

Data Augmentation for Speech Recognition in Maltese: A Low-Resource Perspective

Carlos Mena^{*1}, Andrea DeMarco², Claudia Borg², Lonneke van der Plas³, and Albert Gatt^{2,4}

¹Reykjavik University, Iceland

²University of Malta, Malta

³Idiap Research Institute, Switzerland

⁴Utrecht University, The Netherlands

Abstract

Developing speech technologies is a challenge for low-resource languages for which both annotated and raw speech data is sparse. Maltese is one such language. Recent years have seen an increased interest in the computational processing of Maltese, including speech technologies, but resources for the latter remain sparse. In this paper, we consider data augmentation techniques for improving speech recognition for such languages, focusing on Maltese as a test case. We consider three different types of data augmentation: unsupervised training, multilingual training and the use of synthesized speech as training data. The goal is to determine which of these techniques, or combination of them, is the most effective to improve speech recognition for languages where the starting point is a small corpus of approximately 7 hours of transcribed speech. Our results show that combining the three data augmentation techniques studied here lead us to an absolute WER improvement of 15% without the use of a language model.

Keywords Data Augmentation; Speech Recognition; Maltese Language; Unsupervised Transcriptions; Multilingual Training; Synthesized Speech

1 Introduction

In recent years, the field of Natural Language Processing has seen renewed interest in multilingual models. New techniques of transfer and multi-task learning, as well as the availability of multilingual datasets, have motivated the development of large-scale reusable models (Devlin et al., 2019; Conneau et al.,

^{*}This work was conducted while all the authors were at the University of Malta. Correspondence should be addressed to Carlos Mena: carlosm.ru.is

2020) and NLP benchmarks (e.g. Hu et al., 2020; Liang et al., 2020). Many of these models rely on large quantities of data for pretraining. As a result, ‘under-resourced’ languages, for which such data is less easy to come by, remain under-represented in these developments. To take an example, the large-scale language models currently in use in many language understanding tasks, such as Multilingual BERT (Devlin et al., 2019) or XLM-R (Conneau et al., 2020) are pretrained on multilingual data available in online resources such as Wikipedia or CommonCrawl (Wenzek et al., 2020). This, however, also means that languages with small speaker populations, which also tend to be under-represented on the web, are not represented in these models.

Similar developments are being seen in Automatic Speech Recognition (ASR) (Baevski et al., 2020a). There has been a lot of work focused on transfer learning-based improvements to speech recognition systems in order to make many modern architectures accessible to a wider set of languages. This has taken a number of different forms. For example Wang et al. (2020) use high-to-low resource language machine translation as an intermediate task to create a language model for a low-resource target language. More in line with parallel developments in NLP, recent research has also turned towards large, multilingual, pretrained ASR architectures. For example, Pratap et al. (2020) show that great benefits for low-resource speech recognition can accrue from a multilingual acoustic model. However, the majority of languages used in this work have over 100 hours of training data available. A different line of research has been pursued in Wav2Vec2 (Baevski et al., 2020b), which builds on the earlier Wav2Vec model Schneider et al. (2019). This architecture can learn robust speech representations of sufficient quality to perform recognition even for target languages with a very low volume of labelled data (in the order of one hour), reaching the same performance as previous state-of-the-art systems requiring 100 times more data (100 hours). Wav2Vec2 can also provide good results with just ten minutes of labelled speech data, reaching word error rates as low as 4.8, given substantial pretraining on 53,000 hours of unlabelled data. This work therefore demonstrates that semi- or self-supervised pretraining, similar in spirit to what we have seen in recent years with language models in NLP, can result in robust representations that lower the requirements on labelled data. Once again, however, the key is in ensuring that sufficient data is available for pretraining, in this case in the form of audio. Indeed, Baevski et al. (2020b) describe two settings, one with 960 hours of unlabelled data, and another with 60,000 hours.

In summary, research in NLP in both the textual and speech modalities has convincingly shown the benefits of pretraining with minimal supervision. Yet, from the perspective of under-resourced languages with low presence on the web, in the form of either text or speech, their feasibility is far from guaranteed. Maltese provides a good example of this kind of scenario. Web-scale textual or speech data is harder to obtain for Maltese compared to languages such as English or Mandarin and we expect that similar challenges arise for many other under-resourced languages (Besacier et al., 2014). Furthermore, as recent critiques of large-scale pretraining approaches have emphasised, even where web-scale data is available, there are significant risks arising from its ‘un-

fathomable’ nature, not least that it is likely to be extremely noisy, while not guaranteeing representativeness across demographic or ethnic groups, and/or across language varieties Bender et al. (2021); Rogers (2021). Lastly, the computational resources needed for such experiments are not available to all research teams.

Given these challenges, this paper presents an exploratory analysis of a set of related techniques for improving ASR for Maltese using different data augmentation strategies, relying on the use of smaller, curated datasets in the target language, the use of larger datasets in related languages, and the contribution of artificial data augmentation. Briefly, the scenario in which these experiments were designed consists of the following:

- Approximately 7 hours of labelled speech data in the target language (here: Maltese);
- A small amount of unlabelled speech data in the target language;
- A medium sized corpus of Maltese text (250m tokens);
- A concatenative speech synthesis system whose output, while not of high quality, can be exploited for data augmentation purposes;
- Substantial data in a number of other languages – notably Arabic, Italian and English – which are typologically and historically related to Maltese.

In the remainder of this section, we first introduce some salient features of the Maltese language itself, followed by an explanation of the motivation for these experiments. The section concludes with an outline of the rest of the paper.

1.1 Maltese and other languages

Maltese is spoken as a first language by just under half a million citizens of the Maltese islands (Mena et al., 2020); it is also taught as a second language in a few countries. It is the national language of Malta, but co-exists with two other official languages (English and Maltese Sign Language).

In recent years, there has been significant gain in digital support for the Maltese language. A large (250m tokens) annotated text corpus is available, as are tools for segmentation and morphosyntactic labelling (Gatt and Čéplö, 2013) and electronic lexical repositories (Camilleri, 2013).¹ In addition, some studies have focused on natural language analysis tasks such as morphological labelling (Borg and Gatt, 2017; Ravishankar et al., 2017) and dependency parsing (Tiedemann and van der Plas, 2016; Zammit, 2018).

Despite these advances, Maltese remains under-resourced on a number of fronts, especially where speech is concerned. Concatenative speech synthesis

¹Some of these tools are available online through the Maltese Language Resource Server: <https://mlrs.research.um.edu.mt>.

systems have been developed using classical (e.g. diphone synthesis) methods (Micallef, 1997; Borg et al., 2011), but there are currently no tools for ASR, except for some preliminary experiments done by Mena et al. (2020). The most important reason for this is a lack of resources. In the absence of large repositories of paired speech and text samples, it is not realistic to use current end-to-end methods for speech recognition. As argued above, it is also hard to obtain large samples of unlabelled speech for pretraining purposes.

The present paper focuses on a variety of approaches designed to overcome this bottleneck, conducted in the context of the project *MASRI - Maltese Automatic Speech Recognition*². Starting from a relatively small dataset of Maltese text and speech, we describe experiments deploying a variety of techniques for data augmentation. The broad question we address is therefore the following: *Given an under-resourced language for which a relatively small speech-to-text dataset is available, what data augmentation methods work best to bring ASR performance to a level which provides a suitable, competitive baseline for future development?*

In our experiments, we maintain a focus on data augmentation for end-to-end ASR. While our focus is on the Maltese language, the findings presented here are of relevance to speech technology researchers working on other under-resourced languages.

While a variety of data augmentation methods can be envisaged, our focus on Maltese makes two lines of inquiry particularly relevant. Here, we give an outline of the motivation for each one.

Typological relatedness Maltese bears strong historical relationships to three major languages. From a historical perspective, it has been characterised as having a Semitic/Arabic stratum, a Romance (Italian/Sicilian) superstratum, and an English adstratum (Brincat, 2011). This characterisation reflects the historical development of Maltese, originally a variety of Arabic, which came into intensive contact with Italian due to its geographical proximity and historical relations with that country. This was followed by a period of intensive contact with English (Malta was a British colony from 1800 to 1964), as a result of which, the Maltese population is largely bilingual. The impact of this linguistic history is clearly evident in the language at many levels of analysis, including the lexical and morphological (Borg and Gatt, 2017). From the perspective of the present experiments, this suggests that a promising way to approach the data augmentation problem is to exploit the (much larger) resources available for these languages. We do this by transcribing speech data from one of these languages with a baseline Maltese acoustic model trained on a very small corpus (Section 5.2); and using a ‘language mixture’ approach (Section 5.3), in which an end-to-end system is pretrained on data using mixtures of such languages. In both cases, pretraining is followed by a fine-tuning step.

²<https://www.um.edu.mt/projects/masri/>

Synthesis Given the existence of a concatenative speech synthesis system for Maltese, we also discuss data augmentation experiments involving artificially synthesised training data (Section 5.4).

1.2 Outline of the paper

The rest of this paper is structured as follows. Section 2 reviews data augmentation methods, especially for speech. In Section 3 we describe the datasets we use whilst Section 4 describes the experimental methodology. Section 5 presents a variety of experiments and results, in which we aim to identify the best data augmentation methods for ASR in Maltese. Section 6 concludes with a general discussion and some pointers to future work.

2 Data augmentation for speech recognition

Data augmentation is a technique utilized in the field of machine learning that consists of modifying the existing data in some way, or using data from other sources that are easy and/or cheap to obtain, in order to increase the effective size of the corpus with which a system is trained, obtaining better performance. For example, one of the earliest applications of data augmentation was applied in the field of image processing. It consists in rotating the images of the dataset multiple times (Simard et al., 2003), in order to increase the amount of the effective images in the training set.

In the literature, one may find several types of data augmentation methods used in the field of speech recognition. In this section, we present a review of these methods, from among which we identify those that fit better with the small amount of Maltese language data available and with modest computational resources.

According to the taxonomy presented in Ramirez et al. (2019), the types of data augmentation for ASR can be divided in three groups: semi-supervised training, transformation of acoustic data, and speech synthesis. On the other hand, in Ragni et al. (2014) the authors argue that “the schemes can be split based on the type of produced data, such as unsupervised, synthesized and other language data”. In Cui et al. (2015), it is mentioned that there exists one kind of data augmentation that uses “label-preserving transformations” (Jaitly and Hinton, 2013), where data is altered but not its labels. This latter type of data augmentation can include several subtypes, such as transformation of acoustic data and other-language data - this however does not qualify as unsupervised or semi-supervised training, because in such cases, the labels are generated, and not “preserved” from their original form.

One other issue that comes out of the literature is the different definitions assigned to the different methods of data augmentation; for example in Ragni et al. (2014), the authors argue that “the unsupervised data refers to data which lacks correct transcriptions. This also includes data having only rough transcriptions, such as closed captions”. This definition is problematic because

it overlaps with the definition of semi-supervised training given in Ramirez et al. (2019), which claims: “the semi-supervised training approach assumes the use of the text produced by an automatic speech recognition system to train acoustic models”. The problem is that the transcriptions produced by an ASR system can also be full of errors; more than that, closed captions can also be generated by an ASR system.

Given the somewhat variable nature of terminology and classification in the field, in the remainder of this section, we organise related work using our own data augmentation taxonomy. Our goal is to preserve, as much as possible, the main ideas found in literature, but at the same time, avoid clashes between definitions. Our taxonomy is:

- Semi-supervised Training
- Data Perturbation
- Training with Synthesized Data
- Unsupervised Training
- Multilingual Training
- Training with Text without Audio

2.1 Semi-supervised training

Semi-supervised data refers to data which lacks “gold” transcriptions (Ragni et al., 2014). These transcriptions can be obtained in different ways. For example, they can be supplied manually by non-expert transcribers, or by harvesting closed captions. Semi-supervised training is sometimes also referred to in the literature as “lightly-supervised training” (Gales et al., 2006). Some examples of this can be found in Ragni et al. (2014); Gales et al. (2006); Lamel et al. (2002).

2.2 Data perturbation

In this approach, the speech signal is modified in some way, but assigned the same associated transcription (label). The most common example of this scheme is to add noise to the speech signal (Ramirez et al., 2019; Jimerson et al., 2018). Variation of the speech rate and the pitch (Thai et al., 2019; Vachhani et al., 2018; Kanda et al., 2013) is also a popular technique. Vocal Tract Length Perturbation (VTLP) is a more sophisticated method to perturb the input signal and it is part of this category as well (Billa, 2018; Jaitly and Hinton, 2013; Cui et al., 2015). More recently, SpecAugment (Park et al., 2019a) was proposed, performing rectangular cuts in the log mel spectrogram. We integrated SpecAugment as part of our experiments (see section 4.5).

2.3 Synthesized data

Data augmentation based on synthesized data refers to new, artificially generated data. For example, it is common practice to create synthesized speech from raw text with a Text to Speech (TTS) system, and then combine it with a regular ASR corpus (Rossenbach et al., 2020; Tjandra et al., 2017; Rosenberg et al., 2019). Nevertheless, one can find more sophisticated types of synthesized data, as in Gales et al. (2009), where some Hidden Markov Models (HMMs) are created artificially to help support vector machines (SVMs) to outperform their classification capabilities. In our experiments, we use synthesized speech from a simple TTS system based on HMMs (see section 5.4).

2.4 Unsupervised training

Unsupervised training data refers to data with transcriptions generated only by an ASR system. It can be confused with the semi-supervised approach in the sense that the unsupervised transcriptions can have a lot of errors too, but as discussed in sections 2 and 2.1, the semi-supervised transcriptions can come from a human source, and not necessarily from an ASR system. That is the main difference.

In this paper, we distinguish unsupervised training from the semi-supervised training discussed in section 2.1. For convenience, we reserve the term “semi-supervised” to noisy transcriptions obtained from human annotators. The term “unsupervised” training will refer only to setups involving data from transcriptions generated by an ASR system.

Unsupervised training is commonly associated in the literature with transcribing speech in a language A, using an ASR system trained in a language B, as in Ragni et al. (2014); Löff et al. (2009); Qian et al. (2013). However, this need not be the case. One can transcribe speech in a language A using an ASR system in the same language (Zavaliagkos and Colthurst, 1998; Cucu et al., 2014; Kemp and Waibel, 1999; Wessel and Ney, 2004). In our experiments, we use an ASR system in Maltese to decode other languages as well as our target language (see sections 5.2, 5.3, 5.4).

2.5 Multilingual training

In this scheme of data augmentation, the idea is to take advantage of the speech in a language different from the target, in order to augment the training data available for the desired system. In the literature, related procedures are sometimes referred to as “data borrowing” (Qian et al., 2013) or even “out-of-language data” (Imseng et al., 2014).

However, the way in which this multilingual data is used may vary, depending on factors such as the amount of data available in the target language or the type of system utilized to do the training. For example, a common scenario is when the amount of target language data is small, but then it is combined with data in other languages (Billa, 2018; Yilmaz et al., 2018; Qian et al., 2013;

Imseng et al., 2014). Nevertheless, the scenario of not having any data at all in the target language can also occur, as in Lööf et al. (2009).

Sometimes the original transcriptions in the other, non-target languages are available, so one way to combine all the languages in one system is using a phoneme set in common (Billa, 2018; Lööf et al., 2009; Knill et al., 2013; Vu and Schultz, 2013; Grézl et al., 2014). Sometimes, for various reasons, transcriptions in the non-target languages are not available or they are not usable; in that case, one can use an ASR system in the target language to produce unsupervised transcriptions of the other languages, in order to use them to train an “augmented” acoustic model that will be the combination of the data in the target language plus the data in other languages (Ragni et al., 2014; Qian et al., 2013).

Commonly used techniques for developing such multilingual systems include multi-layer perceptrons (MLP) (Tüske et al., 2014; Qian et al., 2013; Vu and Schultz, 2013; Imseng et al., 2014) and bottleneck features (Tüske et al., 2014; Knill et al., 2013; Tüske et al., 2013). In our particular case, we will be using an End-to-End (E2E) system to produce unsupervised transcriptions from distinct languages which are typologically or historically related to the target language Maltese, such as Arabic, Italian and English (see sections 5.2, 5.3, 5.4). In addition to that, we experiment with the original transcriptions of some languages related to Maltese.

2.6 Text without audio

Unlike other data augmentation techniques discussed above, in this setting the goal is to exploit raw text with no associated audio. This, however, should not be confused with just adding text to the language model of an ASR system, which is a standard procedure in the creation of such systems. (Huang et al., 2001).

The literature contains a number of examples of techniques that benefit from augmentation via raw text, in combination with other resources. For example, Hori et al. (2019) created a text-to-encoder that, in combination with a Text to Speech (TTS) system, is used to generate the hidden states of an end-to-end ASR system; the goal is to be able to use text without associated audio. In a similar fashion, Hayashi et al. (2018) built a text-to-encoder model which learns to predict the hidden states of an E2E-ASR encoder. Tjandra et al. (2017) claim that “while ASR transcribes the unlabeled speech features, TTS attempts to reconstruct the original speech waveform based on the text from ASR. In the opposite direction, ASR also attempts to reconstruct the original text transcription given the synthesized speech”.

We also include the work of Yilmaz et al. (2018) in this category. In this paper, the authors artificially generated text with code-switching to be used as data augmentation in the language model.

As we do not use a language model in our experiments (see section 4.6 for our motivation), we did not have to opportunity to test this particular technique of data augmentation. We will leave this for future work.

3 Data and Resources

The focus of this paper is on improving ASR for Maltese, an under-resourced language, using data augmentation for pretraining. We explore two principal data augmentation strategies, namely (i) creating synthetic data, and (ii) using both unsupervised (noisy) and supervised (gold) data from related languages in a transfer setting. This section gives an overview of the tools and datasets used.

3.1 Speech synthesis

The Maltese Speech Engine³ (Borg et al., 2014) is a diphone-based synthesis system. It relies on a rule-based grapheme to phoneme (G2P) tool, which transforms text into its equivalent IPA representation. Since Maltese is fairly homographic, the rules are primarily based on one-to-one mappings with a list of exceptions.⁴

Although the TTS speech engine produces mostly intelligible Maltese, it is far from natural-sounding and contains several pronunciation errors. This justifies our labelling the data created using this engine as ‘noisy’. Nevertheless, we contend that such noisy data can still prove useful in pretraining, if followed by fine-tuning on carefully curated speech-text pairs. Hence, we used the Maltese Speech Engine to create a synthetic corpus, described in Section 3.2.3 below.

3.2 Maltese Corpora

3.2.1 Korpus Malti text corpus

The Korpus Malti v3.0 (Gatt and Čéplö, 2013) is a text corpus of around 250m tokens in several different genres, including parliamentary debates, news, law, opinion articles, sports articles, culture, academic, literature and religious texts.⁵ Tokens in the corpus are tagged with part of speech, and labelled with lemmas and (where relevant) consonantal root for words of a Semitic origin. We use this text corpus for creating synthesized data as described in section 3.2.3.

3.2.2 The MASRI-HEADSET corpus

The MASRI-HEADSET Corpus (Mena et al., 2020) is, to our knowledge, the first published ASR corpus in Maltese. It comprises 8 hours of read speech, recorded in a low noise environment. The corpus contains 25 different speakers (13 female; 12 male) and includes orthographic transcriptions. We refer to this

³The Maltese Speech Engine is available for download <https://fitamalta.eu/projects/maltese-speech-engine-synthesis-erdf-114/download-maltese-speech-engine/> and runs on MS Windows.

⁴A reimplemented Python version of the G2P tool is available from the Maltese Language Resource Server: <https://mlrs.research.um.edu.mt/index.php?page=api>. A demo is also available on the same server.

⁵The corpus is available on the Maltese Language Resource Server, and can also be searched through an online interface. See: <https://mlrs.research.um.edu.mt/index.php?page=corpora>

corpus as the MASRI-HEADSET version 1. In our experiments, we split the corpus into two unique portions: a training set, with a size of 7 hours and 34 minutes; and a test set of 31 minutes.

Since its publication, MASRI-HEADSET v1 was further modified to improve quality, by removing excessive periods of silence, suppressing clicks and segmenting audio files longer than 11 seconds. This resulted in what we refer to as MASRI-HEADSET version 2, with a total duration of 06h39m, against the original 08h06m.

For the purposes of the present experiments, we used the same amount of utterances (250) in the test split of MASRI-HEADSET v2 as in the original v1 corpus. This resulted in a duration of 6 hours and 19 minutes for the training set and 19 minutes for the test set.

3.2.3 The MASRI-SYNTHETIC corpus

We produced a synthetic speech corpus using the Maltese Speech Engine described above, by feeding the speech engine sentences from the Korpus Malti v3.0. A total of 52,500 sentences were sampled, after filtering according to the following criteria:

- All sentences sampled consisted of 30 words or less;
- No sentences containing non-Maltese orthographic characters (such as *y*, which only occurs in loan words), foreign words or proper names were included.

The resulting speech corpus, which we refer to as the MASRI-SYNTHETIC corpus, consists of 105 male and 105 female voices, produced with varied pitch (21 values of pitch ranging between -20 and 20) and speech rate (5 values between -2 and 2). Each voice was assigned 250 utterances of 13 words each, resulting in utterances of varying length between 2 and 10 seconds duration. The corpus has a total duration of 99 hours and 18 minutes.

3.2.4 Maltese audiobooks

As an additional source of untranscribed data, we use a set of audiobook recordings, obtained from a Maltese publisher. These consist of 12 hours and 24 minutes of recorded speech.

3.3 Speech Corpora from Related Languages

Another aspect of our research is to experiment with transfer learning from related languages. Thus we used a number of corpora described below.

3.3.1 LibriSpeech

LibriSpeech⁶ is an ASR corpus of read speech in English (Panayotov et al., 2015), derived from audiobooks that are part of the LibriVox project⁷. It contains around 1,000 hours of speech sampled at 16 kHz and partitioned into different training, development and test sets according to quality. For our work, we combined all the different training sets into a single pool, irrespective of the quality, and then randomly selected male-female pairs of speakers. The resulting selection from LibriSpeech is a 100-hour corpus with the only constraint that total speech time is perfectly balanced across the two genders.

3.3.2 Arabic Speech MGB-2 corpus

Given the influence of Arabic on the Maltese language, we decided to include Arabic as part of our transfer learning experiments. Arabic Speech⁸ is a collaborative effort aimed at increasing speech resources for Arabic and its dialects. For this research, we selected the MGB-2 corpus, a multi-dialect broadcast news corpus consisting of 1,200 hours, collected from Aljazeera TV programs and made available by the Arabic Speech project. This corpus was used in the 2016 MGB challenge (Ali et al., 2016). For our purposes, we used a 100-hour portion.

This corpus includes transcriptions in both Arabic and the Buckwalter (BW) transliteration (Habash et al., 2007). Since Maltese uses the Roman alphabet, and the models we experiment with rely on latin characters, we could not use Arabic transcriptions. We therefore rely on BW transliterations, noting that these lack vowel annotations, which is a source of potential ambiguity for an ASR system. Given the provenance of the Arabic data, it is worth noting that it is noisier than the other corpora described here. In particular, several audio files contain audible speech from more than one person (e.g. a newscaster, with a simultaneous translator), as well as background noise or music.

We use a 100-hour, random sample of the corpus. Due to the gender imbalance in the corpus as a whole (a far higher representation of male speakers), the sample is also not balanced for male versus female voices.

3.3.3 Italian M-AILABS speech dataset

The M-AILABS Speech Dataset⁹ is an open source corpus combining data from LibriVox and Project Gutenberg¹⁰. The data consists of nearly a thousand hours of audio clips of read speech in eight different languages. Clips vary in length, from 1 to 20 seconds, and include the corresponding transcriptions. In our experiments, we only use the Italian portion of this corpus with a total

⁶LibriSpeech can be downloaded from: <http://www.openslr.org/12>

⁷LibriVox project allows volunteers to read and record chapters from books in the public domain. See: <https://librivox.org/>

⁸<https://arabicspeech.org/>

⁹<https://www.caito.de/2019/01/the-m-ailabs-speech-dataset/>

¹⁰Project Gutenberg is a collection of books that are free from copyright. See: <http://www.gutenberg.org/>

length of 118 hours and 41 minutes. This portion was taken as is, with no gender balancing or any other preprocessing made to the audio or text data.

4 Experimental Setup

In this section, we describe in general terms, how the experiments in Section 5 are performed, as well the features that they have in common.

4.1 Pretraining and Fine-tuning stages

All the experiments reported in section 5.2 onwards, are split into two separate stages: pretraining and fine-tuning. Data augmentation methods are applied during the pretraining stage. The fine-tuning stage uses a pretrained model as a starting point or bootstrap. fine-tuning is always performed with the training portion of MASRI-HEADSET version 2. All experiments utilise the MASRI-HEADSET v2 test set.

4.2 Architectures

In all our experiments, we use Jasper (Just Another Speech Recognizer), a family of End-to-End speech recognizers created by NVIDIA (Li et al., 2019). We use the version included in the Neural Modules (NeMo) framework¹¹ (Kuchaiev et al., 2019).

A Jasper architecture typically consists of n consecutive blocks, each with r convolutional layers or sub-blocks. Each sub-block is composed of a 1D separable convolution followed by batch normalisation, a ReLU layer and dropout. Finally, all Jasper architectures have four additional convolutional blocks: one for pre-processing and three for post-processing. Table 1 summarises the pre-processing parameters used for all experiments.

Table 1: Jasper preprocessing layer parameters

Parameter	Value
Window type	Hanning
Window size	20ms
Window stride	10ms
Padding	16
Feature type	logbank
Normalization scope	Per feature
Number of features	64
Apply Short-Time Fourier Transform (STFT)?	Yes
Fast Fourier Transform (FFT) data points	512
Dither value	0.00001

¹¹<https://github.com/NVIDIA/NeMo>

4.3 Execution Parameters

A number of execution parameters are common to all experiments without exception; these are summarised in Table 2. As shown in the table, we use the Novograd optimizer, which is similar to Adam (Kingma and Ba, 2014), except that its second moments are computed per layer instead of per weight (Li et al., 2019). The loss is calculated based on the Connectionist Temporal Classification (CTC) algorithm (Graves et al., 2006).

Table 2: Common execution parameters of all the experiments

Parameter	Value
Batch size	32
Number of epochs	50
Learning rate	0.02
Optimizer	Novograd
Loss function	CTC

4.4 Generation of Noisy Transcriptions

A number of experiments in the following section describe data augmentation using unsupervised (‘noisy’) transcription. This is always performed in the following way: a set of audio files in some language is transcribed using an acoustic model trained to produce transcriptions in Maltese. These transcriptions can have a real meaning when the input files are in Maltese, but if the input is in a different language, the transcriptions may only contain text that looks like Maltese, but it is not (‘pseudo-Maltese’). It has to be emphasized that all the unsupervised transcriptions produced in the present work are only in Maltese or pseudo-Maltese.

4.5 Noising via SpecAugment

In all experiments, we use SpecAugment (Park et al., 2019b), a technique which has been shown to improve the robustness of end-to-end ASR models, by augmenting data with noise introduced directly into the feature inputs of the model. Table 3 summarises the parameters used.

Table 3: SpecAugment parameters

Parameter	< 100h	≥ 100h
Rectangles placed randomly	5	5
Vertical stripes (time)	2	120
Horizontal stripes (frequency)	2	50

In general, SpecAugment parameters for 100 hours or more are applied in the pretraining stage, while parameters for less than 100 hours are applied during fine-tuning.

4.6 Language Model-less experimentation

None of the experiments reported below incorporates language model (LM) rescoring. There are two reasons for this. First, once an LM is used to rescore and improve ASR output, it results in modifications of the transcriptions produced by an acoustic model. This occurs because an LM adjusts the transcriptions produced by the acoustic model, based on the LM vocabulary.¹² (Huang et al., 2001; Jurafsky et al., 1995). Since our goal is to determine the impact of noisy or synthetic data augmentation on acoustic model quality, the inclusion of LM rescoring would obscure the contribution of these techniques. Clearly, the results reported in the remainder of this paper could be improved by the inclusion of LM rescoring, but this is orthogonal to the present goals.

The second reason is a direct consequence of the small amount of Maltese data available. The only ASR corpus used at the time the present experiments were performed is MASRI-HEADSET (see section 3.2.2), whose sentences were extracted from the Korpus Malti v3.0 (Section 3.2.1). The latter is also the largest currently available corpus on which a language model of reasonable size could be built. Rescoring with an LM built on this corpus would give better results for an acoustic model trained on MASRI-HEADSET (the baseline model described in Section 5.1 below), given the substantial overlap in their vocabulary. On the other hand, filtering the test utterances from Korpus Malti v3.0 prior to creating the LM would introduce a bias against the test utterances of the MASRI-HEADSET which are sourced from the same corpus.

To demonstrate these potential biases, we created two different LMs using the settings described by Mena et al. (2020). The Maltese text was extracted from the Korpus Malti v3.0 and excluded sentences containing extensive borrowing from English, as well as digits and proper names.¹³ Our initial LM (which is biased in favour of the baseline) is a 6-gram model (ca. 9.5MB) which includes test set sentences from MASRI-HEADSET. A second LM ('against') was created by filtering the test set sentences (ca. 8.7MB).

Table 4 shows the results obtained on the MASRI-HEADSET test set, using acoustic models trained only with the training portion of the same corpus (the baselines described in Section 5.1), as well as the best model obtained in our experiments (Section 5.4). We show results both with and without rescoring using either of the two LMs.

As expected, with LM rescoring, the difference between the two models (baseline versus best) is smaller, compared to when the acoustic models are

¹²This is also why, in a classic ASR system based on HMM, both the vocabulary extracted from the LM and from the training transcriptions, have to be included in a pronunciation model such as a pronouncing dictionary.

¹³English words were identified using the CMU Pronunciation Dictionary (Weide, 1998). To identify proper names we used the CIEMPIESS-PNPD (Mena, 2019).

Table 4: Word Error Rate (WER) for baseline and best models, with and without the LMs

	Baseline	Best
No LM	63.71%	48.97%
LM Against	59.04%	54.67%
LM in Favor	20.87%	16.65%

compared in their own right (first row of Table 4). Furthermore, we note that a baseline system tested on data which overlaps with the data on which the LM was obtained (‘LM in favour’) achieves a much better WER than the best model.

These results indicate the best performance we can expect when a language model is incorporated, both with the best augmented model, and with a baseline trained on a very limited (ca. 7h) training corpus. However, the use of an LM also diverts attention from the main question, which focuses on the most effective pretraining methods in low-resource scenarios. Thus, in what follows, all results are reported with no LM rescoring.

5 Experiments

In this section, we turn to the experiments on data augmentation for Maltese ASR. We first describe an initial experiment to validate the architecture which will be used throughout the remainder of the experiments. Each subsequent set of experiments is characterised by a pretraining and a fine tuning stage. Since our emphasis is on data augmentation, it is the pretraining stage that mainly distinguishes the approaches described here. Our goal is to identify to what extent different data augmentation strategies work for an under-resourced language such as Maltese.

5.1 Baselines

As our baselines, we compare two variants of the Jasper architecture described in Section 4.2:

1. Jasper 12×1 SEP: This consists of 12 blocks, each with 1 sub-block and uses separable convolutions.¹⁴ This architecture was originally described in the context of experiments on 100 hours of English data obtained from the LibriSpeech corpus. During training, the model is exposed to audio files sampled at a rate of 16 khz, with a maximum length of 16.7 seconds. Silence is trimmed at the beginning and end of each audio file, and transcriptions are normalised to expand abbreviations, convert numbers

¹⁴See: <http://www.openslr.org/12>

to letters, normalise consecutive whitespace and lowercase all text. This normalisation process was optimised for English.

2. `Jasper_an4_new`: This is an adaptation of Jasper 12×1 to the set of languages and the amount of data available for our experiments. The architecture consists of 4 blocks with 2 repetitions of the 1D convolutional layer, making it smaller than Jasper 12×1 SEP. This is motivated by the empirical observation that a leaner architecture exhibits more stable behaviour (in terms of WER) with less training data, making it our architecture of choice. In this architecture, we apply a dropout rate of 0.2 and place no restriction of length of audio files. Silence is not trimmed and no text normalisation is performed.

First, we establish that the two models achieve comparable performance after training on 100h of English data from LibriSpeech, then we compare them on the low-resource scenario of the present work, training them on both a small sample of English data from LibriSpeech and Maltese data from the MASRI-HEADSET corpus, with approximately 7h of data in each case.

Table 5: Results on varying training set sizes for the two architectures

Architecture	Train	Test	Loss	WER
Jasper 12×1 SEP	LibriSpeech train-clean-100	LibriSpeech dev-clean	44.37	29.54%
	M-HEADSET v1 7h34m	M-HEADSET v1 31m	177.86	99.87%
	LibriSpeech 7h34m	LibriSpeech 31m	190.50	70.06%
jasper_an4_new	LibriSpeech train-clean-100	LibriSpeech dev-clean	37.67	29.21%
	LibriSpeech 7h34m	LibriSpeech 31m	103.88	48.61%
	M-HEADSET v1 7h34m	M-HEADSET v1 31m	69.28	65.61%
	M-HEADSET v2 6h19m	M-HEADSET v2 19m	58.65	63.82%

The results suggest that on relatively large training sets of around 100h, the two architectures achieve comparable results (WER $\sim 29\%$). However, with a much smaller training set of less than 10h, `Jasper_an4_new` outperforms the

other. Though performance on Maltese is far from optimal, these experiments yield a baseline, consisting of a relatively small architecture with minimal data augmentation beyond the random noising strategy incorporated in SpecAugment.

In what follows, we adopt this architecture and consider several ways to improve the results through pretraining. We note once again that all results reported below are for the acoustic models with no language model rescoring, as we are seeking to establish the most effective pretraining strategies to improve over the baseline.

5.2 Unsupervised transcription

In our first set of experiments, we consider unsupervised (noisy) transcription as a technique for data augmentation. The main question we address is whether, given a baseline system trained on a very small dataset (see section 5.1), it is possible to generate noisy data for further pretraining and subsequent fine-tuning. To this end, we experiment with noisy transcription of data in Maltese (limited to around 12h), but also in other languages. In the latter case, we ask whether using a baseline Maltese ASR model to transcribe data from Arabic, Italian or English, can yield suitable pretraining data for further fine-tuning. The rationale for using non-target language data is that, in the presence of limited audio data for Maltese, pretraining using a baseline Maltese ASR system amounts to ‘reinterpreting’ the non-target data using a representation that should bring it closer to the target. At the same time, our selection of data from non-target languages rests on observations about their typological proximity or historical relationship to Maltese (see section 1).

Data The following datasets are used to create noisy pretraining data:

- Maltese: 12h of audiobooks in Maltese (Section 3.2.4)
- English: 100h hours portion of the LibriSpeech corpus (Section 3.3.1)
- Arabic: 100h portion of the MGB-2 Corpus (Section 3.3.2)
- Italian: 118h from the Italian portion of the M-AILABS Speech Dataset (Section 3.3.3)

In addition, we reserve the manually transcribed MASRI-HEADSET v2 corpus for fine-tuning, reserving 250 utterances for testing (Section 3.2.2).

Procedure Pretraining proceeds in cycles, each consisting of a pretraining phase and a fine-tuning stage. Each cycle uses the fine-tuned model from the previous cycle to regenerate the noisy transcriptions for the pretraining. Each experiment consists of three cycles, the optimum having been determined empirically: beyond 3 cycles, we begin to observe diminishing returns and oscillating WER values.

Table 6: Results for different pretraining datasets, over 3 cycles

Dataset	Cycle	LOSS	WER
MALTESE	1	55.82	61.64%
	2	49.35	63.17%
	3	58.95	60.95%
ENGLISH	1	59.01	58.46%
	2	53.43	58.73%
	3	58.30	58.27%
ARABIC	1	64.67	60.26%
	2	64.34	59.15%
	3	60.06	58.15%
ITALIAN	1	57.36	60.22%
	2	60.84	57.12%
	3	59.99	63.06%

Results Table 6 summarises the results on all experiments, reporting the final test-set WER for each language used for data augmentation, over three pretrain-and-fine-tune cycles. Two main conclusions can be drawn from these results. First, pretraining with around 100h of noisy data from non-target languages yields improved WER results. Second, the same noisy transcription strategy applied to a smaller amount of data from the target language itself also yields benefits (see the top panel in Table 6). However, at around 61% WER, this is less than what we observe for pretraining over three cycles with English or Arabic. For both of these languages, we see a gradual WER drop over the three cycles. Similar observations apply to loss. On balance, with a lower loss and comparable WER, it appears that pretraining on English unsupervised (noisy) transcription is the more stable solution compared to Arabic. On the other hand, when pretraining on Italian we observe fluctuations in WER, with the best results obtained on the second cycle, though with a concomitant increase in loss.

Summary Overall, the evidence suggests that non-target audio data in sufficient quantities, processed using a model trained on the target language, can yield modest improvements in target language ASR. Differential benefits are observed for pretraining in the non-target languages, perhaps in part due to the degree of lexical overlap between Maltese and the other three languages. A natural next step is to consider pretraining on data mixtures from different non-target languages, which also enables us to increase the pretraining dataset size. This is the focus of the next set of experiments.

5.3 Language mixtures

The rationale of the experiments in the previous subsection was that reinterpreting non-target data using a baseline ASR system for the target language should map the data to a representation that the target ASR system can benefit from with further training. Under this working assumption, we now consider the effect of mixing data from combinations of the three non-target languages, with a view to increasing the amount of pretraining data. Our goal is to determine if gains in terms of WER increase with dataset size when the data consists of language mixtures.

Data We use the same pretraining data as in the previous subsection, presenting results for pretraining with combined datasets of two or more languages, including the small dataset of Maltese audiobook data. In what follows, we only report results for combinations involving English, alone or with other languages, based on the observation in the previous section that pretraining on English data results in stable gains in both WER and loss.

Procedure Each experiment consists of a single cycle of pretraining followed by fine-tuning, following the same procedure as in Section 5.2.

Table 7: Results of mixed language pretraining using noisy transcription. Language columns indicate the amount of pretraining data in a given experiment.

EN	MT	IT	AR	Loss	WER
100h	12h24m	-	-	55.85	59.84%
100h	-	118h41m	-	56.90	58.96%
100h	-	-	100h	61.33	58.38%
100h	-	118h41m	100h	63.98	59.15%
100h	12h24m	-	100h	58.95	59.04%
100h	12h24m	118h41m	100h	63.71	57.96%

Results Table 7 displays results for different language mixtures. The results suggest that the benefits of this strategy are limited, and no notable gains in WER accrue with an increase in pretraining data; nor are there notable improvements from mixing languages. However, two further observations are worth making. First, mixing a relatively large corpus of English with a small sample of Maltese, the target language, yields WER improvements that outweigh those observed from using Maltese alone (see Table 6). Second, the best results are obtained when pretraining on all four languages; however, the benefits are modest compared to the results observed in the previous section with individual languages.

5.3.1 Pretraining using supervised transcriptions and unsupervised data

Finally, we alter the pretraining strategy, while maintaining the same rationale of augmenting data using language mixtures from related languages. We noted above that mixing English and Maltese in the pretraining yields better WER scores than only using a limited amount of Maltese (compare Table 7 and Table 6). Rather than pretrain using English data transcribed with a baseline Maltese ASR system, we now use the gold English annotations in the LibriSpeech corpus, and mix these with noisy transcriptions in Maltese during pretraining.

Table 8: Mixed language pretraining results with gold annotations for English

EN	MT	Loss	WER
100h (gold)	12h24m	48.62	55.05%

Results Table 8 shows that this strategy yields a WER of around 55%. This exceeds the best results obtained for the language mixture experiment. It shows the benefits of using supervised transcriptions for a language that is related to the target language in addition to noisy unsupervised annotations in the target language.

Summary The experiments in this subsection suggest that not much is gained from mixing languages during pretraining using noisy data, compared to pretraining on individual languages. On the other hand, surprising improvements are obtained from mixing gold English transcriptions with noisy transcription of a small dataset in the target language. In the following section, we examine to what extent increasing the pretraining data in the target language using artificial means is beneficial on its own, as well as in tandem with gold data in English.

5.4 Synthesising Training Data

In this section, we describe experiments using speech synthesis as a data augmentation tool, using the synthesized corpus described in Section 3.2.3. Despite the fact that the synthesis method yields noisy, imperfect data, synthetic training data has the advantage of being in potentially unlimited supply. In line with the conclusions from the previous subsection, we also consider the benefits of combining synthetically augmented target language data with gold data in English. In this way, we compare the trade-off between increasing data in the target language using artificial means (synthesis) versus adding data from a related language.

Data We use the Maltese synthesized speech corpus described in Section 3, as well as the gold transcriptions from the 100h of LibriSpeech English data used previously.

Procedure We present results from two experiments. First, we establish baseline results using only the MASRI-SYNTHETIC Corpus. Next, we extend the findings from the previous subsection, where we find that the inclusion of gold English transcribed data with noisy Maltese data yields the best improvements. Here, we examine the impact of mixing gold English transcribed data with the synthesized data during pretraining. In both cases, as before, we fine-tune on the MASRI-HEADSET v.2 dataset.

Table 9: Results of augmentation using synthesised speech

Pretraining	Loss	WER
MT (synth)	48.63	52.03%
MT (synth) + EN 100h (gold)	42.78	48.97%

Results Results are presented in Table 9. The use of 100h of synthesized speech brings our best WER down to 52%, despite the noise present in the data. Once again, we observe a benefit from the inclusion of gold annotations in English during pretraining, with WER dropping to just under 49%, which is an absolute reduction of 15% in WER from our baseline systems (63.71% WER).

5.5 Summary of experiments and findings

The experiments described in this section addressed different methods for data augmentation to overcome bottlenecks in the training data for an under-resourced language. Broadly speaking, we can characterise the rationale in terms of two distinct strategies: (a) pretraining by combining limited target data with comparatively large datasets from languages which are typologically or historically related to the target language; (b) pretraining using synthesized data. An important variable was the extent to which the data used in pretraining was ‘noisy’: in the case of pretraining with data from other languages, we experiment with both a noisy transcription using a baseline ASR system, and gold annotations. In the case of synthesis, we expect the data to be imperfect in view of the limitations of the available TTS system.

Our experiments suggests that language mixtures yield benefits, even when the transcription is noisy. However, gold annotations are better; in particular, the inclusion of English gold data with noisy Maltese transcriptions yielded significant gains. Furthermore, we observe that pretraining on imperfect synthesized data in the target language also improves performance, with further

gains provided once more by the inclusion of gold English data: an absolute reduction of 15% in WER from the baseline system.

6 Conclusions and future work

The work presented in this paper had one over-arching theme - the analysis of whether ASR performance for under-resourced languages can be improved by various data augmentation methods. We elicited a detailed taxonomy of what types of augmentation exist in the literature, and we also devised a consistent set of experiments that sift through the many possibilities, whilst providing clear results for each. The results obtained are specific to ASR for the Maltese language, but we postulate that similar approaches would yield similar improvements for other under-resourced languages.

In order to assess that we were indeed testing for improvements in acoustic modelling based on data augmentation, this work specifically makes no use of language modelling as part of the ASR pipeline. We do however show that language model rescoring does indeed have a big effect on WER (section 4.6), and in future experimentation, we intend to combine the best methods from this work with further language modelling experiments to obtain state-of-the-art performance for Maltese ASR. That is however, a different research question altogether.

As part of our efforts in maintaining consistency in our experiments, questions were also raised as to whether augmentation is helpful for pretraining methodologies, whilst keeping the network architectures unchanged. To this end we employ training cycles, and show that some measure of performance improvement is obtained, as shown in Section 5.2. We also postulate that this approach is generic enough to be proposed for any ASR setup for an under-resourced target language.

The one aspect of the work we present here that is probably dependent on the target language is in section 5.3. The ancillary language data chosen is not random. This augmentation is based on languages which are somehow related to Maltese - either historically, culturally or linguistically. The results show that supervised transcriptions of some closely related languages provide a substantial improvement in WER (Section 5.3.1). The choice of these languages, however, has to be assessed individually based on the target language.

Furthermore, this paper then analyzed the use of synthesized speech for ASR training (Section 5.4). Whilst data from a speech synthesis system might not always be readily available for under-resourced languages, we highlight the fact that the synthetic speech quality need not be of very good quality for immediate benefits in WER. In fact, the synthesis system used is produces highly unnatural speech with many pronunciation errors, yet we gained the largest improvements from using this method of data augmentation.

In summary, the conclusions we can draw from the work presented is as follows:

- Both gold and noisy transcriptions can be used as data augmentation techniques, up until a limit is reached after various training cycles.
- Noisy transcriptions have poorer performance with respect to gold transcriptions. In the absence of sufficient gold transcription quantities, however, the utilization of noisy transcriptions of the target language shows substantial improvements in WER.
- Mixing pretraining data from non-target languages is useful, especially with a small data batch from the target language. There are however limits to this and adding more data beyond a certain point yields diminishing returns.
- The use of synthesized speech in the target language, as training data, outperforms the use of gold or noisy transcriptions in languages different to the target, even when the synthesized speech has a low quality.

The results we obtain from this work are very encouraging, as we have observed an absolute reduction of 15% in WER from our baseline systems (63.71% WER) to the system resulting from the best setup based on pretraining and fine tuning from an augmented dataset (48.97% WER). This is a remarkable improvement in light of the fact that they revolve around a curated Maltese ASR corpus of less than 7 hours of speech.

We believe this work sets a starting point for all future work in ASR for the Maltese language, and propose that all languages which are similarly under-resourced could follow the methodology we outline in this paper to assess and capitalize on data augmentation methods.

A promising direction for future work is to explore the augmentation techniques presented here with a different neural architecture. Recent research in multilingual NKP has showing the remarkable performance of models in the Transformer paradigm (Vaswani et al., 2017), both for textual, language-understanding tasks with models such as multilingual BERT (Devlin et al., 2019) and XLM-R (Conneau et al., 2020) and for speech recognition with Wav2Vec-2 Baevski et al. (2020b). These approaches provide impetus for further experimentation, with a view to identifying the best strategies for data augmentation to make ASR more feasible for low-resource languages.

Ethics statement

The authors declare that they have no conflict of interest. All data used in this study is publicly available and was previously published in Mena et al. (2020). The data collection and research reported here was screened by the University of Malta Research Ethics Committee (UREC; Application no. 728:5.02.2019) and conforms to the University of Malta Research Code of Practice, available at <https://www.um.edu.mt/research/ethics>. All participants in data collection, and all assistants in the research reported here, were remunerated for their

work at rates specified by the University of Malta. All participants in data collection gave their informed consent prior to participation.

Acknowledgements

The MASRI Project is funded by the University of Malta Research Excellence Fund Awards. We gratefully acknowledge the support of NVIDIA Corporation with the donation of the Titan Xp GPU used for MASRI research. The audiobooks portion of our pretraining data was generously provided by Merlin Publishers Ltd; special thanks are due to Chris Gruppetta for facilitating the acquisition process. The authors gratefully acknowledge the help of Amanda Muscat, Kirsty Azzopardi and Ayrton Brincat with data collection and transcription for the MASRI-HEADSET corpus. We also thank KPMG Microsoft Business Solutions (formerly CrimsonWing) for providing the TTS system used in our experiments.

References

- Ali A, Bell P, Glass J, Messaoui Y, Mubarak H, Renals S, Zhang Y (2016) The mgb-2 challenge: Arabic multi-dialect broadcast media recognition. In: 2016 IEEE Spoken Language Technology Workshop (SLT), IEEE, pp 279–284
- Baevski A, Zhou H, Mohamed A, Auli M (2020a) wav2vec 2.0: A Framework for Self-Supervised Learning of Speech Representations. arXiv 2006.11477:1–19, URL <http://arxiv.org/abs/2006.11477>, 2006.11477
- Baevski A, Zhou H, Mohamed A, Auli M (2020b) wav2vec 2.0: A framework for self-supervised learning of speech representations. NIPS 2020
- Bender EM, Gebru T, McMillan-Major A, Shmitchell S (2021) On the Dangers of Stochastic Parrots: Can Language Models Be Too Big? In: Proceedings of the fourth ACM Conference on Fairness, Accountability, and Transparency (FAccT’21), Association for Computing Machinery, Online
- Besacier L, Barnard E, Karpov A, Schultz T (2014) Automatic speech recognition for under-resourced languages: A survey. *Speech Communication* 56:85–100
- Billa J (2018) Isi asr system for the low resource speech recognition challenge for indian languages. In: INTERSPEECH, pp 3207–3211
- Borg C, Gatt A (2017) Morphological analysis for the Maltese language: The challenges of a hybrid system. In: Proceedings of the 3rd Arabic Natural Language Processing Workshop (WANLP’17), Association for Computational Linguistics, Valencia, Spain, URL <http://aclweb.org/anthology/W17-1304>, <https://arxiv.org/abs/1703.08701>

- Borg M, Bugeja K, Vella C, Mangion G, Gafà C (2011) Preparation of a free-running text corpus for Maltese concatenative speech synthesis. In: Proceedings of the 3rd International Conference on Maltese Linguistics (GhiLM'11), GhiLM, Valletta
- Borg M, Bugeja K, Vella C, Mangion G, Gafà C (2014) Preparation of a free-running text corpus for Maltese concatenative speech synthesis. Perspectives on Maltese Linguistics, *Studia Typologica* 14:297–318, berlin:Akademie Verlag
- Brincat J (2011) Maltese and other languages: A linguistic history of Malta. Midsea Books, Malta
- Camilleri JJ (2013) A Computational Grammar and Lexicon for Maltese. Master's thesis, Chalmers University of Technology, Gothenburg, Sweden
- Conneau A, Khandelwal K, Goyal N, Chaudhary V, Wenzek G, Guzman F, Grave E, Ott M, Zettlemoyer L, Stoyanov V (2020) Unsupervised Cross-lingual Representation Learning at Scale. In: Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics (ACL'20), pp 8440–8451
- Cucu H, Buzo A, Burileanu C (2014) Unsupervised acoustic model training using multiple seed asr systems. In: Spoken Language Technologies for Under-Resourced Languages
- Cui X, Goel V, Kingsbury B (2015) Data augmentation for deep neural network acoustic modeling. *IEEE/ACM Transactions on Audio, Speech, and Language Processing* 23(9):1469–1477
- Devlin J, Chang MW, Lee K, Toutanova K (2019) BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding. In: Proceedings of NAACL-HLT 2019, Association for Computational Linguistics, Minneapolis, MN, pp 4171–4186, URL <http://arxiv.org/abs/1810.04805>
- Gales MJ, Woodland P, Chan HY, Mrva D, Sinha R, Tranter SE, et al. (2006) Progress in the cu-htk broadcast news transcription system. *IEEE Transactions on Audio, Speech, and Language Processing* 14(5):1513–1525
- Gales MJ, Ragni A, AlDamarki H, Gautier C (2009) Support vector machines for noise robust asr. In: 2009 IEEE Workshop on Automatic Speech Recognition & Understanding, IEEE, pp 205–210
- Gatt A, Čéplö S (2013) Digital corpora and other electronic resources for Maltese. In: Proceedings of the International Conference on Corpus Linguistics, Lancaster, UK: University of Lancaster
- Graves A, Fernández S, Gomez F, Schmidhuber J (2006) Connectionist temporal classification: labelling unsegmented sequence data with recurrent neural networks. In: Proceedings of the 23rd international conference on Machine learning, pp 369–376

- Grézl F, Karafiát M, Veselý K (2014) Adaptation of multilingual stacked bottleneck neural network structure for new language. In: 2014 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP), IEEE, pp 7654–7658
- Habash N, Soudi A, Buckwalter T (2007) On arabic transliteration. In: Arabic computational morphology, Springer, pp 15–22
- Hayashi T, Watanabe S, Zhang Y, Toda T, Hori T, Astudillo R, Takeda K (2018) Back-translation-style data augmentation for end-to-end asr. In: 2018 IEEE Spoken Language Technology Workshop (SLT), IEEE, pp 426–433
- Hori T, Astudillo R, Hayashi T, Zhang Y, Watanabe S, Le Roux J (2019) Cycle-consistency training for end-to-end speech recognition. In: ICASSP 2019-2019 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP), IEEE, pp 6271–6275
- Hu J, Ruder S, Siddhant A, Neubig G, Firat O, Johnson M (2020) XTREME: A Massively Multilingual Multi-task Benchmark for Evaluating Cross-lingual Generalization. arXiv 2003.11080, URL <http://arxiv.org/abs/2003.11080>, 2003.11080
- Huang X, Acero A, Hon HW, Reddy R (2001) Spoken language processing: A guide to theory, algorithm, and system development. Prentice hall PTR
- Inseng D, Motlicek P, Boulard H, Garner PN (2014) Using out-of-language data to improve an under-resourced speech recognizer. *Speech communication* 56:142–151
- Jaitly N, Hinton GE (2013) Vocal tract length perturbation (vtlp) improves speech recognition. In: Proc. ICML Workshop on Deep Learning for Audio, Speech and Language, vol 117
- Jimerson R, Simha K, Ptucha R, Prud’hommeaux E (2018) Improving asr output for endangered language documentation. In: The 6th Intl. Workshop on Spoken Language Technologies for Under-Resourced Languages
- Jurafsky D, Wooters C, Segal J, Stolcke A, Fosler E, Tajchaman G, Morgan N (1995) Using a stochastic context-free grammar as a language model for speech recognition. In: 1995 International Conference on Acoustics, Speech, and Signal Processing, IEEE, vol 1, pp 189–192
- Kanda N, Takeda R, Obuchi Y (2013) Elastic spectral distortion for low resource speech recognition with deep neural networks. In: 2013 IEEE Workshop on Automatic Speech Recognition and Understanding, IEEE, pp 309–314
- Kemp T, Waibel A (1999) Unsupervised training of a speech recognizer: Recent experiments. In: Sixth European Conference on Speech Communication and Technology

- Kingma DP, Ba J (2014) Adam: A method for stochastic optimization. arXiv preprint arXiv:14126980
- Knill KM, Gales MJ, Rath SP, Woodland PC, Zhang C, Zhang SX (2013) Investigation of multilingual deep neural networks for spoken term detection. In: 2013 IEEE Workshop on Automatic Speech Recognition and Understanding, IEEE, pp 138–143
- Kuchaiev O, Li J, Nguyen H, Hrinchuk O, Leary R, Ginsburg B, Krizan S, Beliaev S, Lavrukhin V, Cook J, et al. (2019) Nemo: a toolkit for building ai applications using neural modules. arXiv preprint arXiv:190909577
- Lamel L, Gauvain JL, Adda G (2002) Lightly supervised and unsupervised acoustic model training. *Computer Speech & Language* 16(1):115–129
- Li J, Lavrukhin V, Ginsburg B, Leary R, Kuchaiev O, Cohen JM, Nguyen H, Gadde RT (2019) Jasper: An end-to-end convolutional neural acoustic model. arXiv preprint arXiv:190403288
- Liang Y, Duan N, Gong Y, Wu N, Guo F, Qi W, Gong M, Shou L, Jian D, Cao G, Fan X, Zhang R, Agrawal R, Cui E, Wei S, Bharti T, Qiao Y, Chen JH, Wu W, Liu S, Yang F, Campos D, Majumder R, Zhou M (2020) XGluE: A new benchmark dataset for cross-lingual pretraining, understanding and generation. arXiv 2004.01401
- Löf J, Gollan C, Ney H (2009) Cross-language bootstrapping for unsupervised acoustic model training: Rapid development of a polish speech recognition system. In: Tenth Annual Conference of the International Speech Communication Association
- Mena C (2019) The ciempiess proper-names pronouncing dictionary. In: Corpus presented at OpenCor 2019 Conference, Guanajuato City, Mexico.[Available online at <https://opencor.gitlab.io/corpora-list/>]
- Mena CDH, Gatt A, DeMarco A, Borg C, van der Plas L, Muscat A, Padovani I (2020) Masri-headset: A maltese corpus for speech recognition. In: Proceedings of The 12th Language Resources and Evaluation Conference, pp 6381–6388
- Micallef P (1997) A text to speech synthesis system for Maltese. PhD thesis, University of Surrey
- Panayotov V, Chen G, Povey D, Khudanpur S (2015) Librispeech: an asr corpus based on public domain audio books. In: 2015 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP), IEEE, pp 5206–5210
- Park DS, Chan W, Zhang Y, Chiu CC, Zoph B, Cubuk ED, Le QV (2019a) SpecAugment: A simple data augmentation method for automatic speech recognition. arXiv preprint arXiv:190408779

- Park DS, Chan W, Zhang Y, Chiu CC, Zoph B, Cubuk ED, Le QV (2019b) SpecAugment: A Simple Data Augmentation Method for Automatic Speech Recognition. In: Proceedings of Interspeech 2019, ISCA, Graz, Austria, pp 2613–2617, DOI 10.21437/Interspeech.2019-2680, URL <http://arxiv.org/abs/1904.08779>, 1904.08779
- Pratap V, Sriram A, Tomasello P, Hannun A, Liptchinsky V, Synnaeve G, Collobert R (2020) Massively multilingual asr: 50 languages, 1 model, 1 billion parameters. CoRR abs/2007.03001, URL <https://arxiv.org/abs/2007.03001>
- Qian Y, Yu K, Liu J (2013) Combination of data borrowing strategies for low-resource lvsr. In: 2013 IEEE Workshop on Automatic Speech Recognition and Understanding, IEEE, pp 404–409
- Ragni A, Knill K, Rath SP, Gales M (2014) Data augmentation for low resource languages. In: INTERSPEECH, pp 810–814
- Ramirez JM, Montalvo A, Calvo JR (2019) A survey of the effects of data augmentation for automatic speech recognition systems. In: Iberoamerican Congress on Pattern Recognition, Springer, pp 669–678
- Ravishankar V, Tyers F, Gatt A (2017) A Morphological analyser for Maltese. *Procedia Computer Science (Special Issue on Arabic Computational Linguistics)* 117:175–182, DOI 10.1016/j.procs.2017.10.107, URL <https://www.sciencedirect.com/science/article/pii/S1877050917321646>
- Rogers A (2021) Changing the World by Changing the Data. arXiv 2105.13947, 2105.13947
- Rosenberg A, Zhang Y, Ramabhadran B, Jia Y, Moreno P, Wu Y, Wu Z (2019) Speech recognition with augmented synthesized speech. In: 2019 IEEE Automatic Speech Recognition and Understanding Workshop (ASRU), IEEE, pp 996–1002
- Rossenbach N, Zeyer A, Schlüter R, Ney H (2020) Generating synthetic audio data for attention-based speech recognition systems. In: ICASSP 2020-2020 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP), IEEE, pp 7069–7073
- Schneider S, Baevski A, Collobert R, Auli M (2019) wav2vec: Unsupervised Pre-training for Speech Recognition. In: Proceedings of Interspeech 2019, Graz, Austria, pp 1–9, URL <http://arxiv.org/abs/1904.05862>
- Simard PY, Steinkraus D, Platt JC, et al. (2003) Best practices for convolutional neural networks applied to visual document analysis. In: *Icdar*, vol 3
- Thai B, Jimerson R, Arcoraci D, Prud’hommeaux E, Ptucha R (2019) Synthetic data augmentation for improving low-resource asr. In: 2019 IEEE Western New York Image and Signal Processing Workshop (WNYISPW), IEEE, pp 1–9

- Tiedemann J, van der Plas L (2016) Bootstrapping a Dependency Parser for Maltese - A Real-World Test Case. In: Weiling M, Kroon M, Noord GV (eds) From Semantics to Dialectometry : Festschrift in honor of John Nerbonne, College Publications, London
- Tjandra A, Sakti S, Nakamura S (2017) Listening while speaking: Speech chain by deep learning. In: 2017 IEEE Automatic Speech Recognition and Understanding Workshop (ASRU), IEEE, pp 301–308
- Tüske Z, Pinto J, Willett D, Schlüter R (2013) Investigation on cross-and multilingual mlp features under matched and mismatched acoustical conditions. In: 2013 IEEE International Conference on Acoustics, Speech and Signal Processing, IEEE, pp 7349–7353
- Tüske Z, Golik P, Nolden D, Schlüter R, Ney H (2014) Data augmentation, feature combination, and multilingual neural networks to improve asr and kws performance for low-resource languages. In: Fifteenth Annual Conference of the International Speech Communication Association
- Vachhani B, Bhat C, Kopparapu SK (2018) Data augmentation using healthy speech for dysarthric speech recognition. In: Interspeech, pp 471–475
- Vaswani A, Shazeer N, Parmar N, Uszkoreit J, Jones L, Gomez AN, Kaiser L (2017) Attention Is All You Need. In: Proceedings of the 31st Conference on Neural Information Processing Systems (NIPS’17), Long Beach, CA, [arXiv:1706.03762v5](https://arxiv.org/abs/1706.03762v5)
- Vu NT, Schultz T (2013) Multilingual multilayer perceptron for rapid language adaptation between and across language families. In: Interspeech, pp 515–519
- Wang C, Pino J, Gu J (2020) Improving cross-lingual transfer learning for end-to-end speech recognition with speech translation. [2006.05474](https://arxiv.org/abs/2006.05474)
- Weide RL (1998) The CMU pronouncing dictionary. URL <http://svn.code.sf.net/p/cmuspinx/code/trunk/cmudict/>
- Wenzek G, Lachaux MA, Conneau A, Chaudhary V, Guzmán F, Joulin A, Grave E (2020) CCNet: Extracting high quality monolingual datasets from web crawl data. In: Proceedings of the 12th International Conference on Language Resources and Evaluation (LREC’20), European Language Resources Association (ELRA), Marseille, France, pp 4003–4012, [1911.00359](https://arxiv.org/abs/1911.00359)
- Wessel F, Ney H (2004) Unsupervised training of acoustic models for large vocabulary continuous speech recognition. *IEEE Transactions on Speech and Audio Processing* 13(1):23–31
- Yilmaz E, Heuvel Hvd, van Leeuwen DA (2018) Acoustic and textual data augmentation for improved asr of code-switching speech. [arXiv preprint arXiv:1807.10945](https://arxiv.org/abs/1807.10945)

Zammit A (2018) A dependency parser for the Maltese language using deep neural networks. Master's thesis, Department of Artificial Intelligence, University of Malta, Msida, Malta

Zavaliagkos G, Colthurst T (1998) Utilizing untranscribed training data to improve performance. In: DARPA Broadcast News Transcription and Understanding Workshop, Landsdowne