

Dear Reviewer,

Thank you for considering our manuscript for publication and for providing constructive feedback. Hereunder you will find our detailed replies to all your comments. The changes are highlighted in the output of the latexdiff file attached to this cover letter.

Issue 1: *The authors presented the development of a keyword spotting system detecting regions of interest in Polish call centre recordings. The paper is interesting, and it is related to an important field. However, the language must be proofread by a native speaker.*

Answer Thank you for this comment. We have carefully read the text and corrected the manuscript. The changes are depicted in attached latex-diff output.

Issue 2: *1. The abstract must be improved with a stating background sentence;*

Answer We have expanded the abstract with an introductory sentence stating the background. The sentence “Models for keyword spotting in continuous recordings can significantly improve the experience of navigating vast libraries of audio recordings.” was added at beginning of the abstract.

Issue 3: *2. The introduction is not clearly presenting the purpose and scope of the study;*

Answer We have added the research hypothesis as a separate subsection (1.4) to clearly present the purpose and scope of the study.

Issue 4: *3. The material and methods section needs to define the statistical analysis to be performed for the presentation of the results;*

Answer We have added Section 2.4 to our paper, which defines the statistical analysis performed on the results achieved from our experiments.

Issue 5: *4. The results are interesting and compared;*

Answer Thank you for your appreciation of our work.

Issue 6: 5. *The discussion must be improved, and a brief conclusions' section must be added.*

Answer Thank you for highlighting this issue. We agree, we have added a brief conclusion and modified the discussion section.

Sincerely yours,

Lukasz Lepak, Kacper Radzikowski, Robert Nowak, Karol Piczak

Generalisation Gap of Keyword Spotters in a Cross-Speaker Low-Resource Scenario

Łukasz Lepak ^{1,*} , Kacper Radzikowski ^{1,2}, Robert Nowak ¹ , and Karol J. Piczak ^{3,*} 

¹ Warsaw University of Technology, Faculty of Electronics and Information Technology, Institute of Computer Science, Warsaw, Poland

² Waseda University, Graduate School of Information, Production and Systems, Tokyo, Japan

³ Jagiellonian University, Institute of Computer Science and Computational Mathematics, Kraków, Poland

* Correspondence: lukasz.lepak.dokt@pw.edu.pl, karol.piczak@uj.edu.pl

Abstract: Models for keyword spotting in continuous recordings can significantly improve the experience of navigating vast libraries of audio recordings. In this paper, we describe the development of such a keyword spotting system detecting regions of interest in Polish call centre conversations. Unfortunately, in spite of recent advancements in automatic speech recognition systems, human-level transcription accuracy reported on English benchmarks does not reflect the performance achievable in low-resource languages, such as Polish. Therefore, in this work we shift our focus from complete speech-to-text conversion to acoustic similarity matching in the hope of reducing the demand for data annotation. As our primary approach, we evaluate Siamese and prototypical neural networks trained on several datasets of English and Polish recordings. While we obtain usable results in English, our models' performance remains unsatisfactory when applied to Polish speech, both after mono- and cross-lingual training. This performance gap shows that generalisation with limited training resources is a significant obstacle for actual deployments in low-resource languages. As a potential countermeasure, we implement a detector using audio embeddings generated with a generic pre-trained model provided by Google. It has a much more favourable profile when applied in a cross-lingual setup to detect Polish audio patterns. Nevertheless, despite these promising results, its performance on out-of-distribution data is still far from stellar. It would indicate that in spite of the richness of internal representations created by more generic models, such speech embeddings are not entirely malleable to cross-language transfer.

Keywords: keyword spotting; speech embedding; Siamese networks; Automatic Speech Recognition

Citation: Lepak, Ł.; Radzikowski, K.; Nowak, R.; Piczak, K. J. Generalisation Gap of Keyword Spotters in a Cross-Speaker Low-Resource Scenario. *Sensors* **2021**, *1*, 0. <https://doi.org/>

Received:

Accepted:

Published:

Publisher's Note: MDPI stays neutral with regard to jurisdictional claims in published maps and institutional affiliations.

Copyright: © 2021 by the authors. Submitted to *Sensors* for possible open access publication under the terms and conditions of the Creative Commons Attribution (CC BY) license (<https://creativecommons.org/licenses/by/4.0/>).

1. Introduction

1.1. Keyword spotting

The goal of a keyword spotter is to detect words of interest in continuous audio recordings. These recordings can be provided either as prerecorded files of considerable length (offline processing) or real-time streaming data (online processing). Nowadays, keyword spotting is ~~associated more often~~ more often associated with the latter scenario. The purpose of such a real-time detector is to catch utterances of a specific wake word provided by the user and activate a fully functional conversation with a personal digital assistant.

However, offline usage of keyword spotting can be of great help when navigating vast libraries of audio recordings, pinpointing short regions of interest without the need to listen through the entire conversation. This feature is especially applicable to the mundane task of reviewing lengthy call centre recordings.

35 1.2. Paper overview

36 In this paper, we describe a case study of developing such a proof-of-concept solution
37 for spotting keywords in call centre recordings of a Polish bank. The system's goal
38 was to provide rapid localisation of predefined words of interest when reviewing call
39 centre conversations with the client, which would help bank employees process clients'
40 complaints more efficiently.

41 1.3. Approaches to keyword spotting

42 In general, there are two main approaches that we can utilise for creating keyword
43 spotters: speech-to-text conversion or audio similarity matching. Both of them have
44 their strengths and weaknesses.

45 1.3.1. Speech-to-text conversion

46 The first approach consists of a full-scale *automatic speech recognition* (ASR) system
47 transforming incoming audio streams into a textual representation. In this way, searching
48 for regions of interest is converted to a trivial problem of finding words in a text stream.
49 This approach has several advantages.

50 First of all, the generated output has a much more desirable form. Instead of only
51 locating particular keywords, the system provides the user with a complete transcription
52 of the reviewed conversation. Such an output form can then be utilised for numerous
53 additional purposes.

54 A textual representation also helps when extending the dictionary of the system.
55 As long as the speech recogniser generates a high-quality transcription, localising com-
56 pletely new keywords is trivial and can be performed ex-post.

57 Typical ASR systems have another advantage over audio content similarity match-
58 ing, as they combine acoustic and language modelling information. Processing whole
59 sentences enables the model to refine its predictions and eliminate implausible responses
60 based on the probability estimate of ~~a~~-the generated word sequence. Language mod-
61 elling can thus help lower the *word error rate* (WER) in comparison to the processing of
62 words in isolation.

63 Unfortunately, a-the ASR-based approach has several drawbacks. For instance,
64 any errors introduced during the transcription phase are final and cannot be fixed in
65 post hoc analysis by adjusting the detection threshold. Another considerable drawback
66 of training a complete ASR system is the data intensity of this process. Various esti-
67 mates place the requirements on training data availability at 2,000 hours of transcribed
68 recordings for a viable system [1] and as much as 10,000 hours for production-quality
69 results [2].

70 While the high cost of model training on such a scale [2] effectively prohibits any
71 further experimentation with the baseline systems, it is not the most problematic factor.
72 Unfortunately, in mid-2021, there were no publicly available resources of transcribed
73 Polish speech that would remotely meet the required scope of such a project. Conse-
74 quently, using a complete ASR system is still not a viable option when processing Polish
75 speech, despite the numerous advantages of this approach.

76 1.3.2. Audio similarity matching

77 Due to the aforementioned problems with data availability, we have opted for the second
78 option, i.e. developing a keyword detection system based on the perceptual similarity of
79 acoustic fragments. Figure 1 presents a schematic depiction of this approach. The main
80 principle of such a system lies in sequentially processing small fragments of an input
81 recording. Each extracted fragment is then compared to a set of predefined exemplary
82 patterns representing specific keywords. If a significant similarity is found, the system
83 reports detection of a given keyword at a particular point in time.

84 As our similarity detector, we employ a Siamese neural network model. It is a
85 metric-learning model that has been widely used in similarity ranking problems and is




Figure 1. Keyword detection based on audio similarity matching. This schematic depicts a typical pipeline. An input recording is processed sequentially with a short rolling window (e.g. 800 ms). The system compares each extracted fragment to a set of exemplary patterns predefined for each keyword. The model then outputs a decision based on the distance between the fragment and each pattern. If the distance is sufficiently small, a keyword occurrence is emitted for a given timestamp.

specifically tailored to problems where examples of new classes are scarce, i.e. to *few-shot learning* [3]. We also compare Siamese models with prototypical networks [4], another solution popular in the domain of few-shot learning.

The main advantage of such similarity-based approaches lies in more moderate data requirements, both in terms of quantity and quality of annotation. Instead of using detailed transcripts that are costly to procure, we can train detectors by providing pairs of word utterances with a binary “match”/“no match” label, a much more convenient data acquisition setup.

Additionally, focusing on acoustic similarity allows for recognition of proper names that language modelling might filter out due to their rareness. The patterns might even represent data other than voice, as shown by Wang et al. [5]. By computing similarity in the domain much closer to the raw data representation, it ~~is~~ would be also possible to search for fragments based on information usually lost during conversion, such as emotion, the rhythm of speech, and tone of voice. However, in this work we focus solely on phonetic matching of keywords.

1.4. Research hypothesis

Our research hypothesis is that keyword spotting models that employ acoustic similarity matching should allow for keyword retrieval from call centre conversations in a low-resource language setting.

1.5. Research outcome

Our experiments show that with a relatively diverse pool of examples used to train the acoustic embedding model, the performance of acoustic similarity matching is satisfactory when evaluated on synthetic benchmarks. Unfortunately, while such training resources are obtainable for the English language, we were unsuccessful in creating an equivalently robust model for Polish speech, ~~as the project’s resources for annotating additional Polish recordings were limited.~~ This performance gap underlines how problematic it is to create keyword spotting systems for low-resource languages.

However, as the pattern matching in such systems relies solely on acoustic similarity metrics, it should be, at least partially, language-agnostic ~~—~~ as evidenced by the results of Wang et al. [5]. Therefore, a plausible solution would be to train a model on a high-resource language and use it directly on the low-resource language, hoping that such perceptual matching will be sufficiently accurate, despite the language difference. To this end, we perform several experiments by initially training models on various standard English datasets and using Polish patterns in further detection steps. Unfortunately, such models exhibit weak generalisation when switching between languages, even when augmented with additional Polish data during training. This problem could indicate

that the internal representation created by these models is highly specialised and attuned to the particular dataset.

As a way to potentially overcome this issue, we also evaluate an approach that utilises more generic speech embeddings created with a model provided by Google [6] that was pre-trained on 200 million English audio clips from YouTube. Our analysis shows that this embedder creates a more nuanced internal representation of the keyword classes, capturing additional variability factors present in the dataset. It is also better at detecting Polish patterns, despite being trained on non-Polish speech. Although the handicap of an enormous training dataset definitely helps in the few-shot cross-language setup, the results are still far from stellar. A final evaluation on out-of-distribution target recordings shows viability only for a very limited subset of potential Polish keywords.

1.6. Related works

1.6.1. Automatic speech recognition

The history of automatic speech recognition dates as far back as 1952 when Bell Labs released *Audrey*, a software that could recognise simple numbers. The next breakthrough came in the mid-1970s when researchers began to use *hidden Markov models* (HMMs). This approach has led to the development of many popular ASR systems, e.g. *SPHINX* [7]. Nowadays, with the introduction of more extensive datasets, faster computers, and GPU processing, a new ASR paradigm, end-to-end speech recognition with deep learning, has mostly supplanted other methods.

For instance, in 2014, Hannun et al. [8] developed a recurrent neural network system called *Deep Speech*, which outperformed previously published results on the widely studied Switchboard Hub5'00, achieving 16.0% error on the full test set. This concept was refined in further instantiations of the model [2,9].

In 2016, Oord et al. [10] introduced *WaveNet*, a probabilistic and autoregressive generative model designed for speech synthesis. However, it can also be used discriminatively for speech recognition. The authors report a *phone error rate* (PER) of 18.8% on the TIMIT dataset.

A year later, in 2017, Microsoft combined convolutional neural networks with bi-directional-LSTM layers in a conversational speech recognition system [11]. This setup yielded a 5.1% word error rate on the 2000 Switchboard evaluation set.

Nowadays, fully convolutional end-to-end acoustic models, like *Jasper* [12] or *QuartzNet* [13], report a WER below 3.0% on LibriSpeech *test-clean* evaluation when combined with transformer-based language models.

An even more recent approach, *Conformer* [14], proposes a convolution-augmented transformer model for speech recognition, achieving state-of-the-art accuracies on the LibriSpeech benchmarks.

While these results on English benchmarks are remarkable, the problem with current approaches lies in their data intensity. The referenced models were trained on datasets containing hundreds or even thousands of hours of transcribed English speech. Simple adoption of these techniques to Polish ASR systems is thus impossible.

1.6.1. Keyword spotting

Since the advent of deep learning models in speech processing, researchers have also proposed numerous variants of such approaches for keyword spotting. Many of these models have been smaller-footprint versions of typical video and audio processing solutions. Unfortunately, most initial evaluations of models employing convolutional neural networks [15] and recurrent neural networks [16] used proprietary recordings for training and validation. More recently, by introducing the Speech Commands dataset [17], Warden established a popular public framework for benchmarking keyword spotting systems.

The most prevalent evaluation protocol for Speech Commands v1 employs an 80:10:10 split between training, validation, and testing data, selecting 10 of the 30 words

in the dataset as target commands. The remaining 20 words create a single negative category (*unknown word*), while additional ambient recordings serve as a silent class, resulting in 12 classification options. Standard training data augmentation consists of random shifts and noise injection.

The baseline method provided by the author [17] attains a top-1 classification accuracy of 85.4% with a simple convolutional architecture [15] operating on 40-dimensional log-mel filter bank features. Slightly better performance is achievable with convolutional neural networks using 40-dimensional MFCC features [18]. By combining a broad hyperparameter sweep with depthwise separable convolutions, Zhang et al. were able to bring the accuracy up to 95% even with resource-constrained architectures [19].

In a similar vein, de Andrade et al. [20] applied a recurrent neural attention model to the task of speech command recognition achieving an accuracy of 94.1% on Speech Commands v1 with 80-band mel spectrograms. Zeng and Xiao [21] reported similar levels for baseline LSTM/CNN architectures, which they were able to improve up to 96.2% through the application of BiLSTM layers to a DenseNet model. What is even more noteworthy, these kinds of results are attainable even with low-footprint models suitable for speech recognition at the edge with less than 100,000 parameters (96.3% with *EdgeSpeechNets* [22], 96.1% reported for CNNs with temporal convolution over spectrograms [23], 96.4% with parametrized sinc-convolutions [24], and 97.5% with a bigger version of *MatchboxNet* [25]). When the number of parameters is not an issue, current state-of-the-art approaches can achieve even higher accuracy levels with *MoEx* (*moment exchange*) feature regularisation on wide ResNet models (98.0% with WRN-28-10 consisting of 36.5 million parameters) [26].

Considering these stellar results, it might seem that keyword spotting is a solved problem. However, the accuracy of classifying isolated excerpts from the Speech Commands dataset might not translate to actual performance when working on continuous audio streams. Real scenarios involve unknown word boundaries, more variance in noisy conditions, and a highly imbalanced class distribution between trigger words and background information. This fact was brought up even in the original Speech Commands paper [17]. It has shown that a baseline model, with 88.2% accuracy of classifying isolated words from version 2 of the dataset, detects keywords with a 46.0% true positive rate when applied to a synthetically generated one-hour audio stream. While some papers were able to report much more promising results in this area of streaming keyword spotting, they were either limited to a detection of a single wake-word [27] or used proprietary datasets of unprecedented scope, e.g. 1 million training utterances of “OK/Hey Google” [28,29] and recordings of several hundred thousand subjects interacting with *Alexa* devices [30].

1.6.2. Speech recognition in low-resource settings

Over the years, many successful approaches to English ASR [2,8–14] have been proposed in the form of various deep learning models. Unfortunately, these models were trained on datasets containing hundreds or even thousands of hours of transcribed English speech. Simple adoption of these techniques to Polish ASR systems is thus impossible.

There were some previous attempts to create speech recognition systems ~~for~~ specifically for the Polish language. However, they are mostly several years old and based on non-neural network methods. Some examples include systems based on hidden Markov models [31], k-nearest neighbours [32], and speech n-grams [33]. Results achieved with these methods are not satisfactory, therefore they were not considered further in our project.

A recent development in global ASR research is the introduction of models trained on large amounts of unlabelled data. These methods swap the effort in data collection for computational time. By analysing vast collections of raw recordings, they can create more robust and powerful speech representations that can also facilitate knowledge transfer to low-resource languages.

227 An example of this approach is the *wav2vec 2.0* framework introduced by Face-
228 book [34]. It uses a contrastive self-supervised learning procedure to create latent speech
229 representations. After pre-training on unlabelled data, the *Connectionist Temporal Classifi-*
230 *cation* (CTC) loss is used to refine the model further on downstream tasks with labelled
231 data. The main advantage of this approach is the tremendous drop in the requirements
232 for data annotation. Wav2vec 2.0 can achieve a WER of 2.9% on LibriSpeech *test-clean*
233 with as little as one hour of labelled data. For such an outcome to occur, a high price
234 has to be paid when training the model. The most performant variant consisting of 24
235 transformer blocks needs the equivalent of roughly 2 GPU years of training on a V100 to
236 process the 53,200 hours of LibriVox recordings.

237 An extension of this concept has led to the introduction of *XLSR-53* [35], a large
238 model pre-trained on 56,000 hours of recordings in 53 languages. It has shown re-
239 markable reductions in phone error rates when evaluated on various languages of the
240 Common Voice dataset, with as little as one hour of labelled data required for fine-tuning
241 the final model. Based on these results, XLSR-53 seems to be the most promising ap-
242 proach for speech recognition in low-resource languages to date. Due to its complexity,
243 recreating such a model requires a considerable implementation effort and access to a
244 GPU cluster. Unfortunately, Facebook released the pre-trained model only near the very
245 end of 2020, after the conclusion of our project's development. Nevertheless, for future
246 endeavours in Polish ASR, it is an interesting possibility to explore.

247 However, in our work, we have opted for the help of a different pre-trained solution.
248 In 2020, a team from Google Research released a speech embedding model trained on
249 200 million 2-second audio clips from YouTube [6]. This convolutional model aims to
250 convert an audio stream into a stream of 96-dimensional features vectors. Each generated
251 embedding vector encodes speech content in windows of approximately 800 ms, every
252 80 ms. To ensure the reusability of these embeddings, the model was pre-trained on
253 an arbitrary set of 5,000 keywords split randomly over 125 keyword spotters sharing
254 the embedder backbone. After 40 GPU days of training, the embedding backbone
255 has been released on TensorFlow Hub for reuse. The authors have shown that these
256 generic embeddings could limit the amount of data needed for creating a robust Speech
257 Commands keyword classifier. They were also successful in partially substituting real
258 training data with synthesised speech. Therefore, we have chosen this approach as a
259 potential solution for our data availability problem.

260 1.6.3. Voice datasets

261 Various voice datasets suitable for speech recognition and keyword spotting tasks are
262 used for training neural network models. While many datasets are strictly proprietary, a
263 handful of them are distributed with more permissive licenses.

264 One of the most popular in this area is the already mentioned Speech Commands
265 by Google [36]. Its second version consists of 35 different words in English, spoken by
266 2,618 different speakers, totalling 105,829 utterances. The Speech Commands dataset is
267 available for download from the TensorFlow Datasets catalogue.

268 Common Voice [37], an initiative of the Mozilla Foundation, is on the way to become
269 the most significant publicly available dataset of voice recordings. It contains various
270 sentences read by volunteers in many languages and is freely distributed under the
271 Creative Commons Zero license. Everyone can contribute to the development of this
272 dataset by recording their voice through a browser or a mobile phone. The quantity of
273 sentences varies between languages, with English having the most significant share.

274 Initially, the dataset incorporated new languages when a sufficient number of
275 recordings was amassed. Hence, the Polish version was publicly distributed for the
276 first time only in Common Voice Corpus 5.1, in the later stages of our project's timeline.
277 Nevertheless, the total length of validated recordings, reaching approximately 100 hours,
278 would be insufficient for training a complete ASR system either way. In mid-2021,

more than 70 languages are publicly available. However, the distribution of recordings between languages remains very uneven.

Common Voice Single Word Target Segment was a spin-off dataset published alongside Corpus 5.1, resembling an extension of the Speech Commands concept across multiple languages. It contains recordings of spoken digits, as well as the words *yes*, *no*, *hey* and *Firefox*. Number of speakers varies greatly by language.

Another popular voice dataset is LibriSpeech [38]. It is a collection of around 1,000 hours of sentences in English. Recordings are annotated on a sentence rather than word level, making it more suitable for speech recognition tasks than keyword spotting.

The Spoken Wikipedia Corpus [39] is a collection of Wikipedia articles read by volunteers. It contains 182 hours of word-aligned transcriptions in English, 249 hours in German, and 79 hours in Dutch. Based on this dataset, Wang et al. [5] have successfully employed prototypical networks for few-shot keyword detection, also in a cross-language evaluation. However, their problem formulation assumed that, in the detection phase, the keyword spotter is provided with exemplary patterns belonging to the same speaker as the speaker in the analysed recording. This availability of very similar patterns is in contrast to our project's assumptions. We expect keyword searching to function equally well in a cross-speaker regime, generalising to completely new speakers not encountered in the training set.

When deploying speech recognition models, some problematic aspects concerning voice datasets have to be considered. For instance, although the datasets described in this section use permissive licenses, many other datasets may come with limits on their research and commercial usage.

Another very significant problem is language availability. Numerous datasets are available in English. However, finding a good dataset in other languages is often not easy, as they tend to have much less data or insufficient quality. This deficiency makes speech recognition and keyword spotting tasks outside English very difficult. Although several authors, apart from the Common Voice project [37], have prepared Polish voice corpora, they are either limited in recording length [40–42], not easily downloadable for offline use [43,44] or proprietary [45]. ~~An interesting dataset that we encountered only at the time of writing~~ A potentially interesting dataset is SNUV (Spelling and Numbers Voice database) [46] that contains 220 hours of Polish speakers reading numbers and spelling words. However, due to the lack of complete word utterances (only spelling), its usability in training keyword spotters might be limited.

2. Materials and Methods

During the development of our keyword spotting system, we have conducted several experiments with various training and evaluation datasets, both in English and Polish. All the experiments in keyword detection adhere to the same pattern.

First, we prepare the datasets by preprocessing the recordings into mel spectrograms and dividing the data into three parts, described in section 2.1:

- *training data* - single-word utterances used for training the acoustic similarity model,
- *search patterns* - keyword examples used as templates for comparison with fragments of the analysed evaluation recording,
- *evaluation recordings* - longer recordings used to assess how well the model detects keywords.

Following this step, we train the acoustic similarity model on single-word utterances. In our work, we consider two different types of models that we train from scratch: a Siamese convolutional neural network (*Siamese*) and a prototypical network (*Prototypical*). We also use a pre-trained speech embedder (*Google*) for comparison. Section 2.2 describes in more detail the architectures of these models and the training procedure.

After training, the similarity models work by calculating distances between a provided recording fragment and predefined keyword patterns. For each time step, this mode of operation provides distinct similarity values for each of the analysed keywords.

Therefore, as the last step, we use a detector to aggregate all this information into actual detection decisions based on the configured similarity threshold level. In this phase, we also introduce some temporal smoothing and filtering to make the predictions more robust. We provide more detailed detector configurations in section 2.3.

2.1. Datasets and data preprocessing

The main goal of our system was to detect 22 pre-selected keywords in Polish call centre recordings provided by our industrial partner. These recordings came as an unlabelled collection of couple hundred conversations of differing lengths (from minutes to more than an hour) registered with standard call centre equipment at 8 kHz. Unfortunately, privacy concerns with this dataset resulted in stringent local access policies that proved to be problematic in combination with the ongoing pandemic restrictions. Therefore, we could effectively use this dataset only to evaluate the final system, with potentially significant domain shift.

In this less than ideal scenario, we had to resort to several other datasets, with different characteristics, for training and interim validation. Some of the datasets contain real audio recordings, while others are generated as synthetic mixtures. The original sampling rates of the datasets are varied, but we initially downsample all data to 8 kHz, which is the sample rate of the target recordings. We provide specific parameters of spectrogram preprocessing for each of the model types in Table 1.

Table 1. Parameters of mel spectrogram preprocessing for each model type.

Model type	Siamese	Prototypical	Google
Sample rate	8 kHz	8 kHz	16 kHz
Segment duration	0.8 s	0.8 s	0.775 s
Mel bands	40	40	32
FFT window	512	512	400
Hop length	160	160	160
Centring	false	false	false

It is worth noting that the Google model has its processing sample rate set to 16 kHz. This setting aims to match the original training configuration [6]. However, as all our source recordings are initially downsampled to 8 kHz, the available information is the same for all models. Moreover, this discrepancy is also negligible as, effectively, the pre-trained speech embedder only uses log-mel features up to 3.8 kHz [6].

2.1.1. Training data

This section provides short descriptions of the training datasets. We use training folds for creating the similarity model and validation data for interim monitoring. Test folds either serve as the basis for evaluation mixtures or are discarded. We initially downsample all recordings to 8 kHz. Where needed, we also refactor the datasets for the purpose of prototypical training by further splitting the data into *query* and *support* subsets [by selecting speakers with a 75:25 ratio](#). Table 2 presents the aggregated statistics for the datasets.

Speech Commands (SC 1, SC 2, ΔSC)

Speech Commands datasets [17] are defined as follows: SC 1 contains words from 30 classes, SC 2 is a superset of SC 1 extended to 35 classes, while ΔSC contains words from the five classes that are present exclusively in SC 2. All recordings contain an utterance of a single word. We follow the original structure of the dataset by using the pre-defined data splits.

Table 2. Statistics of the training datasets.

Dataset	Language	Samples	Classes	Speakers	Acronym
Speech Commands v1		64,727	30	1,881	 SC 1
Speech Commands v2		105,829	35	2,618	 SC 2
Speech Commands “delta”		7,967	5	663	 ΔSC
Common Voice Single Word		26,070	14	3,519	 CV
Common Voice Single Word		898	14	86	 CV
Spoken Wikipedia Corpus		429,354	4,657	372	 SWC
Warsaw University of Technology		1,058	36	29	 WUT
Text-to-Speech		39,809	5,687	7	 TTS

Common Voice Single Word (CV, CV)

Common Voice Single Word datasets [37] are available both in English and Polish. Each recording contains a single word. We prepare the dataset by splitting it into *train*, *validation* and *test* folds with a 80:10:10 ratio.

Spoken Wikipedia Corpus (SWC)

Spoken Wikipedia Corpus [39] contains English Wikipedia articles read by volunteers. Based on the provided annotations, we extract single word fragments that fit our constraints: they have a sample duration between 200 ms and 1,000 ms, a keyword length of at least three characters, at least five unique speakers for a given keyword, and at least ten samples in total. We limit the number of samples per single keyword/speaker combination to five occurrences. We then split the created collection into 80:20 training and validation folds.

Warsaw University of Technology (WUT)

Warsaw University of Technology dataset is a small collection of single word utterances based on our list of target Polish keywords extended with 14 classes with high acoustic similarity. The word examples are recorded with personal devices by employees of the Artificial Intelligence Division at the Warsaw University of Technology and the personnel of mBank SA. The dataset is used exclusively for training purposes.

Text-to-Speech (TTS)

Text-to-Speech keywords dataset contains Polish spoken words generated with Azure and Google text-to-speech services. We use two forms of this dataset. For pattern matching, we limit the keywords to the 22 classes of the target evaluation. For training purposes, we extend the vocabulary with additional words found while scraping the bank’s website. This process results in a total of 5,687 classes generated with seven TTS voices.

Combined datasets (ALL, ALL)

Combined datasets, as the name suggests, are composed from the datasets described earlier. For English language, we merge Speech Commands v2 ( SC 2), English Common Voice Single Word ( CV) and Spoken Wikipedia Corpus ( SWC). For Polish language, the combined dataset consists of the Polish Common Voice Single Word ( CV), Warsaw University of Technology ( WUT) and Text-to-Speech recordings ( TTS).

2.1.2. Search patterns

Herein, we briefly describe the collections of recordings that we use in our experiments as keyword templates in the matching process. For each keyword we use 10 recordings as templates.

407 **Speech Commands (🇺🇸 SC1, 🇵🇸 ΔSC)**

408 Speech Commands search patterns are generated directly from the respective SC1 and
 409 ΔSC datasets. For each keyword, we select ten random examples from the datasets'
 410 merged training and validation folds. The datasets consist of 30 and 5 classes, respec-
 411 tively.

412 **Target keywords (🇵🇸 KW)**

413 We create these search patterns based on the Warsaw University of Technology and
 414 Text-to-Speech recordings. The dataset contains ten examples for each of the 22 target
 415 keywords, although only 19 keywords are actually present in the evaluation recordings.

416 **2.1.3. Evaluation recordings**

417 This section describes the datasets we use for evaluation purposes.

418 **Synthetic mixtures (🇺🇸 SC1^{Mix}, 🇵🇸 ΔSC^{Mix}, 🇵🇸 KW^{Mix})**

419 We create synthetic evaluation recordings by combining the utterances from the Speech
 420 Commands datasets (🇺🇸 SC1 and 🇵🇸 ΔSC) as a single continuous audio stream. The
 421 utterances come from the respective test folds. In each case, we place 20 keyword
 422 samples per recording with random delays between each occurrence. We generate 50
 423 random evaluation mixtures in this way. For the Polish language, we generate these
 424 mixtures based on 1,136 hand-annotated keyword examples from YouTube audio clips.
 425 All these synthetic mixtures contain only keywords detected in a given scenario, and
 426 are merged into continuous audio stream by us.

427 **Speech Commands overlay on VoxCeleb (semi-synthetic) (🇺🇸 SC1^{Vox})**

428 This evaluation dataset consists of Speech Commands keywords overlaid on various
 429 backgrounds in the form of VoxCeleb [47] conversations. For each evaluation recording,
 430 we use a single keyword occurrence and two conversation fragments with a total length
 431 of approximately 15 seconds. We generate 20 recordings for each keyword.

432 **Target keywords in YouTube recordings (authentic) (🇵🇸 KW^{Real})**

433 As our most realistic evaluation protocol, we use authentic fragments of continuous
 434 speech extracted from various YouTube audio clips. Each fragment contains at least
 435 one target keyword and lasts from a couple of seconds to more than a minute. We
 436 annotated these fragments by hand. The recordings are separate from the synthetic
 437 mixtures (🇵🇸 KW^{Mix}). These recordings are real, and contain numerous background
 438 words which are not the keywords we wish to detect. In total, we use 344 recordings
 439 with 511 keyword occurrences.

440 **Call centre target keywords (🇵🇸 CC^{Real})**

441 This dataset was provided by our industrial partner. It consists of several hundred
 442 recordings from the bank's call centre, with conversation lengths varying from seconds
 443 to hours, sparsely distributed occurrences of keywords and noise typical for call centre
 444 recordings. The recordings are provided in a typical call centre quality, an 8 kHz sample
 445 rate, and are unlabelled. During the project's time frame, the keyword labelling was
 446 performed for about 20% of the provided recordings. The availability of the dataset
 447 was strictly limited due to stringent privacy regulations, so we only tested our final
 448 approaches on it, as it required much coordination and effort to launch it on the bank's
 449 infrastructure.

450 **2.2. Similarity ranking models**

451 In this part, we briefly provide the specifications for our models and training procedures.




Figure 2. The architecture of the embedding model. Layer types denoted as: D – dropout, BN – batch normalisation, C – convolutional, MP – max-pooling, LReLU – LeakyReLU. Both the Siamese and prototypical models use the same embedder architecture, albeit with different loss functions and input formulation.

2.2.1. Siamese convolutional neural network

As our primary similarity ranking model, we use a Siamese neural network. The goal of this model is to map pairs of examples into pairs of embedding vectors. During optimisation, embeddings of examples belonging to the same class are placed close together in the embedding space.

The network consists of two identical convolutional branches serving as a speech embedder. In practice, we employ only one instantiation of the convolutional embedder as both branches are identical clones, with shared parameters.

The convolutional embedder uses a typical-VGG-like architecture, depicted in Figure 2, with four convolutional blocks, two dense layers, and a linear embedding. Each block combines 3×3 padded convolutions with batch normalisation, Leaky ReLU activations, max pooling, and 10% dropout. We use 32-64-128-64 filters, accordingly. After each dense layer (128 and 256 neurons), we add batch normalisation, Leaky ReLU, and 10% dropout. The final embedding has an output size of 128 features. We normalise the input recordings with 10% dropout and single-channel batch normalisation.

The Siamese model training is based on the contrastive loss, i.e. for pairs of training examples $\langle \mathbf{X}_i, \mathbf{X}_j \rangle$ created with a given selection strategy $\langle (i, j) \in \mathbb{P} \rangle$ we calculate the loss value according to the following formula:

$$\mathcal{L}(\mathbf{W}, \mathbf{X}, \mathbf{Y}) = \frac{1}{|\mathbb{P}|} \sum_{(i,j) \in \mathbb{P}} \left((1 - Y_{(i,j)}) \cdot D_{\mathbf{W}}(\mathbf{X}_i, \mathbf{X}_j) + Y_{(i,j)} \cdot \max\left(0, m - \sqrt{D_{\mathbf{W}}(\mathbf{X}_i, \mathbf{X}_j)}\right) \right)^2 \quad (1)$$

where $\mathbf{W}$ specifies the weights of the embedder, $Y_{(i,j)}$ term defines if the examples $\langle \mathbf{X}_i, \mathbf{X}_j \rangle$ are similar ($Y_{(i,j)} = 0$) or dissimilar ($Y_{(i,j)} = 1$), $D_{\mathbf{W}}$ is the similarity (distance) function defined for the embeddings $\langle G_{\mathbf{W}}(\mathbf{X}_i), G_{\mathbf{W}}(\mathbf{X}_j) \rangle$, and m is the margin value. In our case we use squared Euclidean distances, i.e.:

$$D_{\mathbf{W}}(\mathbf{X}_i, \mathbf{X}_j) = \|G_{\mathbf{W}}(\mathbf{X}_i) - G_{\mathbf{W}}(\mathbf{X}_j)\|_2^2 \quad (2)$$

This loss formulation keeps the embeddings of the same class closer together. On the other hand, the embeddings of different keyword classes are pushed away, so that they do not fall inside the margin, as illustrated in Figure 3.

We train the model for 100 epochs, using the Adam optimiser with default hyperparameters and a learning rate of 0.001, by optimising the contrastive loss based on squared Euclidean distances with a contrastive loss margin of 1.0. We use 50,000 training examples per epoch. In each epoch, we generate a number of 200 training episodes, consisting of 25 classes per episode and 10 samples per class. Our pair selection procedure uses the *hard negative* variant, i.e. we first create all possible positive pairs (matching keywords) from the samples in the current episode, and then we generate a corresponding an equal number of negative pairs selecting examples with the smallest distances.




Figure 3. Overview of the Siamese and prototypical training approaches. Siamese networks use a contrastive loss to make pairwise comparisons between examples from the dataset X . Examples from the same class are attracted to each other, while the examples from different classes are pushed back if their distance is lower than the defined margin m . The prototypical approach uses additional support examples S . Their mean vector defines the class prototype. Query examples Q of the same class are attracted to this prototype, while queries of different classes are pushed back in the same way as with Siamese training.

2.2.2. Prototypical network

In prototypical network experiments, we use the same embedder architecture as in section 2.2.1. ~~We train the~~ The main difference between the Siamese and prototypical approaches lies in the introduction of class prototypes. In Siamese models, we calculate the distances for similar/dissimilar pairs of individual examples from the dataset, whereas prototypical networks divide the training data into *support* (S) and *query* (Q) subsets. In each episode of training, a class prototype is created for each of the C selected classes by computing the mean value of the K vectors from the support subset:

$$\bar{S}_c = \frac{1}{K} \sum_{i=1}^K S_{c_i}, \quad (3)$$

where S_c denotes support examples of class c . The pair comparison is then performed between such class prototypes and all query examples selected in the given episode, resulting in a loss function:

$$\mathcal{L}(W, S, Q, Y, C) = \frac{1}{C|Q|} \sum_{c=1}^C \sum_{i=1}^{|Q|} \left((1 - Y_{c_i}) \cdot D_W(\bar{S}_c, Q_i) + Y_{c_i} \cdot \max\left(0, m - \sqrt{D_W(\bar{S}_c, Q_i)}\right)^2 \right), \quad (4)$$

where Y_{c_i} defines if the query example Q_i is of the same class c as the prototype $\bar{S}_c$ ($Y_{c_i} = 0$) or different ($Y_{c_i} = 1$). A more intuitive depiction of this procedure is presented in Figure 3.

Using this loss formulation, we train the prototypical model for 200 epochs comprising 100 episodes each, by optimising the prototypical loss with squared Euclidean distances and a margin of 1.0. In each prototypical training episode, we take 5 support samples and 15 query samples per class. Each episode consists of 25 classes. This setup results in the same total number of parameter updates for the Siamese and prototypical approaches.

2.2.3. Google speech embedder

In the experiments involving the Google speech embedder, we use the pre-trained model provided as version 1 on the TensorFlow Hub¹. When fine-tuning the network, we use the same training approach as in section 2.2.1, but with 50 epochs and 10 classes per episode.

¹ https://tfhub.dev/google/speech_embedding/1

493 2.3. Detector settings

494 Using the similarity models described in section 2.2, we generate distances between
 495 consecutive fragments of the analysed recording and each of the provided keyword
 496 templates (we use 10 templates per keyword class). This way, we obtain multiple values
 497 for each single time step, telling us how closely the current fragment resembles these
 498 various patterns. To aggregate such data into more meaningful detection decisions, we
 499 use various detection policies that we briefly describe in this section.

500 2.3.1. Common detection pipeline

501 In most of the experiments, we use the same detection pipeline. We analyse each frame
 502 of the recording, i.e. we employ a step size of 1. The actual frame size in milliseconds is
 503 determined by the spectrogram processing settings defined in Table 1. We then transform
 504 the generated distances with a median filter with a window length of 5 frames, obtaining
 505 smoothed similarity scores for each search pattern.

506 Based on these scores, we find the lowest average across all the analysed classes.
 507 If the average score falls below the threshold value defined for a given policy, we emit
 508 a detection marker for a given keyword at this particular time step. These marker
 509 emissions are then smoothed with a median filter applied across 15 frames.

510 Finally, if any keyword generates a continuous sequence of markers exceeding our
 511 minimum length of 25 consecutive frames, we return a detection at a given time step.
 512 We also introduce a minimal distance of 10 frames between two consecutive occurrences
 513 of the same keyword.

514 When evaluating the results of the system, we count the detection as a true positive
 515 if it falls at a time step representing the middle of the keyword occurrence, with a collar
 516 of 1 second. We adjust the policy threshold values based on the specifics of each model.
 517 For instance, for the Siamese model, we evaluate all values from 0 to 1 with a step size of
 518 0.05, as this range allows us to generate a complete precision-recall curve.

519 2.3.2. Additional post-processing for the Google speech embedder

520 In standard experiments with the pre-trained embedder model, we use the same ap-
 521 proach as described in section 2.3.1, albeit with a minimum length of consecutive matches
 522 reduced to 10 frames. ~~Nevertheless~~However, our experiments show that the similarity
 523 scores returned by direct distance calculations on speech embeddings generated by
 524 the Google model have a ~~distinct characteristic~~different scale and are less uniformly
 525 distributed across the different keyword classes.

526 Therefore, we introduce an additional post-processing step to our detection pipeline
 527 in the form of an *exciter* module. The role of this module is to make the similarity values
 528 more uniform across different keywords and filter out superfluous detections. To achieve
 529 this goal, we perform several operations.

530 First, we standardise the values of the distances in the detection matrix, visualised
 531 as input data in Figure 4. We perform this standardisation per ~~pattern (row)~~each of the 10
 532 patterns (rows). Then, we filter out all the values above the 5th percentile, leaving only
 533 the responses for the time steps with the closest matches. This way, we obtain several
 534 time series representing filtered and standardised similarity scores for ~~a single pattern~~
 535 single patterns.

536 After that, we use an envelope follower to perform excitation filtering on each
 537 of the time series. Therefore, after encountering a distance value below our threshold
 538 (signifying a potentially close match for a given template), we sustain this information
 539 in time with a decaying impulse response for up to 50 steps. When we find another
 540 close match, the decaying response is restarted. The result of this process is depicted in
 541 Figure 4.




Figure 4. Post-processing of the similarity scores generated by the pre-trained embedder. First matrix, denoted as *input*, shows an example of raw distances between the embeddings of keyword patterns and recording fragments. Each row corresponds to a single pattern, while each column represents a single time step. Brighter colours indicate smaller distances (closer matches). The second matrix shows the same values after row-wise standardisation and filtering through the *exciter* module.

2.4. Statistical analysis of the results

After running the experiments, we conduct a statistical analysis of the obtained results. For every experiment, we report training, patterns and evaluation datasets that were used. As the performance metrics, we use the Area Under Precision-Recall Curve (AUPRC) and F-score. The AUPRC is calculated based on the precision and recall values for different detection thresholds. Precision tells us how many of the reported keyword detections are correct, while recall shows us how many of the actual keyword utterances are retrieved. The AUPRC is provided with both micro- and macro-averaging. Micro-averaging takes into account the sizes of every keyword class we wish to detect by averaging over all the examples as a whole. This approach might be better suited to our problem, as we mainly deal with imbalanced datasets. On the other hand, macro-averaging, which aggregates the AUPRC values by first calculating them separately for each class, may also be helpful when interpreting the results. The F-score is defined as $F = \frac{\text{precision} + \text{recall}}{2}$, and it combines the precision and recall results into one metric, with higher values considered better.

3. Results

3.1. Isolated words classification benchmark

Before proceeding with the actual keyword detection experiments, we performed a baseline verification to make sure that our embedder architecture (*Siamese*), described in section 2.2.1, is sufficiently performant when processing audio data. To this end, we recreate the Speech Commands evaluation protocol [17] of classifying single-word utterances as one of the 12 possible classes. This is the only experiment where we do not downsample the input recordings. Instead, we maintain the original 16 kHz sampling rate, adjusting the spectrogram preprocessing settings accordingly.

Our convolutional model achieves a 94.5% top-1 classification accuracy on the Speech Commands v1 dataset. This performance is on par with similar models processing mel spectrograms, as presented in Table 3. While more sophisticated models can achieve better classification results, we deemed the differences not significant enough to warrant the trade-off these models introduce in terms of complexity and training time.

We also verify the same architecture on a downsampled version of the Speech Commands recordings. As expected, this impairs the accuracy of the evaluated model. However, the results show that 8 kHz recordings can still provide sufficient information for proper classification.

Table 3. Top-1 classification accuracy for isolated words of Speech Commands v1.

Authors	Method	Accuracy
(Warden, 2018) [17]	<i>Baseline CNN model</i>	85.4%
(Tang, 2017) [18]	CNN with MFCC features	90.2%
(de Andrade, 2018) [20]	CNN with 80-band mel spectrograms	94.1%
(Zhang, 2017) [19]	Depthwise separable CNN	95.4%
(Choi, 2019) [23]	CNN with temporal convolution	96.1%
(Zeng, 2019) [21]	DenseNet with BiLSTM layers	96.2%
(Lin, 2018) [22]	<i>EdgeSpeechNets</i>	96.3%
(Mittermaier, 2020) [24]	Parametrized sinc-convolutions	96.4%
(Majumdar, 2020) [25]	<i>MatchboxNet</i>	97.5%
(Li, 2020) [26]	<i>Wide-ResNet with MoEx</i>	98.0%
<i>Ours (Siamese)</i>	CNN with dropout, 40-band mel spectrograms	94.5%
<i>Ours (Siamese), 8 kHz</i>	<i>Same as above, downsampled recordings</i>	92.0%

575 3.2. Keyword detection in a monolingual setup

576 As a first step in evaluating keyword detectors on continuous audio data streams, we
 577 analyse our Siamese approach in monolingual scenarios, i.e. by training and validating
 578 the model on recordings of the same language. For selected experiments, we also
 579 compare its performance to a prototypical network.

580 In each evaluation setting, we create a separate *precision-recall curve* (PRC) for each
 581 of the analysed keywords. The PRC shows detection performance at different values
 582 of the similarity threshold. We aggregate this information across classes using either
 583 micro- (*instance*) or macro- (*class*) averaging, a common approach in multi-class problems [48]

584 ~

585 To summarise the performance of a model with a single value, we report the *Area*
 586 *Under Precision-Recall Curve* (AUPRC). We also highlight the best F-score achieved by
 587 the model across different threshold values. Table 4 presents the results obtained in the
 588 analysed monolingual scenarios.

589 Our initial verification, described in section 3.1, has shown that a Siamese convolutional
 590 neural network can effectively differentiate between utterances of different keywords of
 591 the Speech Commands dataset. We also confirm this capability in a streaming evaluation
 592 by employing the Siamese model with standard detection settings on a synthetic mixture
 593 of Speech Commands keywords (🇺🇸 SC1^{Mix}). When the vocabulary present in the
 594 target recording consists solely of the expected keywords, the model can achieve an
 595 outstanding performance of 91.3% micro-AUPRC and an F-score of 0.94. This result
 596 shows that our detection approach allows for properly recognising keywords with
 597 shifted word boundaries.

Table 4. Keyword detection performance for models trained and evaluated in monolingual scenarios. The Area Under Precision-Recall Curve is based on the keyword classes present in the search patterns. We report it either with micro- or macro-averaging. The F-score column represents the best result achieved by the model across different settings of the detection threshold. All models use the default detector configuration.

Model	Training ¹ (# KW)	Patterns ²	Evaluation ³ (# KW)	AUPRC		F-score
				Micro	Macro	
English 🇺🇸 → English 🇺🇸						
Siamese	🇺🇸 SC 1 (30)	🇺🇸 SC 1	🇺🇸 SC 1 ^{Mix} (30)	91.3%	91.3%	0.94
Siamese	🇺🇸 SC 2, CV (37)	🇺🇸 SC 1	🇺🇸 SC 1 ^{Mix} (30)	87.9%	88.0%	0.92
Siamese	🇺🇸 SC 1, SWC (4K+)	🇺🇸 SC 1	🇺🇸 SC 1 ^{Mix} (30)	69.4%	65.6%	0.79
Siamese	🇺🇸 ALL (4K+)	🇺🇸 SC 1	🇺🇸 SC 1 ^{Mix} (30)	70.8%	67.4%	0.81
Prototypical	🇺🇸 SC 1 (30)	🇺🇸 SC 1	🇺🇸 SC 1 ^{Mix} (30)	80.2%	81.5%	0.88
Siamese	🇺🇸 SC 1 (30)	🇺🇸 SC 1	🇺🇸 SC 1 ^{Vox} (30)	60.4%	65.0%	0.62
Siamese	🇺🇸 SC 2, CV (37)	🇺🇸 SC 1	🇺🇸 SC 1 ^{Vox} (30)	55.2%	62.4%	0.57
Siamese	🇺🇸 SC 1, SWC (4K+)	🇺🇸 SC 1	🇺🇸 SC 1 ^{Vox} (30)	38.0%	43.4%	0.47
Siamese	🇺🇸 ALL (4K+)	🇺🇸 SC 1	🇺🇸 SC 1 ^{Vox} (30)	43.7%	46.9%	0.50
Prototypical	🇺🇸 SC 1 (30)	🇺🇸 SC 1	🇺🇸 SC 1 ^{Vox} (30)	40.8%	50.2%	0.54
Siamese	🇺🇸 SC 1 (30)	🇺🇸 ΔSC	🇺🇸 ΔSC ^{Mix} (5)	48.3%	53.1%	0.65
Prototypical	🇺🇸 SC 1 (30)	🇺🇸 ΔSC	🇺🇸 ΔSC ^{Mix} (5)	39.6%	42.3%	0.59
Polish 🇵🇱 → Polish 🇵🇱						
Siamese	🇵🇱 ALL (5K+)	🇵🇱 KW	🇵🇱 KW ^{Mix} (22)	8.8%	1.3%	0.20
Prototypical	🇵🇱 ALL (5K+)	🇵🇱 KW	🇵🇱 KW ^{Mix} (22)	5.9%	1.9%	0.22
Siamese	🇵🇱 ALL (5K+)	🇵🇱 KW	🇵🇱 KW ^{Real} (22)	0.0%	0.0%	0.02
Prototypical	🇵🇱 ALL (5K+)	🇵🇱 KW	🇵🇱 KW ^{Real} (22)	0.0%	0.0%	0.01

¹ Training data consists of utterances from Speech Commands ( SC 1,  SC 2), Common Voice ( CV), and Spoken Wikipedia Corpus ( SWC). We also use a combined dataset ( ALL). For Polish language we use all accessible training data ( ALL), i.e.:  CV,  WUT and  TTS. Number of keyword classes (# KW) is denoted with a subscript.

² English search patterns are extracted from Speech Commands training data ( SC 1) or from the subset present only in the second version of the dataset ( ΔSC). Polish templates come from the target keywords dataset ( KW). We use 10 examples for each keyword class.

³ English models are evaluated with 30 keyword classes, either on fully synthetic mixtures of Speech Commands utterances ( SC 1^{Mix}) or Speech Commands keywords overlaid on VoxCeleb recordings ( SC 1^{Vox}). We also show the performance on the *delta* dataset mixtures ( ΔSC^{Mix}, i.e. 5 classes). Polish evaluations assess 22 keywords either in synthetic mixtures ( KW^{Mix}) or actual YouTube audio streams ( KW^{Real}).

As a first step in evaluating keyword detectors on continuous audio data streams, we analyse our Siamese approach in monolingual scenarios, i.e. by training and validating the model on recordings of the same language. For selected experiments, we also compare its performance to a prototypical network.

In each evaluation setting, we create a separate *precision-recall curve* (PRC) for each of the analysed keywords. The PRC shows detection performance at different values of the similarity threshold. We aggregate this information across classes using either micro- (instance) or macro- (class) averaging, a common approach in multi-class problems [48].

To summarise the performance of a model with a single value, we report the *Area Under Precision-Recall Curve* (AUPRC). We also highlight the best F-score achieved by the model across different threshold values. Table 4 presents the results obtained in the analysed monolingual scenarios.

Our initial verification, described in section 3.1, has shown that a Siamese convolutional neural network can effectively differentiate between utterances of different keywords of

the Speech Commands dataset. We also confirm this capability in a streaming evaluation by employing the Siamese model with standard detection settings on a synthetic mixture of Speech Commands keywords (Δ). When the vocabulary present in the target recording consists solely of the expected keywords, the model can achieve an outstanding performance of 91.3% micro-AUPRC and an F-score of 0.94. This result shows that our detection approach allows for properly recognising keywords with shifted word boundaries.

To measure the robustness of the detector to background distractors, we evaluate it on Speech Commands keywords mixed into fragments of conversations of the VoxCeleb dataset (Δ SC 1^{Vox}). In this case, the detection accuracy drops to a level of 60.4% micro-AUPRC. Nevertheless, the performance is satisfactory for a cross-speaker keyword spotter as the system can correctly highlight more than $\frac{2}{3}$ of keyword occurrences while maintaining a precision of 55%. Results at this level would be usable for prospective users of such a tool. Unfortunately, we have to admit that the semi-synthetic nature of this evaluation oversimplifies the task presented to the detector, making it an upper bound on achievable accuracy. Due to the lack of adequate data, we could not verify how such detectors would cope with more natural keyword occurrences and diverse recording conditions.

On the other hand, our problem setting assumes that the detector can only be trained on copious amounts of generic keywords. After that, it should cope well with limited examples of target keywords, especially since the end-user can extend the vocabulary after the system's deployment. We assess this aspect with an experiment using keywords not occurring in the original training data, isolated from the second version of the Speech Commands dataset (Δ SC^{Mix}). This evaluation scenario confirms that few-shot learning is quite difficult. A drop of the F-score to 0.65 on synthetic mixtures indicates that when combined with more natural evaluation settings, detection of completely new keywords might be problematic. A more reasonable approach would require at least some retraining with the extended vocabulary.

Interestingly, across all the experiments, the extension of the Speech Commands training data with other datasets proves detrimental to the model's performance. We analyse this phenomenon more closely in Figure 5 by visualising the embedding space created with Siamese models trained solely on the Speech Commands data and in combination with the Spoken Wikipedia Corpus recordings. The model trained only on Speech Commands utterances maps the examples from this dataset to groups with much clearer separability between the particular keywords. Apart from some stray confusions, the only intermixing of keywords happens for the "three-tree" pair, showing that the model is indeed focusing on the acoustic similarity of the provided samples. The inclusion of more diverse training examples from the Spoken Wikipedia Corpus prohibits the model from learning an equally discriminative mapping for the Speech Commands keywords. The created groups of examples show much more bleed between keywords. Unfortunately, we were unable to devise a simple mitigation technique for this issue. It is possible that a more nuanced training procedure could create a more robust representation using all of the available data.

We also extend our investigation with additional experiments employing a prototypical network model instead of the Siamese embedder. In the analysed monolingual settings, this approach proves to be less performant than the Siamese counterpart. In Table 4, we report only the results for training with Speech Commands data, but the tendency remains unchanged with different dataset setups.

Finally, we conclude with evaluations performed on Polish datasets. Unfortunately, the experiments confirm our initial concerns. Models trained solely on such limited datasets are entirely unusable in detecting keywords in continuous recordings.

² We employ a zero minimum distance between embedded points. Other visualisation settings use standard values of the *umap-learn* Python package, i.e. 15 neighbours with an Euclidean metric for Uniform Manifold Approximation and Projection.




Figure 5. Comparison of the UMAP visualisations² of the Speech Commands training examples ( SC 1) processed through the Siamese embedder: (a) embeddings generated with a model trained only on the Speech Commands data; (b) embeddings generated with a model trained on recordings both from the Speech Commands and the Spoken Wikipedia Corpus datasets.

3.3. Keyword detection in a cross-lingual setup

Following the expected failure of models trained exclusively on Polish recordings, we try to approach the problem of detecting Polish keywords with models trained on English datasets, i.e. in a cross-lingual mode. We hope that the general audio processing capabilities acquired by training on more extensive English datasets will allow for a successful transfer of knowledge to Polish recognition tasks, despite the inherent differences between acoustic features of the languages.

Table 5 summarises our findings in cross-lingual scenarios. When trained solely on the Speech Commands recordings, both Siamese and prototypical models are better in detecting Polish keywords than their counterparts trained on limited Polish data. Nevertheless, this improvement is still insufficient to achieve satisfactory performance. The Siamese model, which outperforms the prototypical approach, achieves a micro-AUPRC of only 20.3% on synthetic mixtures of Polish keywords ( KW^{Mix}). In practical terms, this result means that we can roughly achieve a precision of 70% at a 25% recall rate. However, the outcomes for individual keywords vary widely. For instance, the best performing one, *umowa*, has a recall of 67% with 96% precision. Unfortunately, although we can find several other classes with potentially usable results, many keywords have a near-zero detection rate.

Table 5. Keyword detection performance for models trained and evaluated in cross-lingual scenarios. Metric values are reported in the same manner as in Table 4.

Model	Training ¹ (# KW)	Patterns ²	Evaluation ³ (# KW)	AUPRC		F-score
				Micro	Macro	
English 🇺🇸 → Polish 🇵🇱						
Siamese	🇺🇸 SC 1 ₍₃₀₎	🇵🇱 KW	🇵🇱 KW ^{Mix} ₍₂₂₎	20.3%	10.6%	0.39
Siamese	🇺🇸 SC 2, CV ₍₃₇₎	🇵🇱 KW	🇵🇱 KW ^{Mix} ₍₂₂₎	10.3%	8.9%	0.31
Siamese	🇺🇸 SC 1, SWC _(4K+)	🇵🇱 KW	🇵🇱 KW ^{Mix} ₍₂₂₎	15.5%	11.6%	0.34
Siamese	🇺🇸 ALL _(4K+)	🇵🇱 KW	🇵🇱 KW ^{Mix} ₍₂₂₎	7.3%	10.7%	0.26
Prototypical	🇺🇸 SC 1 ₍₃₀₎	🇵🇱 KW	🇵🇱 KW ^{Mix} ₍₂₂₎	13.4%	2.6%	0.27
Siamese	🇺🇸 SC 1 ₍₃₀₎	🇵🇱 KW	🇵🇱 KW ^{Real} ₍₂₂₎	0.2%	0.2%	0.03
Siamese	🇺🇸 SC 2, CV ₍₃₇₎	🇵🇱 KW	🇵🇱 KW ^{Real} ₍₂₂₎	0.1%	0.2%	0.02
Siamese	🇺🇸 SC 1, SWC _(4K+)	🇵🇱 KW	🇵🇱 KW ^{Real} ₍₂₂₎	0.1%	0.1%	0.02
Siamese	🇺🇸 ALL _(4K+)	🇵🇱 KW	🇵🇱 KW ^{Real} ₍₂₂₎	0.1%	0.1%	0.01
Prototypical	🇺🇸 SC 1 ₍₃₀₎	🇵🇱 KW	🇵🇱 KW ^{Real} ₍₂₂₎	0.3%	0.2%	0.04
Combined 🇺🇸 + 🇵🇱 → Polish 🇵🇱						
Siamese	🇺🇸 ALL, 🇵🇱 ALL _(9K+)	🇵🇱 KW	🇵🇱 KW ^{Mix} ₍₂₂₎	6.5%	7.2%	0.23
Prototypical	🇺🇸 ALL, 🇵🇱 ALL _(9K+)	🇵🇱 KW	🇵🇱 KW ^{Mix} ₍₂₂₎	7.6%	2.7%	0.22
Siamese	🇺🇸 ALL, 🇵🇱 ALL _(9K+)	🇵🇱 KW	🇵🇱 KW ^{Real} ₍₂₂₎	0.1%	0.1%	0.02
Prototypical	🇺🇸 ALL, 🇵🇱 ALL _(9K+)	🇵🇱 KW	🇵🇱 KW ^{Real} ₍₂₂₎	0.2%	0.0%	0.04

¹ Training data is denoted in the same way as in Table 4.² All search patterns come from the Polish target keywords dataset ( KW).³ Evaluation is performed on Polish mixtures ( KW^{Mix}) and real recordings ( KW^{Real}) with 22 keywords.

Following the expected failure of models trained exclusively on Polish recordings, we try to approach the problem of detecting Polish keywords with models trained on English datasets, i.e. in a cross-lingual mode. We hope that the general audio processing capabilities acquired by training on more extensive English datasets will allow for a successful transfer of knowledge to Polish recognition tasks, despite the inherent differences between acoustic features of the languages.

Table 5 summarises our findings in cross-lingual scenarios. When trained solely on the Speech Commands recordings, both Siamese and prototypical models are better in detecting Polish keywords than their counterparts trained on limited Polish data. Nevertheless, this improvement is still insufficient to achieve satisfactory performance. The Siamese model, which outperforms the prototypical approach, achieves a micro-AUPRC of only 20.3% on synthetic mixtures of Polish keywords (). In practical terms, this result means that we can roughly achieve a precision of 70% at a 25% recall rate. However, the outcomes for individual keywords vary widely. For instance, the best performing one, *umowa*, has a recall of 67% with 96% precision. Unfortunately, although we can find several other classes with potentially usable results, many keywords have a near-zero detection rate.

When looking at possible training dataset extensions, we observe a similar situation as with monolingual models. Training solely on the Speech Commands dataset proves to be the most efficient way to achieve acoustically discriminative embedders. Additional recordings are similarly detrimental in creating an embedding space appropriate for cross-lingual transfer to Polish patterns.

This observation also holds for further extensions with Polish training recordings, though not without some caveats. Initially, we hoped to fill potential gaps in the generated embedding space by introducing additional training examples more closely




Figure 6. Comparison of the UMAP visualisations of Polish keywords processed through the Siamese embedder: (a) embeddings generated with a model trained only on the Speech Commands data; (b) embeddings generated with a model trained on all the available data (both English and Polish); Polish keywords are extracted directly from YouTube videos. Synthetic mixtures of these keywords are denoted as ■ KW^{Mix} throughout the results section.

resembling the phonetic structure of the target patterns. However, the outcome of this process was somewhat ambiguous.

On the one hand, if we compare the 2D representations of the embedding space presented in Figure 6, the second model, supplied with extended training data, groups Polish keywords into much tighter clusters. Although low-dimensionality mappings of complex embedding spaces might be misleading at times, this visual difference most probably indicates that the second model can more effectively discriminate between the keyword classes.

On the other hand, this capability does not translate to an advantage when comparing the detection performance of both models. While we do not have a definite explanation for this phenomenon, we hypothesise that a more dispersed representation might be actually beneficial in our scenario. In contrast to the evaluation on Speech Commands keywords, where the evaluation and search patterns come from the same distribution, Polish patterns used as keyword templates differ in the recording conditions from the YouTube evaluation fragments. Therefore, a broader, less regularised representation might expose more potential points of contact to find close neighbours that could match actual keyword occurrences in audio streams. Such an increased coverage could be significant since the data manifold of Polish search patterns is quite limited to begin with.

This representational problem is accentuated by evaluations performed on longer, authentic Polish audio streams (■ KW^{Real}). In this scenario, all systems fail to provide any hint of usable results. We can devise a two-fold explanation for this behaviour.

First of all, the embedding space learnt by our models does not capture the acoustic differences at a detailed enough level. Therefore, these models cannot handle utterances outside of their limited vocabulary. Non-keyword audio content easily derails the detectors, which is confirmed by numerous false positives.

The second factor is associated with the difficulty of the problem itself. In contrast to our English setups, the evaluation performed on Polish recordings uses longer fragments




Figure 7. Comparison of the UMAP visualisations of the Speech Commands training examples (🇺🇸 SC 1) processed through the Google speech embedder: (a) embeddings generated with a pre-trained model without fine-tuning; (b) embeddings generated with a pre-trained model fine-tuned on the Speech Commands dataset.

of naturally sounding speech from diverse recording conditions, thus being the closest to an actual environment in which these kinds of systems might be deployed. Consequently, many keyword occurrences are less perceptible than in semi-synthetic mixtures.

Based on all these observations, we suspect that a self-supervised approach to training acoustic models could be promising in solving similar problems as described in this paper. Self-supervision should help create rich representations, more robustly capturing the differences between various words present in the recording – all without the need for extra hand-labelling. Unfortunately, devising a sensible self-supervised approach is not a trivial task, and each iteration of such an experiment requires a significant computational effort to train the actual model. Therefore, due to the constrained timeline of our project, we were unable to explore this option further.

3.4. Keyword detection with generic speech embeddings

Instead of investigating self-supervised techniques, in this last experimental section, we concentrate on models trained traditionally, in a fully supervised manner, but on much bigger datasets. Based on our assumption that the main factor limiting the performance of our models is the lack of a more generic and robust internal speech representation, we replace our previously analysed similarity models with a pre-trained speech embedding model provided by Google [6]. It has been trained on more than 100,000 hours of English audio clips, which we expect should cover a big part of possible recording conditions and variants of speech. We present the results of applying this model in Table 6, divided into four different approaches.

Table 6. Keyword detection performance for the pre-trained speech embedder. We use it either with default detection settings or with an additional post-processing procedure. We also compare the embeddings generated directly from the pre-trained model and from a model fine-tuned on selected datasets. Metric values are reported in the same manner as in Table 4.

Model	Training ¹ (# KW)	Patterns ²	Evaluation ³ (# KW)	AUPRC		F-score
				Micro	Macro	
Speech embedder 🇺🇸, pre-trained						
Google	—	🇺🇸 SC 1	🇺🇸 SC 1 ^{Mix} ₍₃₀₎	15.3%	71.5%	0.41
Google	—	🇺🇸 SC 1	🇺🇸 SC 1 ^{Vox} ₍₃₀₎	3.7%	38.5%	0.07
Google	—	🇺🇸 ΔSC	🇺🇸 ΔSC ^{Mix} ₍₅₎	4.6%	4.3%	0.65
Google	—	🇬🇧 KW	🇬🇧 KW ^{Mix} ₍₂₂₎	27.2%	62.5%	0.41
Google	—	🇬🇧 KW	🇬🇧 KW ^{Real} ₍₂₂₎	3.4%	21.0%	0.05
Speech embedder 🇺🇸, pre-trained, with post-processing						
Google (post)	—	🇺🇸 SC 1	🇺🇸 SC 1 ^{Mix} ₍₃₀₎	84.8%	87.7%	0.84
Google (post)	—	🇺🇸 SC 1	🇺🇸 SC 1 ^{Vox} ₍₃₀₎	40.6%	49.8%	0.48
Google (post)	—	🇺🇸 ΔSC	🇺🇸 ΔSC ^{Mix} ₍₅₎	86.5%	87.8%	0.85
Google (post)	—	🇬🇧 KW	🇬🇧 KW ^{Mix} ₍₂₂₎	46.2%	53.9%	0.69
Google (post)	—	🇬🇧 KW	🇬🇧 KW ^{Real} ₍₂₂₎	6.6%	9.2%	0.17
Speech embedder 🇺🇸, fine-tuning						
Google	🇺🇸 SC 1 ₍₃₀₎	🇺🇸 SC 1	🇺🇸 SC 1 ^{Mix} ₍₃₀₎	15.0%	30.1%	0.30
Google	🇺🇸 SC 1 ₍₃₀₎	🇺🇸 SC 1	🇺🇸 SC 1 ^{Vox} ₍₃₀₎	1.2%	1.8%	0.03
Google	🇬🇧 WUT ₍₃₆₎	🇬🇧 KW	🇬🇧 KW ^{Mix} ₍₂₂₎	18.4%	19.9%	0.36
Google	🇬🇧 WUT ₍₃₆₎	🇬🇧 KW	🇬🇧 KW ^{Real} ₍₂₂₎	0.7%	2.8%	0.02
Speech embedder 🇺🇸, fine-tuning, with post-processing						
Google (post)	🇺🇸 SC 1 ₍₃₀₎	🇺🇸 SC 1	🇺🇸 SC 1 ^{Mix} ₍₃₀₎	24.4%	25.7%	0.42
Google (post)	🇺🇸 SC 1 ₍₃₀₎	🇺🇸 SC 1	🇺🇸 SC 1 ^{Vox} ₍₃₀₎	1.3%	1.8%	0.04
Google (post)	🇬🇧 WUT ₍₃₆₎	🇬🇧 KW	🇬🇧 KW ^{Mix} ₍₂₂₎	15.5%	16.0%	0.37
Google (post)	🇬🇧 WUT ₍₃₆₎	🇬🇧 KW	🇬🇧 KW ^{Real} ₍₂₂₎	1.3%	1.5%	0.05

¹ The speech embedding model [6] is pre-trained on English YouTube audio clips. In most experiments, we use the generated embeddings directly, without any further training of the model. In fine-tuning experiments, we use utterances from the Speech Commands dataset ( SC 1) or our own recordings ( WUT).

^{2,3} Search patterns and evaluation recordings are denoted in the same way as in Table 4.

755 Instead of investigating self-supervised techniques, in this last experimental section,
 756 we concentrate on models trained traditionally, in a fully supervised manner, but on
 757 much bigger datasets. Based on our assumption that the main factor limiting the performance
 758 of our models is the lack of a more generic and robust internal speech representation, we
 759 replace our previously analysed similarity models with a pre-trained speech embedding
 760 model provided by Google [6]. It has been trained on more than 100,000 hours of
 761 English audio clips, which we expect should cover a big part of possible recording




Figure 8. Comparison of the UMAP visualisations of Polish keywords processed through the Google speech embedder: (a) embeddings generated with a pre-trained model without fine-tuning; (b) embeddings generated with a pre-trained model fine-tuned on ■ WUT training examples.

~~conditions and variants of speech. We present the results of applying this model in Table 6, divided into four different approaches.~~

First, we use the pre-trained embedder directly, without any fine-tuning, with the standard detector. This setup means that the pre-trained model returns coordinates in the embedding space for each analysed fragment of the evaluation recording and the keyword patterns. We then compute the distances to each provided keyword example and use properly adjusted thresholds with our primary detection approach involving aggregation and filtering. The macro-results of this system are good on the Speech Commands dataset (🇺🇸 SC 1^{Mix}), although not as good as the specifically trained models. However, the system incorporating the pre-trained embedder is much better with Polish keywords (■ KW^{Mix}). It is also the first system to obtain non-zero results on the most interesting, authentic Polish evaluation (■ KW^{Real}).

Comparing Figures 5(a) and 7(a), we can observe some significant changes in the embedding space generated for the Speech Commands examples. The Google speech embedder mapping of keyword classes is more complex than in the Siamese model's case. Instead of densely packed clusters, the pre-trained model distributes the encountered examples across more elongated shapes. This change indicates that the speech embedder captures more variability factors in the data, apart from the keyword class, creating a more nuanced representation. Concurrently, it still creates a distinct separation between the classes, although with some overlap in the middle of the plot. The behaviour in this central region might explain why the raw performance of the pre-trained embedder on Speech Commands data might be slightly worse in direct comparison to a Siamese model.

Despite this robust representation of speech fragments generated by the Google model, a direct application of the returned similarity distances proves to be relatively ineffective in a detection setting. While most keywords produce pretty accurate results, a select few (*eight*, *off*, and *up*) create a tremendous number of false positives. These errors drive down the model's average performance, indicated by the significant discrepancy between the micro- and macro- values of the reported metrics. This problem is presented more explicitly in Figure A1, in the Appendix section A.

Table 7. Keyword detection performance on the final evaluation dataset (■ CC^{Real}).

Model (detector)	Micro			Macro		
	Precision	Recall	F-score	Precision	Recall	F-score
Google	0.4%	18.5%	0.01	14.2%	25.6%	0.18
Google (post)	3.6%	17.0%	0.06	8.2%	22.2%	0.12

We mitigate this issue by introducing the post-processing approach described in section 2.3.2. By adjusting the detector setting in this way, we can almost entirely filter out these massive false detections reported by the model, as shown in Figure A2. This modification comes at a small cost of slightly worse outcomes on some of the high-quality keywords, but on the whole, it creates a detection system with a much more practical behaviour.

The performance of the post-processing version of the system on Speech Commands data (■ SC 1^{Mix}) is comparable to our Siamese and prototypical models trained from scratch. The results for semi-synthetic evaluation with ■ SC 1^{Vox} are slightly worse, but the performance on the “delta” classes (■ ΔSC^{Mix}) is on par with the original Speech Commands keywords. This improvement shows that the Google speech embedder has been exposed to an extensive range of potential English keywords, making it easier to adjust the operating vocabulary on the fly.

The most important change from our perspective is the improvement in Polish keyword detection. A good representation of the Polish keywords, as depicted in Figure 8(a), allows the model to approach much more sensible levels even with the baseline detector. However, the post-processing variant brings the best F-score on the ■ KW^{Mix} dataset up to 0.69, significantly outperforming all the other approaches analysed in this paper. The macro-AUPRC values on authentic Polish recordings (■ KW^{Real}) are better with the standard detector, but post-processing improves both the micro-AUPRC and the F-score value. Unfortunately, it is still somewhat discouraging, achieving a level of 0.17.

Finally, we also evaluate fine-tuned versions of the Google speech embedder, hoping to combine the advantages of both worlds – pre-trained generic representations and dataset-specific mapping. [The fine-tuning procedure is described in section 2.2.3.](#) Unfortunately, our efforts very swiftly prove to be destructible to the intricate representational capabilities of the original model. As evidenced by Figure 7(b), fine-tuning on the Speech Commands dataset creates embeddings with some tightly packed clusters, similar to the behaviour of the Siamese model. Still, most of the keywords become intermixed after this procedure. This degradation is also confirmed quantitatively, as the results for the fine-tuned models are comparatively worse across the board. This drop in performance also pertains to fine-tuning on Polish data, although the embedding space shown in Figure 8(b) seems to be less impacted.

3.5. Final evaluation

As a final step in evaluating our keyword detection systems, we analyse their behaviour on the actual target recordings of call centre conversations (■ CC^{Real}). Due to organisational impediments, we could perform this procedure only once, with a limited number of systems. Therefore, we employ the pre-trained embedder solution with our Polish keyword patterns, as such a combination presents the most promising results on continuous Polish recordings. We choose both variants of the system, with the standard detection pipeline and post-processing. The pre-trained embedder is used directly, without fine-tuning.

Table 7 summarises the findings of our final evaluation. The detection pipeline was executed only once with a predetermined detection threshold. Therefore, instead of full AUPRC numbers, we present actual metric values obtained at this sensitivity point.

Both models perform poorly, especially when looking at the micro-aggregation scheme. As the results of our interim validations on Polish datasets were relatively poor, we did not expect a much better outcome, especially since we had to cope with a domain shift on entirely unseen data.

A more detailed analysis shows that the *Google* model returns a very high number of false positives for two short keywords, “*blik*” and “*link*”. As the utterances of these keywords contain only one syllable, it is understandable that matching based only on acoustic features might be unsuccessful. Although our post-processing method can correctly suppress these erroneous detections, it introduces its own biases, keeping the final performance still at a low level.

Concluding our evaluation, we must admit that the quality of predictions generated by the system with the pre-trained embedder proved to be disappointing. This approach was insufficient in creating a general robust keyword spotter for the Polish language. However, there can still be a small added value of such a system when employed as a support tool to highlight specific keywords. Although it will not recall all the occurrences, it can still help people performing the reviewing work if its precision is sufficiently high. We were able to find a couple of characteristic multi-syllable keywords that exhibited promising results in this regard. For instance, words such as “*reklamacja*” and “*transakcja*” had a precision of over 75% combined with a recall rate of 10–20%. While it is not exactly what we have hoped for when devising the system, our proof-of-concept solution has shown that the main focus when creating robust Polish keyword spotting systems should lie on the mundane task of data annotation.

4. Discussion

In this paper, we have explored the problem of cross-speaker keyword spotting. ~~The goal of our work was to create a proof-of-concept solution that could effectively detect Polish keywords in call-centre recordings. Unfortunately, working in a~~ for the low-resource setting of the Polish language, ~~we~~. We were limited by the lack of datasets suitable for training production-quality automatic speech recognition systems.

Therefore, we have focused on spotting keywords with detectors based on acoustic similarity. These approaches are generally less demanding on the data annotation side. We evaluated two similarity ranking models, i.e. Siamese and prototypical networks. Our experiments with English datasets have shown that these methods can create acoustically discriminative representations of processed recordings when provided with sufficiently diverse training examples.

Unfortunately, although the perceptual principles of comparing two audio fragments remain the same on the fundamental level, such models were unable to generalise to other languages. The acoustic differences between languages and recording conditions proved to be too big for such a cross-lingual transfer to succeed.

Due to the data scarcity problem, we could not create robust keyword spotters on Polish data either. Therefore, we have evaluated a different approach by utilising a generic speech embedding model provided by Google, extensively trained on thousands of hours of English speech. The advantage provided by a very comprehensive training dataset could be seen in more complex representations of the speech samples and much better adaptability to cross-lingual transfer. Although the evaluation on Polish synthetic recordings was quite promising, even with this pre-trained embedder, we still could not create a system that would be fully functioning in realistic scenarios and could effectively process naturally sounding continuous audio streams.

Based on our research findings, we reckon that acoustic similarity comparisons can be a viable approach in various audio matching problems. Nevertheless, the task of creating a robust generic embedding space for speech recordings is not easy. Our evaluations indicate that simply extending the scope of training data is not always profitable. However, it is also possible that we were unsuccessful in finding more

effective training methods, more suitable for mixed datasets, [such as domain adaptation of embeddings, that has proved successful in NLP tasks \[49\]](#).

Looking at the visualisations of the generated embeddings, we see that similarity models can usually maintain correct separation between various classes. Therefore, we expect that some adjustments to the post-processing schemes could improve the quality of final detections. However, to achieve this, more robust validation procedures and datasets would be needed.

This problem is particularly relevant since the disparity between the performance of keyword classifiers and keyword detectors is striking. In fact, during our work we could not verify the performance of our systems on representative English audio streams. We think that establishing new, more realistic evaluation protocols for keyword spotters would be an interesting extension for future work.

Additionally, recent developments in self-supervised training of audio representations create exciting opportunities for low-resource languages, such as Polish. Unfortunately, these types of models were not yet available for use at the onset of our project. If we were to recommence this work, we expect that solutions such as XLSR-53 would prove helpful in bridging the generalisation gap that we encounter in low-resource scenarios.

5. [Conclusions](#)

[The goal of our work was to create a proof-of-concept solution that could effectively detect Polish keywords in low-quality call centre recordings. Based on our research hypothesis, we developed keyword detectors employing few-shot acoustic similarity models. The models have a satisfactory accuracy in English and for selected Polish keywords, but they fail for many shorter Polish utterances. Effectively, the created software system enables navigation in call centre recordings only for a limited subset of Polish keywords. However, such functionality can still reduce the processing time in the complaint processes.](#)

Author Contributions: Conceptualisation, K.P. and R.N.; methodology, K.P., Ł.L., and K.R.; software, Ł.L., K.P., and K.R.; validation, K.P. and Ł.L.; formal analysis, K.P. and Ł.L.; investigation, K.P., Ł.L., and K.R.; resources, Ł.L., K.P., K.R., and R.N.; data curation, K.P., Ł.L., and K.R.; writing—original draft preparation, K.P., Ł.L., and K.R.; writing—review and editing, K.P., Ł.L., K.R., and R.N.; visualisation, K.P.; supervision, K.P. and R.N.; project administration, K.P. and R.N.; funding acquisition, R.N. All authors have read and agreed to the published version of the manuscript.

Funding: This research was funded by the Warsaw University of Technology statutory research grant in 2021 and by the mBank SA research project in 2020. KP was co-funded by the Foundation for Polish Science (grant no POIR.04.04.00-00-14DE/18-00 carried out within the Team-Net program co-financed by the European Union under the European Regional Development Fund).

Data Availability Statement: Publicly available datasets:

Speech Commands v1: http://download.tensorflow.org/data/speech_commands_v0.01.tar.gz

Speech Commands v2: http://download.tensorflow.org/data/speech_commands_v0.02.tar.gz

Mozilla Common Voice: <https://commonvoice.mozilla.org/en/datasets>

Spoken Wikipedia Corpus: <https://nats.gitlab.io/swc>

Acknowledgments: We would like to thank Karol Chęciński, Piotr Gawrysiak and Kamil Żbikowski for their support.

Conflicts of Interest: The authors declare no conflict of interest.


934 **Appendix A**

Figure A1. Example of a detection report for the Google speech embedding model with standard detection settings, without post-processing. The upper part shows a timeline of the Speech Commands mixture file with a corresponding spectrogram visualisation. The lower part contains an excerpt from the detection reports for three selected keywords. Blue regions indicate actual keyword occurrences with an acceptable detection collar. The main lanes below each keyword show similarity values for each of the provided keyword patterns (i.e. multiple rows per keyword). Darker values correspond to closer matches. Desaturated parts indicate discarded fragments with distances above the threshold. Orange markers denote possible detections. If a sufficiently long streak of detections is generated, actual occurrences are emitted – denoted either by a dot (correct detection) or a cross sign (incorrect).




Figure A2. Same detection report for the Google speech embedding model, with post-processing applied to the similarity values.

References

1. Thomas, S.; Suzuki, M.; Huang, Y.; Kurata, G.; Tuske, Z.; Saon, G.; Kingsbury, B.; Picheny, M.; Dibert, T.; Kaiser-Schatzlein, A.; others. English broadcast news speech recognition by humans and machines. ICASSP 2019-2019 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP). IEEE, 2019, pp. 6455–6459.
2. Amodei, D.; Ananthanarayanan, S.; Anubhai, R.; Bai, J.; Battenberg, E.; Case, C.; Casper, J.; Catanzaro, B.; Cheng, Q.; Chen, G.; others. Deep speech 2: End-to-end speech recognition in english and mandarin. International conference on machine learning. PMLR, 2016, pp. 173–182.
3. Koch, G.; Zemel, R.; Salakhutdinov, R.; others. Siamese neural networks for one-shot image recognition. ICML deep learning workshop. Lille, 2015, Vol. 2.
4. Snell, J.; Swersky, K.; Zemel, R.S. Prototypical networks for few-shot learning. *arXiv preprint arXiv:1703.05175* **2017**.
5. Wang, Y.; Salamon, J.; Bryan, N.J.; Bello, J.P. Few-shot sound event detection. ICASSP 2020-2020 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP). IEEE, 2020, pp. 81–85.
6. Lin, J.; Kilgour, K.; Roblek, D.; Sharifi, M. Training keyword spotters with limited and synthesized speech data. ICASSP 2020-2020 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP). IEEE, 2020, pp. 7474–7478.
7. Lee, K.; Hon, H.. Large-vocabulary speaker-independent continuous speech recognition using HMM. ICASSP-88., International Conference on Acoustics, Speech, and Signal Processing, 1988, pp. 123–126 vol.1. doi:10.1109/ICASSP.1988.196527.
8. Hannun, A.; Case, C.; Casper, J.; Catanzaro, B.; Diamos, G.; Elsen, E.; Prenger, R.; Satheesh, S.; Sengupta, S.; Coates, A.; Ng, A.Y. Deep Speech: Scaling up end-to-end speech recognition, 2014, [[arXiv:1412.5567](#)].
9. Battenberg, E.; Chen, J.; Child, R.; Coates, A.; Li, Y.G.Y.; Liu, H.; Satheesh, S.; Sriram, A.; Zhu, Z. Exploring neural transducers for end-to-end speech recognition. 2017 IEEE Automatic Speech Recognition and Understanding Workshop (ASRU). IEEE, 2017, pp. 206–213.
10. van den Oord, A.; Dieleman, S.; Zen, H.; Simonyan, K.; Vinyals, O.; Graves, A.; Kalchbrenner, N.; Senior, A.; Kavukcuoglu, K. WaveNet: A Generative Model for Raw Audio, 2016, [[arXiv:1609.03499](#)].
11. Xiong, W.; Wu, L.; Allewa, F.; Droppo, J.; Huang, X.; Stolcke, A. The Microsoft 2017 Conversational Speech Recognition System, 2017, [[arXiv:1708.06073](#)].
12. Li, J.; Lavrukhin, V.; Ginsburg, B.; Leary, R.; Kuchaiev, O.; Cohen, J.M.; Nguyen, H.; Gadde, R.T. Jasper: An end-to-end convolutional neural acoustic model. *arXiv preprint arXiv:1904.03288* **2019**.
13. Krizan, S.; Beliaev, S.; Ginsburg, B.; Huang, J.; Kuchaiev, O.; Lavrukhin, V.; Leary, R.; Li, J.; Zhang, Y. Quartznet: Deep automatic speech recognition with 1d time-channel separable convolutions. ICASSP 2020-2020 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP). IEEE, 2020, pp. 6124–6128.
14. Gulati, A.; Qin, J.; Chiu, C.C.; Parmar, N.; Zhang, Y.; Yu, J.; Han, W.; Wang, S.; Zhang, Z.; Wu, Y.; Pang, R. Conformer: Convolution-augmented Transformer for Speech Recognition, 2020, [[arXiv:2005.08100](#)].
15. Sainath, T.N.; Parada, C. Convolutional neural networks for small-footprint keyword spotting. Sixteenth Annual Conference of the International Speech Communication Association, 2015.
16. Lengerich, C.; Hannun, A. An end-to-end architecture for keyword spotting and voice activity detection. *arXiv preprint arXiv:1611.09405* **2016**.
17. Warden, P. Speech commands: A dataset for limited-vocabulary speech recognition. *arXiv preprint arXiv:1804.03209* **2018**.
18. Tang, R.; Lin, J. Honk: A pytorch reimplementation of convolutional neural networks for keyword spotting. *arXiv preprint arXiv:1710.06554* **2017**.
19. Zhang, Y.; Suda, N.; Lai, L.; Chandra, V. Hello edge: Keyword spotting on microcontrollers. *arXiv preprint arXiv:1711.07128* **2017**.
20. de Andrade, D.C.; Leo, S.; Viana, M.L.D.S.; Bernkopf, C. A neural attention model for speech command recognition. *arXiv preprint arXiv:1808.08929* **2018**.
21. Zeng, M.; Xiao, N. Effective combination of DenseNet and BiLSTM for keyword spotting. *IEEE Access* **2019**, 7, 10767–10775.
22. Lin, Z.Q.; Chung, A.G.; Wong, A. Edgespeechnets: Highly efficient deep neural networks for speech recognition on the edge. *arXiv preprint arXiv:1810.08559* **2018**.
23. Choi, S.; Seo, S.; Shin, B.; Byun, H.; Kersner, M.; Kim, B.; Kim, D.; Ha, S. Temporal convolution for real-time keyword spotting on mobile devices. *arXiv preprint arXiv:1904.03814* **2019**.
24. Mittermaier, S.; Kürzinger, L.; Waschneck, B.; Rigoll, G. Small-footprint keyword spotting on raw audio data with sinc-convolutions. ICASSP 2020-2020 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP). IEEE, 2020, pp. 7454–7458.
25. Majumdar, S.; Ginsburg, B. MatchboxNet–1D Time-Channel Separable Convolutional Neural Network Architecture for Speech Commands Recognition. *arXiv preprint arXiv:2004.08531* **2020**.
26. Li, B.; Wu, F.; Lim, S.N.; Belongie, S.; Weinberger, K.Q. On feature normalization and data augmentation. *arXiv preprint arXiv:2002.11102* **2020**.
27. Coucke, A.; Chlieh, M.; Gisselbrecht, T.; Leroy, D.; Poumeyrol, M.; Lavril, T. Efficient keyword spotting using dilated convolutions and gating. ICASSP 2019-2019 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP). IEEE, 2019, pp. 6351–6355.
28. Razi, A.; Hyun-Jin, P. End-to-end streaming keyword spotting. *arXiv preprint arXiv:1812.02802* **2018**.

29. Mazzawi, H.; Gonzalvo, X.; Kracun, A.; Sridhar, P.; Subrahmanya, N.; Lopez-Moreno, I.; Park, H.J.; Violette, P. Improving Keyword Spotting and Language Identification via Neural Architecture Search at Scale. *INTERSPEECH*, 2019, pp. 1278–1282.
30. Guo, J.; Kumatani, K.; Sun, M.; Wu, M.; Raju, A.; Ström, N.; Mandal, A. Time-delayed bottleneck highway networks using a DFT feature for keyword spotting. 2018 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP). IEEE, 2018, pp. 5489–5493.
31. Kubanek, M. Method of speech recognition and speaker identification using audio-visual of polish speech and hidden markov models. In *Biometrics, Computer Security Systems and Artificial Intelligence Applications*; Springer, 2006; pp. 45–55.
32. Ziółko, M.; Gałka, J.; Ziółko, B.; Jadczyk, T.; Skurzok, D.; Masior, M. Automatic speech recognition system dedicated for Polish. Twelfth Annual Conference of the International Speech Communication Association. Citeseer, 2011.
33. Pohl, A.; Ziółko, B. Using part of speech n-grams for improving automatic speech recognition of Polish. *International Workshop on Machine Learning and Data Mining in Pattern Recognition*. Springer, 2013, pp. 492–504.
34. Baevski, A.; Zhou, H.; Mohamed, A.; Auli, M. wav2vec 2.0: A framework for self-supervised learning of speech representations. *arXiv preprint arXiv:2006.11477* **2020**.
35. Conneau, A.; Baevski, A.; Collobert, R.; Mohamed, A.; Auli, M. Unsupervised cross-lingual representation learning for speech recognition. *arXiv preprint arXiv:2006.13979* **2020**.
36. Warden, P. Speech Commands: A Dataset for Limited-Vocabulary Speech Recognition. *ArXiv e-prints* **2018**, [arXiv:cs.CL/1804.03209].
37. Ardila, R.; Branson, M.; Davis, K.; Henretty, M.; Kohler, M.; Meyer, J.; Morais, R.; Saunders, L.; Tyers, F.M.; Weber, G. Common voice: A massively-multilingual speech corpus. *arXiv preprint arXiv:1912.06670* **2019**.
38. Panayotov, V.; Chen, G.; Povey, D.; Khudanpur, S. Librispeech: an asr corpus based on public domain audio books. 2015 IEEE international conference on acoustics, speech and signal processing (ICASSP). IEEE, 2015, pp. 5206–5210.
39. Baumann, T.; Köhn, A.; Hennig, F. The Spoken Wikipedia Corpus collection: Harvesting, alignment and an application to hyperlistening. *Language Resources and Evaluation* **2019**, 53, 303–329.
40. Żelasko, P.; Ziółko, B.; Jadczyk, T.; Skurzok, D. AGH corpus of Polish speech. *Language Resources and Evaluation* **2016**, 50, 585–601.
41. Koržinek, D.; Marasek, K.; Brocki, Ł.; Wołk, K. Polish read speech corpus for speech tools and services. *arXiv preprint arXiv:1706.00245* **2017**.
42. Pęzik, P. Increasing the Accessibility of Time-Aligned Speech Corpora with Spokes Mix. *Proceedings of the Eleventh International Conference on Language Resources and Evaluation (LREC 2018)*, 2018.
43. Pęzik, P. Spokes-a search and exploration service for conversational corpus data. *Selected papers from the CLARIN 2014 Conference*, 2014, pp. 99–109.
44. Demenko, G.; Grochowski, S.; Klessa, K.; Ogórkiewicz, J.; Wagner, A.; Lange, M.; Śledziński, D.; Cylwik, N. Jurisdic: Polish speech database for taking dictation of legal texts. *Proceedings of the Sixth International Conference on Language Resources and Evaluation (LREC'08)*, 2008.
45. Szwełnik, Tomasz.; Kawalec, Jacek.; Gutowska, Dorota. Polish Speech Database, 2019. doi:10.35111/TWQH-F096.
46. Polish & English language corpora for research & applications. <http://pelcra.pl/new/snuv>. Accessed: 2021-10-07.
47. Nagrani, A.; Chung, J.S.; Xie, W.; Zisserman, A. Voxceleb: Large-scale speaker verification in the wild. *Computer Science and Language* **2019**.
48. Mesaros, A.; Heittola, T.; Virtanen, T. Metrics for polyphonic sound event detection. *Applied Sciences* **2016**, 6, 162.
49. Kruspe, A. A simple method for domain adaptation of sentence embeddings. *arXiv preprint arXiv:2008.11228* **2020**.