and SIAMESERNN [32]. We re-implemented both existing models and compared our model in the same experimental setup as Kamper et al. [23]. On a word-discrimination task measuring the intrinsic quality of the AWEs, the CONTRASTIVERNN significantly outperformed the CAE-RNN and SIAMESERNN following the unsupervised monolingual training strategy, reaching up to 17.8% absolute increase in average precision on one of the evaluation languages. In the supervised multilingual setting, we observed marginal performance increases from the CONTRASTIVERNN compared to the CAE-RNN and SIAMESERNN.

In Chapter 4, we proposed a new training strategy that extends the multilingual transfer strategy. After training a multilingual model, instead of only applying it to

8.1. MAIN FINDINGS

an unseen target language, we first fine-tune the multilingual model using discovered word-like pairs obtained from applying an unsupervised term discovery system (UTD) to unlabelled speech data in the target language. For each of the three multilingual models, we adapt a model to every zero-resource language before applying it. Here we observe that the CONTRASTIVERNN generally adapts better to a target language compared to the CAE-RNN and SIAMESE-RNN.

In Chapter 5, we investigated the impact a particular choice of training languages have when training an AWE model for a specific target language. We specifically want to see the impact of including training languages that belong to the same language family as the target language. We performed multiple experiments on data from languages in South Africa where we train multilingual AWEs with different language combinations while controlling for the amount of data. We show that including related languages over unrelated languages are beneficial, evaluated on an intrinsic word discrimination task. We also evaluated the multilingual AWEs in a downstream zero-resource speech task, query-by-example, supporting the results of the intrinsic evaluation task. To summarise the findings in this chapter, including training data from related languages, even in small quantities, are beneficial over using unrelated languages. For instance, a multilingual model trained on approximately 6 hours of related languages, outperform a model trained on approximately 60 hours of unrelated languages. The findings in this chapter would be of great interest to practitioners who want to develop speech applications in a zero-resource setting.

In Chapter 6, we developed keyword spotting systems to monitor hate speech in radio broadcast audio in a collaborative effort with a company, VoxCroit. Our goal was to compare existing ASR-based KWS systems (requiring labelled training data from the target language) to a new multilingual AWE-based KWS system. We first developed our systems in a controlled environment where we used parallel audio-text data from experimental datasets to train and test our models. In this in-domain environment, an ASR model trained on only five minutes of transcribed audio outperforms our AWE-based KWS system. We then applied the models, without further calibrations, to the real radio broadcast audio scraped from radio stations in Kenya. In this in-the-wild test, our AWE-based KWS system proves to be more robust against domain mismatch and performs similar to an ASR-KWS system using 30 hours of labelled target audio. Our findings in this chapter suggest multilingual AWEs can be used as an alternative to existing ASR systems for rapidly deploying KWS systems for languages with limited transcribed audio.

In Chapter 7, we shifted our attention to a new form of representation learning of spoken word segment: AWEs that reflect word-type information as well as word meaning. We specifically proposed using a pre-trained multilingual AWE model to assist in this task compared to existing approaches that only use unlabelled data from the target language [63, 64]. One of our models, CLUSTER+SKIP-GRAM, show great improvement

over previous models in an intrinsic word similarity task, measuring semantic relatedness. For the first time, we applied these semantic AWEs in a downstream semantic speech retrieval task. Our results show that semantic information is indeed being captured in the embeddings produced by our new models, indicating promise for future downstream applications.

8.2 FUTURE WORK

Throughout this thesis, we identified multiple opportunities to expand on in future research.

- All our AWE models up to Chapter 5 took hand-engineered MFCCs as input features. In Chapter 6 and Chapter 7 we incorporated self-supervised speech features as an alternative to MFCCs for input to our AWE models. These learned speech features showed improved performance in a variety of speech tasks [76–79]. We also observed the benefits in AWE modelling using these features: improved word discrimination (Section 6.7.1) and improved semantic representations (Table 7.4). However, we did not systematically investigate this but merely replaced features without considering alterations to our models or training conditions. We hypothesise further improvement by optimising the training setup to accommodate these speech features.
- We notice a large performance gap between the unsupervised monolingual models trained on discovered word pairs (top of Table 3.1) and the supervised monolingual models trained on true word segments (top of Table 4.2). A more accurate unsupervised segmentation of unlabelled speech data, delivering higher quality discovered pairs, could reduce the gap between the unsupervised and supervised training strategies. Kamper et al. [93] presented a full-coverage segmentation and clustering system—unlabelled speech data are segmented and clustered into unknown word types. Their system employs primitive acoustic word embeddings. By replacing the AWEs with higher quality AWEs, for instance, embeddings produced by a multilingual AWE model, better segmentation can be expected. Applying this segmentation system to unlabelled speech data in a zero-resource language could potentially yield higher quality discovered word pairs compared to the UTD system of Jansen and Van Durme [10] used in this thesis.
- Up until now, the choice of training languages in a multilingual transfer setting have not been considered, specifically for AWEs. We showed that a more careful selection of training languages can improve performance when developing a multilingual AWE model for a specific target language by including training languages that belong to the same language family as the target. We showed that the performance increase is not simply due to the overlap of words in related languages. However, a more

8.2. FUTURE WORK

in-depth analysis of the phonetic and syntactic properties common to languages from the same language family could give more insight assisting the development of multilingual transfer learning in general.

- In Chapter 6, we identified the need for a more optimal threshold value to improve KWS performance. We hypothesise that a per-keyword threshold could be determined based on factors including the length of the speech segment, the number of phones present, and the relative distances between other unknown word segments in the acoustic space. In the work of Matuskevych et al. [104], they analyse an AWE space and discover that the embeddings encode absolute duration and speaker information. Additionally, they observe that the distance between word embeddings increases with phonetic dissimilarity. With such knowledge we could explore potential correlations between these factors stored in AWEs and the distance between instances in the acoustic space.
- In Chapter 7, we introduced a novel approach for learning speech representations that reflect semantic meaning instead of acoustic similarity. Although our approach exhibits improved performance compared to existing methods, certain limitations remain: Similar to others [63, 64], we used true word boundaries for segmenting unlabelled audio. Future work should consider unsupervised segmentation methods [136–138], or explore an end-to-end semantic AWE approach that eliminates the need for explicit word boundaries.

Additionally, we performed our experiments within a constrained setup involving a limited vocabulary and a smaller number of words when compared to the millions or even billions used for training textual word embeddings. Future work should therefore investigate the scalability of the semantic approaches we have introduced.

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] Eberhard, M. David, G. F. Simons, and C. D. Fenning, “Ethnologue: Languages of the world,” 2021. [Online]. Available: <https://www.ethnologue.com>
- [3] M. Versteegh, X. Anguera, A. Jansen, and E. Dupoux, “The Zero Resource speech challenge 2015: Proposed approaches and results,” in *Proc. SLTU*, 2016.
- [4] 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.-y. Lee, K. Levin, A. Norouzian, 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.
- [5] O. Räsänen, “Computational modeling of phonetic and lexical learning in early language acquisition: Existing models and future directions,” *Speech Commun.*, vol. 54, no. 9, pp. 975–997, 2012.
- [6] K. Levin, A. Jansen, and B. Van Durme, “Segmental acoustic indexing for zero resource keyword search,” in *Proc. ICASSP*, 2015.
- [7] S.-F. Huang, Y.-C. Chen, H.-y. Lee, and L.-s. Lee, “Improved audio embeddings by adjacency-based clustering with applications in spoken term detection,” *arXiv preprint arXiv:1811.02775*, 2018.
- [8] Y. Yuan, C.-C. Leung, L. Xie, H. Chen, B. Ma, and H. Li, “Learning acoustic word embeddings with temporal context for query-by-example speech search,” in *Proc. Interspeech*, 2018.
- [9] 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.

REFERENCES

-
- [10] A. Jansen and B. Van Durme, “Efficient spoken term discovery using randomized algorithms,” in *Proc. ASRU*, 2011.
 - [11] L. Ondel, H. K. Vydana, L. Burget, and J. Černocký, “Bayesian subspace hidden markov model for acoustic unit discovery,” in *Proc. Interspeech*, 2019.
 - [12] O. Räsänen and M. A. C. Blandón, “Unsupervised discovery of recurring speech patterns using probabilistic adaptive metrics,” *arXiv preprint arXiv:2008.00731*, 2020.
 - [13] H. Kamper, K. Livescu, and S. Goldwater, “An embedded segmental K-means model for unsupervised segmentation and clustering of speech,” in *Proc. ASRU*, 2017.
 - [14] S. Seshadri and O. Räsänen, “SylNet: An adaptable end-to-end syllable count estimator for speech,” *IEEE Signal Process. Letters*, vol. 26, no. 9, pp. 1359–1363, 2019.
 - [15] F. Kreuk, J. Keshet, and Y. Adi, “Self-supervised contrastive learning for unsupervised phoneme segmentation,” in *Proc. Interspeech*, 2020.
 - [16] L. Rabiner, A. Rosenberg, and S. Levinson, “Considerations in dynamic time warping algorithms for discrete word recognition,” *IEEE Trans. Acoust., Speech, Signal Process.*, vol. 26, no. 6, pp. 575–582, 1978.
 - [17] K. Levin, K. Henry, A. Jansen, and K. Livescu, “Fixed-dimensional acoustic embeddings of variable-length segments in low-resource settings,” in *Proc. ASRU*, 2013.
 - [18] N. Holzenberger, M. Du, J. Karadayi, R. Riad, and E. Dupoux, “Learning word embeddings: Unsupervised methods for fixed-size representations of variable-length speech segments,” in *Proc. Interspeech*, 2018.
 - [19] Y.-A. Chung, C.-C. Wu, C.-H. Shen, H.-Y. Lee, and L.-S. Lee, “Unsupervised learning of audio segment representations using sequence-to-sequence autoencoder,” in *Proc. Interspeech*, 2016.
 - [20] H. Kamper, W. Wang, and K. Livescu, “Deep convolutional acoustic word embeddings using word-pair side information,” in *Proc. ICASSP*, 2016.
 - [21] H. Kamper, “Truly unsupervised acoustic word embeddings using weak top-down constraints in encoder-decoder models,” in *Proc. ICASSP*, 2019.
 - [22] H. Kamper, Y. Matuskevych, and S. Goldwater, “Multilingual acoustic word embedding models for processing zero-resource languages,” in *Proc. ICASSP*, 2020.

REFERENCES

-
- [23] —, “Improved acoustic word embeddings for zero-resource languages using multilingual transfer,” *IEEE Trans. Audio, Speech, Language Process.*, vol. 29, pp. 1107–1118, 2021.
 - [24] Y. Hu, S. Settle, and K. Livescu, “Acoustic span embeddings for multilingual query-by-example search,” in *Proc. SLT*, 2021.
 - [25] —, “Multilingual jointly trained acoustic and written word embeddings,” in *Proc. Interspeech*, 2020.
 - [26] S. Ruder, “Neural transfer learning for natural language processing,” PhD diss., NUI Galway, 2019.
 - [27] D. Rumelhart, G. E. Hinton, and R. J. Williams, “Learning representations by back-propagating errors,” *Nature*, vol. 323, pp. 533–536, 1986.
 - [28] T. Mikolov, I. Sutskever, K. Chen, G. Corrado, and J. Dean, “Distributed representations of words and phrases and their compositionality,” in *Proc. NeurIPS*, 2013.
 - [29] J. Pennington, R. Socher, and C. Manning, “Glove: Global vectors for word representation,” in *Proc. EMNLP*, 2014.
 - [30] C. Jacobs, Y. Matuskevych, and H. Kamper, “Acoustic word embeddings for zero-resource languages using self-supervised contrastive learning and multilingual adaptation,” in *Proc. SLT*, 2021.
 - [31] G. Chen, C. Parada, and T. N. Sainath, “Query-by-example keyword spotting using long short-term memory networks,” in *Proc. ICASSP*, 2015.
 - [32] S. Settle and K. Livescu, “Discriminative acoustic word embeddings: Recurrent neural network-based approaches,” in *Proc. SLT*, 2016.
 - [33] C. Doersch and A. Zisserman, “Multi-task self-supervised visual learning,” in *Proc. ICCV*, 2017.
 - [34] Y. M. Asano, C. Rupprecht, and A. Vedaldi, “A critical analysis of self-supervision, or what we can learn from a single image,” in *Proc. ICLR*, 2020.
 - [35] C. Doersch, A. Gupta, and A. A. Efros, “Unsupervised visual representation learning by context prediction,” in *Proc. ICCV*, 2015.
 - [36] M. Noroozi and P. Favaro, “Unsupervised learning of visual representations by solving jigsaw puzzles,” in *Proc. ECCV*, 2016.

REFERENCES

-
- [37] S. Gidaris, P. Singh, and N. Komodakis, “Unsupervised representation learning by predicting image rotations,” in *Proc. ICLR*, 2018.
 - [38] S. Pascual, M. Ravanelli, J. Serrà, A. Bonafonte, and Y. Bengio, “Learning problem-agnostic speech representations from multiple self-supervised tasks,” 2019.
 - [39] G. Synnaeve, Q. Xu, J. Kahn, T. Likhomanenko, E. Grave, V. Pratap, A. Sriram, V. Liptchinsky, and R. Collobert, “End-to-end ASR: From supervised to semi-supervised learning with modern architectures,” in *Proc. ICML*, 2020.
 - [40] A. Baevski, S. Schneider, and M. Auli, “vq-wav2vec: Self-supervised learning of discrete speech representations,” in *Proc. ICLR*, 2020.
 - [41] A. Baevski, M. Auli, and A. Mohamed, “Effectiveness of self-supervised pre-training for speech recognition,” in *Proc. ICASSP*, 2020.
 - [42] W. Wang, Q. Tang, and K. Livescu, “Unsupervised pre-training of bidirectional speech encoders via masked reconstruction,” in *Proc. ICASSP*, 2020.
 - [43] T. Chen, S. Kornblith, M. Norouzi, and G. Hinton, “A simple framework for contrastive learning of visual representations,” in *Proc. ICML*, 2020.
 - [44] K. Sohn, “Improved deep metric learning with multi-class N-pair loss objective,” in *Proc. NeurIPS*, 2016.
 - [45] R. van der Merwe, “Triplet entropy loss: Improving the generalisation of short speech language identification systems,” *arXiv preprint arXiv:2012.03775*, 2020.
 - [46] J. Yi, J. Tao, Z. Wen, and Y. Bai, “Language-adversarial transfer learning for low-resource speech recognition,” *IEEE/ACM Trans. Audio, Speech, Language Process.*, vol. 27, no. 3, pp. 621–630, 2019.
 - [47] E. van der Westhuizen, T. Padhi, and T. Niesler, “Multilingual training set selection for ASR in under-resourced Malian languages,” in *Proc. SPECOM*, 2021.
 - [48] T. J. Hazen, W. Shen, and C. White, “Query-by-example spoken term detection using phonetic posteriorgram templates,” in *Proc. ASRU*, 2009.
 - [49] Y. Zhang and J. R. Glass, “Unsupervised spoken keyword spotting via segmental DTW on Gaussian posteriorgrams,” in *Proc. ASRU*, 2009.
 - [50] A. Jansen and B. V. Durme, “Indexing raw acoustic features for scalable zero resource search,” in *Proc. Interspeech*, 2012.

REFERENCES

-
- [51] A. Anastasopoulos, D. Chiang, and L. Duong, “An unsupervised probability model for speech-to-translation alignment of low-resource languages,” in *Proc. EMNLP*, 2016.
 - [52] S. Settle, K. Levin, H. Kamper, and K. Livescu, “Query-by-example search with discriminative neural acoustic word embeddings,” in *Proc. Interspeech*, 2017.
 - [53] Y. Yuan, C.-C. Leung, L. Xie, H. Chen, and B. Ma, “Query-by-example speech search using recurrent neural acoustic word embeddings with temporal context,” *IEEE Access*, vol. 7, pp. 67 656–67 665, 2019.
 - [54] D. Ram, L. Miculicich, and H. Bourlard, “Neural network based end-to-end query by example spoken term detection,” *IEEE Trans. Audio, Speech, Lang. Process.*, vol. 28, no. 1, pp. 1416–1427, 2019.
 - [55] A. Saeb, R. Menon, H. Cameron, W. Kibira, J. Quinn, and T. Niesler, “Very low resource radio browsing for agile developmental and humanitarian monitoring,” in *Proc. Interspeech*, 2017.
 - [56] R. Menon, A. Saeb, H. Cameron, W. Kibira, J. Quinn, and T. Niesler, “Radio-browsing for developmental monitoring in Uganda,” in *Proc. ICASSP*, 2017.
 - [57] R. Menon, H. Kamper, J. Quinn, and T. Niesler, “Fast ASR-free and almost zero-resource keyword spotting using DTW and CNNs for humanitarian monitoring,” in *Proc. Interspeech*, 2018.
 - [58] R. Menon, H. Kamper, E. v. d. Westhuizen, J. Quinn, and T. Niesler, “Feature exploration for almost zero-resource ASR-free keyword spotting using a multilingual bottleneck extractor and correspondence autoencoders,” in *Proc. Interspeech*, 2019.
 - [59] “United Nations strategy and plan of action on hate speech.” [Online]. Available: https://www.un.org/en/genocideprevention/documents/advising-and-mobilizing/Action_plan_on_hate_speech_EN.pdf
 - [60] M. Larson and G. Jones, “Spoken content retrieval: A survey of techniques and technologies,” *Found. Trends Inform. Retrieval*, pp. 235–422, 2012.
 - [61] A. Mandal, K. R. Prasanna Kumar, and P. Mitra, “Recent developments in spoken term detection: A survey,” *Int. J. of Speech Technol.*, vol. 17, pp. 183–198.
 - [62] E. van der Westhuizen, H. Kamper, R. Menon, J. Quinn, and T. Niesler, “Feature learning for efficient ASR-free keyword spotting in low-resource languages,” *Comput. Speech Lang.*, vol. 71, p. 101275, 2022.

REFERENCES

-
- [63] Y.-A. Chung and J. Glass, “Speech2Vec: A sequence-to-sequence framework for learning word embeddings from speech,” in *Proc. Interspeech*, 2018.
 - [64] Y.-C. Chen, S.-F. Huang, C.-H. Shen, H.-y. Lee, and L.-s. Lee, “Phonetic-and-semantic embedding of spoken words with applications in spoken content retrieval,” in *Proc. SLT*, 2019.
 - [65] I. Goodfellow, Y. Bengio, and A. Courville, *Deep learning*. MIT Press, 2016.
 - [66] Y. Bengio, P. Simard, and P. Frasconi, “Learning long-term dependencies with gradient descent is difficult,” *IEEE Trans. Neural Networks*, vol. 5, no. 2, pp. 157–166, 1994.
 - [67] A. Shewalkar, D. Nyavanandi, and S. Ludwig, “Performance evaluation of deep neural networks applied to speech recognition: RNN, LSTM and GRU,” *J. Artif. Intell. Soft Comput.*, vol. 9, pp. 235–245, 2019.
 - [68] S. Hochreiter and J. Schmidhuber, “Long short-term memory,” *Neural Computation*, vol. 9, no. 8, pp. 1735–1780, 1997.
 - [69] P. Baldi, “Autoencoders, unsupervised learning, and deep architectures,” in *Proc. ICML*, 2012.
 - [70] H. Kamper, M. Elsner, A. Jansen, and S. Goldwater, “Unsupervised neural network based feature extraction using weak top-down constraints,” in *Proc. ICASSP*, 2015.
 - [71] J. Bromley, I. Guyon, Y. LeCun, 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.
 - [72] J. Wang, Y. song, T. Leung, C. Rosenberg, J. Wang, J. Philbin, B. Chen, and Y. Wu, “Learning fine-grained image similarity with deep ranking,” in *Proc. CPVR*, 2014.
 - [73] G. Chechik, V. Sharma, U. Shalit, and S. Bengio, “Large scale online learning of image similarity through ranking,” *J. Mach. Learn. Res.*, vol. 11, pp. 1109–1135, 2010.
 - [74] A. Hermans, L. Beyer, and B. Leibe, “In defense of the triplet loss for person re-identification,” *arXiv preprint arXiv:1703:07737*, 2017.
 - [75] E. Hoffer and N. Ailon, “Deep metric learning using triplet network,” in *Proc. SIMBAD*, 2015.
 - [76] A. Baevski, Y. Zhou, A. Mohamed, and M. Auli, “wav2vec 2.0: A framework for self-supervised learning of speech representations,” in *Proc. NeurIPS*, 2020.

REFERENCES

-
- [77] A. Conneau, A. Baevski, R. Collobert, A. Mohamed, and M. Auli, “Unsupervised cross-lingual representation learning for speech recognition,” in *Interspeech 2021*, 2021.
 - [78] W.-N. Hsu, B. Bolte, Y.-H. H. Tsai, K. Lakhota, R. Salakhutdinov, and A. Mohamed, “HuBERT: Self-supervised speech representation learning by masked prediction of hidden units,” *IEEE/ACM Trans. Audio, Speech, Lang. Process.*, vol. 29, pp. 3451–3460, 2021.
 - [79] S. Chen, C. Wang, Z. Chen, Y. Wu, S. Liu, Z. Chen, J. Li, N. Kanda, T. Yoshioka, X. Xiao, J. Wu, L. Zhou, S. Ren, Y. Qian, Y. Qian, J. Wu, M. Zeng, X. Yu, and F. Wei, “WavLM: Large-scale self-supervised pre-training for full stack speech processing,” *IEEE J. Sel. Top. Signal Process.*, vol. 16, no. 6, pp. 1505–1518, 2022.
 - [80] D. R. H. Miller, M. Kleber, C.-L. Kao, O. Kimball, T. Colthurst, S. A. Lowe, R. M. Schwartz, and H. Gish, “Rapid and accurate spoken term detection,” in *Proc. Interspeech*, 2007.
 - [81] M. Saraclar and R. Sproat, “Lattice-based search for spoken utterance retrieval,” in *Proc. HLT-NAACL*, 2004.
 - [82] P. Yu, K. Chen, C. Ma, and F. Seide, “Vocabulary-independent indexing of spontaneous speech,” *IEEE Trans. Speech, Audio Process.*, vol. 13, no. 5, pp. 635–643, 2005.
 - [83] K. Ng and V. Zue, “Subword-based approaches for spoken document retrieval,” *Speech Commun.*, 2000.
 - [84] N. Rajput and F. Metze, “Spoken web search,” in *Proc. MediaEval Workshop*, 2011.
 - [85] E. Barnard, M. Davel, C. van Heerden, X. Anguera, G. Gravier, and N. Rajput, “The spoken web search task,” in *Proc. MediaEval Workshop*, 2012.
 - [86] F. Metze, A. Buzo, I. Szoke, and L. J. Rodriguez-Fuentes, “The spoken web search task,” in *Proc. MediaEval Workshop*, 2013.
 - [87] X. Anguera, L. Javier, I. Szöke, A. Buzo, and F. Metze, “Query by example search on speech at Mediaeval 2014,” in *Proc. MediaEval Workshop*, 2014.
 - [88] I. Szoke, F. Metze, L. Javier, J. Proenca, A. Buzo, M. Lojka, X. Anguera, and X. Xiong, “Query by example search on speech at Mediaeval 2015,” in *Proc. MediaEval Workshop*, 2015.

REFERENCES

-
- [89] G. Mantena and X. Anguera, “Speed improvements to information retrieval-based dynamic time warping using hierarchical K-Means clustering,” in *Proc. ICASSP*, Oct. 2013.
 - [90] Y. Zhang and J. Glass, “A piecewise aggregate approximation lower-bound estimate for posteriorgram-based dynamic time warping,” in *Proc. Interspeech*, 2011.
 - [91] H. Kamper, A. Anastassiou, and K. Livescu, “Semantic query-by-example speech search using visual grounding,” in *Proc. ICASSP*, 2019.
 - [92] 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.
 - [93] —, “Unsupervised word segmentation and lexicon discovery using acoustic word embeddings,” *IEEE Trans. Audio, Speech, Language Process.*, vol. 24, no. 4, pp. 669–679, 2016.
 - [94] L. van Staden, “Improving unsupervised acoustic word embeddings using segment- and frame-level information,” Thesis, Stellenbosch University, Stellenbosch, 2021.
 - [95] S. J. Pan and Q. Yang, “A survey on transfer learning,” *IEEE Trans. Knowl. Data Eng.*, vol. 22, no. 10, pp. 1345–1359, 2009.
 - [96] M. A. Carlin, S. Thomas, A. Jansen, and H. Hermansky, “Rapid evaluation of speech representations for spoken term discovery,” in *Proc. Interspeech*, 2011.
 - [97] R. Algayres, M. S. Zaiem, B. Sagot, and E. Dupoux, “Evaluating the reliability of acoustic speech embeddings,” in *Proc. Interspeech*, 2020.
 - [98] B. M. Abdullah, M. Mosbach, I. Zaitova, B. Möbius, and D. Klakow, “Do acoustic word embeddings capture phonological similarity? An empirical study,” in *Proc. Interspeech*, 2021.
 - [99] F. Schroff, D. Kalenichenko, and J. Philbin, “FaceNet: A unified embedding for face recognition and clustering,” in *Proc. CPVR*, 2015.
 - [100] A. Medela and A. Picon, “Constellation Loss: Improving the efficiency of deep metric learning loss functions for optimal embedding,” *J. Pathol. Informat.*, vol. 11, no. 1, p. 38, 2020.
 - [101] P. Bachman, R. D. Hjelm, and W. Buchwalter, “Learning representations by maximizing mutual information across views,” in *Proc. NeurIPS*, 2019.

REFERENCES

-
- [102] T. Schultz, T. Vu, and T. Schlippe, “GlobalPhone: A multilingual text & speech database in 20 languages,” in *Proc. ICASSP*, 2013.
 - [103] D. P. Kingma and J. Ba, “Adam: A method for stochastic optimization,” in *Proc. ICLR*, 2017.
 - [104] Y. Matuskevych, H. Kamper, and S. Goldwater, “Analyzing autoencoder-based acoustic word embeddings,” in *BAICS Workshop ICLR*, 2020.
 - [105] L. v. d. Maaten and G. Hinton, “Visualizing data using t-SNE,” *J. Mach. Learn. Res.*, vol. 9, no. 86, pp. 2579–2605, 2008.
 - [106] R. Bedyakin and N. Mikhaylovskiy, “Low-resource spoken language identification using self-attentive pooling and deep 1D time-channel separable convolutions,” *arXiv preprint arXiv:2106.00052*, 2021.
 - [107] I. Orife, J. Kreutzer, B. Sibanda, D. Whitenack, K. Siminyu, L. Martinus, J. T. Ali, J. Abbott, V. Marivate, S. Kabongo, M. Meressa, E. Murhabazi, O. Ahia, E. van Biljon, A. Ramkilowan, A. Akinfaderin, A. Öktem, W. Akin, G. Kioko, K. Degila, H. Kamper, B. Dossou, C. Emezue, K. Ogueji, and A. Bashir, “Masakhane – Machine translation for Africa,” in *Proc. ICLR*, 2020.
 - [108] N. Society, “Family of language,” 2020. [Online]. Available: <http://www.nationalgeographic.org/encyclopedia/family-language/>
 - [109] E. Barnard, M. Davel, C. van Heerden, F. Wet, and J. Badenhorst, “The NCHLT speech corpus of the South African languages,” in *Proc. SLTU*, 2014.
 - [110] T. Probert and M. d. Vos, “Word recognition strategies amongst isiXhosa/English bilingual learners: The interaction of orthography and language of learning and teaching,” *Reading & Writing*, vol. 7, no. 1, p. 10, 2016.
 - [111] F. Chalk, “Hate Radio in Rwanda,” in *The Path of a Genocide: The Rwanda Crisis from Uganda to Zaire*, 1999.
 - [112] K. Somerville, “Violence, hate speech and inflammatory broadcasting in Kenya: The problems of definition and identification,” *Ecquid Novi: African Journalism Studies*, vol. 32, no. 1, pp. 82–101, 2011.
 - [113] E. I. Odera, “Radio and hate speech: a comparative study of Kenya (2007 pev) and the 1994 Rwanda genocide,” Ph.D. dissertation, University of Nairobi, 2015.
 - [114] G. Chen, C. Parada, and G. Heigold, “Small-footprint keyword spotting using deep neural networks,” in *Proc. ICASSP*, 2014.

REFERENCES

-
- [115] K. Audhkhasi, A. Rosenberg, A. Sethy, B. Ramabhadran, and B. Kingsbury, “End-to-end ASR-free keyword search from speech,” in *Proc. ICASSP*, 2017.
 - [116] E. T. Mekonnen, A. Brutti, and D. Falavigna, “End-to-end low resource keyword spotting through character recognition and beam-search re-scoring,” in *Proc. ICASSP*, 2022.
 - [117] E. Gauthier, L. Besacier, S. Voisin, M. Melese, and U. P. Elingui, “Collecting resources in sub-Saharan African languages for automatic speech recognition: A case study of Wolof,” in *Proc. LREC*, 2016.
 - [118] R. Ardila, M. Branson, K. Davis, M. Henretty, M. Kohler, J. Meyer, R. Morais, L. Saunders, F. M. Tyers, and G. Weber, “Common Voice: A massively-multilingual speech corpus,” in *Proc. LREC*, 2020.
 - [119] A. Babu, C. Wang, A. Tjandra, K. Lakhotia, Q. Xu, N. Goyal, K. Singh, P. von Platen, Y. Saraf, J. Pino, A. Baevski, A. Conneau, and M. Auli, “XLS-R: Self-supervised cross-lingual speech representation learning at scale,” in *Proc. Interspeech*, 2022.
 - [120] A. Graves, S. Fernandez, F. Gomez, and J. Schmidhuber, “Connectionist temporal classification: Labelling unsegmented sequence data with recurrent neural networks,” in *Proc. ICML*, 2006.
 - [121] M. McAuliffe, M. Socolof, S. Mihuc, M. Wagner, and M. Sonderegger, “Montreal forced aligner: Trainable text-speech alignment using Kaldi,” in *Proc. Interspeech*, 2017.
 - [122] R. Sanabria, H. Tang, and S. Goldwater, “Analyzing acoustic word embeddings from pre-trained self-supervised speech models,” in *Proc. ICASSP*, 2023.
 - [123] M. L. Seltzer, D. Yu, and Y. Wang, “An investigation of deep neural networks for noise robust speech recognition,” in *Proc. ICASSP*, 2013.
 - [124] T. Likhomanenko, Q. Xu, V. Pratap, P. Tomasello, J. Kahn, G. Avidov, R. Collobert, and G. Synnaeve, “Rethinking evaluation in ASR: Are our models robust enough?” *arXiv preprint arXiv:2010.11745*, 2021.
 - [125] W.-N. Hsu, A. Sriram, A. Baevski, T. Likhomanenko, Q. Xu, V. Pratap, J. Kahn, A. Lee, R. Collobert, G. Synnaeve, and M. Auli, “Robust wav2vec 2.0: Analyzing domain shift in self-supervised pre-training,” in *Proc. Interspeech*, 2021.
 - [126] L. McInnes, J. Healy, N. Saul, and L. Großberger, “UMAP: Uniform manifold approximation and projection,” *Journal of Open Source Software*, vol. 3, no. 29, p. 861, 2018.

REFERENCES

-
- [127] T. Mikolov, K. Chen, G. Corrado, and J. Dean, “Efficient estimation of word representations in vector space,” in *Proc. ICLR*, 2013.
 - [128] D. Harwath and J. Glass, “Deep multimodal semantic embeddings for speech and images,” in *Proc. ASRU*, 2015.
 - [129] H. Kamper, S. Settle, G. Shakhnarovich, and K. Livescu, “Visually grounded learning of keyword prediction from untranscribed speech,” in *Proc. Interspeech*, 2017.
 - [130] B. M. Abdullah, B. Möbius, and D. Klakow, “Integrating form and meaning: A multi-task learning model for acoustic word embeddings,” in *Proc. Interspeech*, 2022.
 - [131] G. Chen and Y. Cao, “A reality check and a practical baseline for semantic speech embeddings,” in *ICASSP*, 2023.
 - [132] H. Kamper, G. Shakhnarovich, and K. Livescu, “Semantic speech retrieval with a visually grounded model of untranscribed speech,” *IEEE/ACM Trans. Audio, Speech, and Lang. Process.*, vol. 27, no. 1, pp. 89–98, 2019.
 - [133] C. Rashtchian, P. Young, M. Hodosh, and J. Hockenmaier, “Collecting image annotations using Amazon’s mechanical turk,” in *Proc. NAACL HLT*, 2010.
 - [134] F. Hill, R. Reichart, and A. Korhonen, “SimLex-999: Evaluating semantic models with (genuine) similarity estimation,” *Comput. Linguist.*, vol. 41, no. 4, pp. 665–695, 2015.
 - [135] L. van Staden and H. Kamper, “A comparison of self-supervised speech representations as input features for unsupervised acoustic word embeddings,” in *Proc. SLT*, 2021.
 - [136] T. S. Fuchs, Y. Hoshen, and J. Keshet, “Unsupervised word segmentation using K nearest neighbors,” in *Proc. Interspeech*, 2022.
 - [137] S. Cuervo, M. Grabias, J. Chorowski, G. Ciesielski, A. Łańcucki, P. Rychlikowski, and R. Marxer, “Contrastive prediction strategies for unsupervised segmentation and categorization of phonemes and words,” in *Proc. ICASSP*, 2022.
 - [138] H. Kamper, “Word segmentation on discovered phone units with dynamic programming and self-supervised scoring,” *IEEE/ACM Trans. Audio, Speech, Lang. Process.*, vol. 31, pp. 684–694, 2023.