

# THE ZERO RESOURCE SPEECH CHALLENGE 2017

Ewan Dunbar<sup>1</sup>, Xuan Nga Cao<sup>1</sup>, Juan Benjumea<sup>1</sup>, Julien Karadayi<sup>1</sup>, Mathieu Bernard<sup>1</sup>,  
Laurent Besacier<sup>2</sup>, Xavier Anguera<sup>3</sup>, Emmanuel Dupoux<sup>1</sup>

<sup>1</sup>ENS, EHESS, PSL Research University, CNRS, INRIA

<sup>2</sup> LIG, Univ. Grenoble Alpes, CNRS, Grenoble INP, France

<sup>3</sup> Elsa Corp., Lisbon, Portugal

## ABSTRACT

We describe a new challenge aimed at discovering subword and word units from raw speech. This challenge is the followup to the Zero Resource Speech Challenge 2015. It aims at constructing systems that generalize across languages and adapt to new speakers. The design features and evaluation metrics of the challenge are presented and the results of seventeen models are discussed.

**Index Terms**— zero resource speech technology, subword modeling, acoustic unit discovery, unsupervised term discovery

## 1. INTRODUCTION

Traditional speech technology relies on linguistic expertise and textual information to build acoustic and language models. Recent work has begun to mitigate the need for expert knowledge, having succeeded at training supervised systems using textual transcripts only [1, 2]. But for the majority of the world’s languages, even textual transcripts are sparse or non-existent, and for many languages, the lack of standard orthography would reduce the utility of such transcripts greatly. This is why there is increasing interest in a radically different approach, the so-called *zero resource setting* [3, 4]. The aim is to construct speech systems without any textual or linguistic resource. The fact that young children become fluent communicators in their native language(s) before they learn to read and write, and with minimal or no expert supervision, demonstrates that such a system is possible.

Systems constructed with zero expert resources have at least three kinds of applications. They can form the basis for speech services for the perhaps billions of speakers of languages without readily available linguistic or textual resources in sufficient quantities. They can also help field linguists document endangered languages, by providing tools to semi-automatically analyze and annotate audio recordings using automatically discovered linguistic units (phonemes, lexicon, grammar) [5]. In this case, a realistic scenario is to process raw speech and its translation in a well-resourced language [6, 7, 8]. Finally, to the extent that they mimic language

development, these models can be used to evaluate scientific claims about development, and act as quantitative tools for psychologists and clinicians interested in the impact of sociolinguistic variations in input on subsequent normal or abnormal language and cognitive development [9].

As the previous Zero Resource Speech Challenge [10, 11], the 2017 installment consists of two different subchallenges. Participants can submit systems that do either *subword modelling* (track one) or *spoken term discovery* (track two). Subword modelling means constructing a representation of speech sounds; we say “subword” instead of “phoneme” or “sound” because the criterion for success does not demand any particular temporal window size. The requirement is to have a representation of the speech signal that is robust to talker variation, and makes the distinctions necessary to distinguish words. Spoken term discovery means discovering recurring speech fragments in the audio: the system must output time-stamps delimiting stretches of speech, associated with class labels, ideally corresponding to real words in the language. The systems must be able to do their task without linguistic resources. The 2017 Zero Resource Speech Challenge<sup>1</sup> has two main innovations compared to the 2015 iteration: (1) it tests how well systems’ architectures and hyperparameters generalize to unseen languages (the systems are developed with three languages and tested on two ‘surprise’ languages); (2) it tests how well the trained systems’ parameters adapt to unseen talkers (varying the amount of speech available for each talker).

## 2. THE CHALLENGE

### 2.1. Track one: Subword modelling

The goal of subword modelling is to learn representations appropriate to the sounds of the language (phones or phonemes). As adults, we easily perceive differences between individual phonemes in our native language, substantially better than in non-native languages [12]. This so-called *perceptual tuning* develops substantially during the first year of life [13, 14].

<sup>1</sup><http://zerospeech.com/2017>

The input to subword modelling is a set of acoustic features encoding continuous speech. The result is a transformation that applies to new samples of running speech, yielding a new representation (continuous or discrete) of that speech. After training in a particular language from unlabelled speech, the representation is computed for a test set in the same language on novel speakers, to be evaluated on phoneme discrimination based on gold-standard annotations.

## 2.2. Track two: Spoken term discovery

Intuitively, the goal of spoken term discovery is to find words in the speech stream—just as the infant learns the words of its language by listening. The input to candidate systems is a series of speech features. The output is a set of boundaries delimiting the start and end of proposed word tokens discovered in the speech, and category labels indicating proposed word types. These boundaries may, but need not, constitute an exhaustive parse of the speech. The evaluation we apply is a set of scores measuring different aspects of the alignment with the words in the gold-standard transcription. As is customary in the field of word segmentation, we do not provide a separate test set for this track; we rely on the surprise languages to assess possible hyperparameter overfitting.

## 2.3. Generalization across languages

This challenge is explicitly set up as a ‘learning to learn’ or meta-learning problem: participants select the best architectures and tune their hyperparameters on three development languages on which the evaluation software is provided for a test set. The higher level test is performed on two new surprise languages for which the evaluation is withheld. Participants submit the output of their previously optimized systems for these two corpora, and the organizers perform the evaluation. Each submission (maximum five per team) is automatically evaluated and made public.

## 2.4. Adaptation to new speakers

The subword modeling task is explicitly set up as a talker adaptation problem. The training data is constructed from a skewed distribution of speakers, and the test set consists only of new speakers appearing in files of various durations (from 1 second to two minutes). This enables us to evaluate models’ ability to generalize and adapt to new speakers on the basis of limited evidence.

# 3. DATA SETS

Two groups of data sets are provided as part of the challenge: the *development data* and the *surprise data*. The development data consists of corpora from three different languages (English, French and Mandarin). Each corpus comes with

software that does the evaluation for the two tracks. Challenge participants are encouraged to use these resources to tune their hyperparameters using a cross validation approach to maximize generalizability. Each corpus is split into a train and a test set; the subword modelling evaluations are run on the test sets, and the spoken term discovery evaluations on the training sets. The participants then have to submit their systems and their output on all the data sets for independent evaluation (run automatically upon submission). The surprise data consists of corpora from two new languages, the identity of which was not revealed to participants. Participants receive only the speech corpora, with no additional resources. As before, each contains a training set (which participants can use to train their systems) and a test set for track one evaluation.

The amount of data in the training part of the development data sets varies from 2.5 to 45 hours (see Table 1), to ensure that systems can work both with limited data and with large data sets. The statistics of the two surprise languages fall between these two extremes. The distribution of speakers in the training sets is shaped to reflect what is typically found in natural language acquisition settings: there is a “family”—a small number of speakers that make up a large proportion of the total speech—and a set of “outsiders”—a larger number of speakers that each appear in smaller proportions. The test sets consist of many short files, are organized into subsets of differing length (1s, 10s and 120s).

The English and French corpora were taken from the LibriVox collection of audio books<sup>2</sup> and phone force-aligned using Kaldi [15]. The Mandarin corpus is the one described in [16], force-aligned using Kaldi. The first surprise language is German, and was taken from LibriVox and force-aligned using Kaldi as well. The second surprise language is Wolof, and is the corpus described in [17].

# 4. EVALUATION

## 4.1. Evaluation of subword modelling

Unsupervised subword modeling is a representation learning problem: the system must find speech features that emphasize linguistically relevant properties (phoneme structure) and de-emphasize the linguistically irrelevant ones (speaker identity, emotion, channel, etc). We use a minimal pair ABX task [18, 19], which does not require any training, and only requires a dissimilarity metric between speech tokens. The ABX task is inspired by match-to-sample tasks used in human psychophysics and is a simple way to measure discriminability between two sound categories (where the sounds  $A$  and  $B$  belong to different categories  $x$  and  $y$ , respectively, and the task is to decide whether the sound  $X$  belongs to one or the other). We define the ABX-discriminability of category  $x$  from category  $y$  as the probability that  $A$  and  $X$  are further apart than  $B$  and  $X$  when  $A$  and  $X$  are from category  $x$  and  $B$

<sup>2</sup><http://librivox.org/>

**Table 1.** Corpus statistics

	Training						Test	
	Relatives		Outsiders		Total		number of files	duration (min)
	#speakers	duration/ speaker	#speakers	duration/ speaker	duration	#words		
<b>Development</b>								
English	9	165-220min	60	10min	45h	370k	30658	1634
French	10	110-195min	18	10min	24h	220k	23765	1061
Mandarin	4	20-25min	8	10min	2h30	20k	25383	1522
<b>Surprise</b>								
S1 – German	10	85-150min	20	10min	25h	213k	15243	687
S2 – Wolof	4	37-42min	10	10min	4h	31k	7201	354

**Table 2.** Track 1 within and between talker ABX scores for three test files durations, for the development (average) and surprise languages, for the 14 submitted systems (the mean performance on the surprise set is in square brackets). The three supervised systems are listed separately. See <http://zerospeech.com/2017> for the full table including development languages.

		Avg. Development			S1 – German			S2 – Wolof		
		1 s	10 s	2 min	1 s	10 s	2 min	1 s	10 s	2 min
<b>Baseline</b> [18.84]	Within	12	12.1	12.1	14.2	13.3	12.9	14.1	14.3	14.1
	Between	23.3	23.4	23.3	27.5	27.3	27.1	30	29.5	29.5
<b>Topline</b> [7.19]	Within	8	5.4	5.3	8.7	7.1	7.0	6.6	4.6	3.4
	Between	9.6	7.2	6.9	12.8	10.5	10.4	7.1	3.6	4.3
<b>H</b> [9.16]	Within	8.5	7.6	7.4	6.5	5.6	5.3	10.9	8.8	8.4
	Between	10.8	9.3	9.0	11.9	10.0	9.7	13.0	10.0	9.9
<b>P1</b> [14.96]	Within	10.9	8.8	8.7	8.9	6.7	6.4	13.3	11.9	11.8
	Between	17.5	15.8	15.7	19.4	16.2	15.9	22.8	23.1	23.1
<b>P2</b> [14.9]	Within	10.8	8.7	8.5	8.8	6.6	6.3	13.1	11.7	11.7
	Between	17.5	15.8	15.7	19.2	16.3	16.0	23.3	23.3	23.1
<b>C1</b> [11.95]	Within	10.0	8.4	8.3	7.6	6.22	6.3	11.7	9.9	9.8
	Between	14.5	12.8	12.5	15.5	12.9	12.7	17.6	16.9	16.3
<b>C2</b> [11.24]	Within	10.1	8.5	8.4	7.6	6.2	6.1	11.6	9.8	9.6
	Between	13.9	11.9	11.7	14.7	11.7	11.6	16.9	14.7	14.4
<b>A1</b> [12.3]	Within	8.8	7.9	7.8	6.9	6.1	6.0	9.9	9.2	9.1
	Between	15.2	14.2	14.0	16.9	14.7	14.7	18.8	17.7	17.7
		Avg. Development			S1 – German			S2 – Wolof		
		1 s	10 s	2 min	1 s	10 s	2 min	1 s	10 s	2 min
<b>A2</b> [11.95]	Within	8.8	7.8	7.7	6.8	6.0	6.0	10.1	9.6	9.6
	Between	14.5	13.4	13.3	16.0	14.0	13.9	17.9	16.9	16.6
<b>A3</b> [11.93]	Within	9.5	8.3	8.2	7.3	6.2	6.1	11.1	10.3	10.2
	Between	14.5	13.3	13.2	15.5	13.5	13.4	17.6	16.0	16.0
<b>A4</b> [12.17]	Within	9.7	8.5	8.5	7.6	6.4	6.2	11.6	10.9	10.7
	Between	14.5	13.3	13.2	15.7	13.7	13.5	17.5	16.1	16.1
<b>Y1</b> [12.43]	Within	10.8	8.4	8.4	8.1	6.0	6.0	12.6	10.0	9.9
	Between	15.3	13	12.6	16.2	12.9	12.6	19.5	17.1	16.6
<b>Y2</b> [12.29]	Within	10.7	8.4	8.2	8.2	6.2	6.2	12.7	10.1	9.9
	Between	15.1	12.7	12.4	16.4	13.3	13.0	19.2	17.3	16.7
<b>YS</b> [12.32]	Within	10.7	8.3	8.1	8.0	6.0	5.9	12.9	10.8	10.6
	Between	14.7	12.3	12.1	15.8	12.4	12.3	18.7	17.4	17.0
<b>S1</b> [8.54]	Within	8.7	8.3	7.3	6.3	5.8	5.0	9.0	8.7	7.2
	Between	11.4	10.4	9	11.6	9.9	8.7	11.5	10.2	8.6
<b>S2</b> [6.99]	Within	7.1	6.7	6.3	5.2	4.9	4.5	6.9	7.0	6.3
	Between	9.0	8.6	7.8	9.3	8.6	7.8	8.3	7.9	7.2

is from category  $y$ , according to a dissimilarity function  $d$ . We obtain a symmetric discriminability score by taking the average of the ABX discriminability of  $x$  from  $y$  and of  $y$  from  $x$ . The dissimilarities provided in this challenge are based on dynamic time warping, the underlying frame-level dissimilarity being either the cosine distance or KL-divergence. For most systems (signal processing, embeddings) the cosine distance usually gives good results, and for others (posteriorgrams) the KL divergence is more appropriate. Participants are allowed to supply their own dissimilarity as long as it was not obtained through supervised training.

We focus on phone triplet minimal pairs: sequences of 3 phonemes that differ in the central sound (not necessarily real words, e.g., “beg”–“bag”, “api”–“ati”, etc). Our compound measure sums over all minimal pairs of this type found in the corpus in a structured manner, that depends on the task. For the *within-speaker* task, all of the phone triplets belong to the same speaker (e.g.  $A = \text{beg}_{T_1}$ ,  $B = \text{bag}_{T_1}$ ,  $X = \text{bag}'_{T_1}$ ). The scores for a given minimal pair are first averaged across all of the speakers for which this minimal pair exists. The resulting scores are then averaged over all found contexts for a given pair of central phones (e.g. for the pair /a/-/e/, average the scores for the existing contexts such as /b\_g/, /r\_d/, /f\_s/, etc.). Finally the scores for every pair of central phones are averaged and subtracted from 1 to yield the reported within-talker ABX error rate. For the *across-speaker*

task,  $A$  and  $B$  belong to the same speaker, and  $X$  to a different one.  $A = \text{beg}_{T_1}$ ,  $B = \text{bag}_{T_1}$ ,  $X = \text{bag}_{T_2}$ . The scores for a given minimal pair are first averaged across all of the pairs of speakers for which this contrast can be made. As above, the resulting scores are averaged over all contexts over all pairs of central phones and converted to an error rate.

## 4.2. Evaluation of spoken term discovery

“Spoken term discovery” is a broad label. We use it to refer to systems that do one or more of three tasks. Systems that do *matching* find fragments that are matched pairwise as being instances of the same sequence of phonemes, attempting to find as many as possible. They can be evaluated based on how similar the matched fragments are, and how much of the corpus they cover. Systems that do *lexicon discovery* group these fragments into sets (rather than just matching pairwise), with the goal of finding a high-quality lexicon of types. They can be evaluated based on how well the sets match on the sequence of phonemes, and how well the sets match the gold-standard lexicon of word types. Systems that do *word segmentation* attempt to find fragments that are aligned with the gold-standard word transcription. By setting out three different types of criteria, the intention is to be open to various types of “spoken term discovery” systems, all of which in some sense “find words.” The result is that we do three

(non-independent) types of evaluations in track two. All of these evaluations are done at the level of the phonemes. Using the aligned phoneme transcription, we convert any discovered fragment of speech into its transcribed string. If the left or right edge of the fragment contains part of a phoneme, that phoneme is included in the transcription if it corresponds to more than 30ms or more than 50% of its duration.

#### 4.2.1. Evaluation as matching system

The evaluation of spoken term discovery systems as matching systems consists of two scores, **NED** and **coverage**. **NED**, is the average, over all matched pairs, of the Levenshtein distance between their phonemic transcriptions, divided by the max of their phonemic length. The second score is the **coverage**, the fraction of the discoverable part of the corpus that is covered by all the discovered fragments. The discoverable part of the corpus is found by computing the union of all of the intervals corresponding to all of the pairs of ngrams (with  $n$  between 3 and 20). This is almost all of the corpus, except for unigram and bigram hapaxes.

#### 4.2.2. Evaluation as lexicon discovery system

Six scores are used to evaluate the performance of a spoken term discovery system in terms of lexicon discovery. The first three are **grouping precision**, **recall** and **F-score**. These are defined in terms of  $P_{\text{clus}}$ , the set of all pairs of fragments that are grouped in the same cluster, and  $P_{\text{goldclus}}$ , the set of all non-overlapping pairs of fragments which are both discovered by the system (not necessarily in the same cluster) and have exactly the same gold transcription.

$$\text{Prec: } \sum_{t \in \text{types}(P_{\text{clus}})} w(t, P_{\text{clus}}) \frac{|\text{occ}(t, P_{\text{clus}} \cap P_{\text{goldclus}})|}{|\text{occ}(t, P_{\text{clus}})|} \quad (1)$$

$$\text{Rec: } \sum_{t \in \text{types}(P_{\text{goldclus}})} w(t, P_{\text{goldclus}}) \frac{|\text{occ}(t, P_{\text{clus}} \cap P_{\text{goldclus}})|}{|\text{occ}(t, P_{\text{goldclus}})|} \quad (2)$$

Where  $t$  ranges over the types of fragments (defined by the transcription) in a cluster  $C$ , and  $\text{occ}(t, C)$  is the number of occurrences of that type,  $w$  the number of occurrences divided by the size of the cluster. In other words, Prec is a weighted measure of cluster purity and Rec, of the inverse of the cluster's fragmentation. The other three scores are **type precision**, **recall**, and **F-score**. Type precision is the probability that discovered types belong to the gold set of types (real words), whereas type recall is the probability that gold types are discovered. We restrict both sets to words between three and twenty segments long.

#### 4.2.3. Evaluation as word segmentation system

Six scores are used to evaluate the performance of a spoken term discovery system as a system for word segmentation.

The first three are **token precision**, **recall**, and **F-score**. Token precision is the probability that discovered fragment tokens are in the gold set of word fragments, and token recall the probability that the gold fragments are discovered. Again, tokens are restricted to be between three and twenty segments long. Finally, **boundary precision**, **recall**, and **F-score** are similarly defined with respect to the set of discovered and gold word boundaries.

## 5. DESCRIPTION OF SUBMITTED SYSTEMS

We received sixteen systems for track one, and three for track two. For two track one submissions, the authors declared that the systems were not intended as submissions, and are thus not included in the present overview, although the results can be found on the leaderboard on the Challenge website ([www.zerospeech.com/2017](http://www.zerospeech.com/2017)).

### 5.1. Track 1

The baseline features are the MFCC features provided to participants; they are evaluated using the ABX score using the cosine similarity. The topline system is a supervised HMM-GMM phone recognition system, with a bigram language model, trained with a Kaldi recipe [15]. The phone posteriors are evaluated using the KL divergence.

The submitted systems follow four strategies. The first strategy consists in performing a bottom-up frame-level clustering, following up on the success of [20] with this simple strategy. In system [H], from Heck et al. [21], PLP representations are used, as well as the output of three learned feature transformations, to be the input to the same frame-level clustering algorithm used in [20] (the parallel Dirichlet process Gaussian mixture model sampler of [22]). The transformations (LDA, LDA+MLLT; LDA+MLLT+fMLLR) are aimed at making the representation talker invariant. The training requires labels, which are obtained from an initial clustering using the DPGMM algorithm. The submitted system is a combination of the raw posteriors (not smoothed by any language model) from the four DPGMMs obtained by clustering frames in each of the three transformed representations, plus the raw PLP features. System [P1], from Pellegrini et al. [23], uses  $k$ -means clustering on ZCA-transformed (whitened using PCA and rotated back to the original axes) MFCC features. In system [P2], this is followed by a step of centroid re-estimation: the centroids are re-estimated based only on the data points for which the label is the same as that of its left and right neighbors. The distance from point to all cluster centroids is the feature vector in both.

A second strategy constructs language-independent embeddings by training neural networks in an unsupervised way. The systems of Chen et al. [24] first cluster frames separately on each language, using the same DPGMM algorithm as [21] and [20]. The cluster labels are then used as tar-

**Table 3.** Track 2 results. Scores are averages over evaluations done on subparts of the corpus: the list of corpus gold utterances is split in groups of 1000 elements/utterances in order to compute the scores in parallel. Thus, the  $F$ -scores presented are not directly related to the precision and recall scores.

					Grouping			Type			Token			Boundary		
	Words	Pairs	NED	Cov	Pr	Re	F	Pr	Re	F	Pr	Re	F	Pr	Re	F
	Avg. Training (English, French, Mandarin)															
Baseline	4948	5842	30.0	4.1	48.7	86.6	52.3	5.5	0.3	0.6	4.4	0.2	0.2	34.1	1.5	2.9
Topline	5779	2742085	0.0	100.0	99.9	100.0	100.0	41.3	45.9	43.4	47.8	58.4	52.2	77.9	94.4	84.9
K	24724	3775685	76.0	105.0	6.7	8.3	6.7	4.6	9.0	6.1	6.3	7.1	6.7	41.7	47.6	44.4
G1	51377	465997	71.5	60.5	3.9	67.8	7.4	3.9	6.0	4.4	3.1	2.7	2.9	25.3	31.8	27.8
G2	51354	467369	71.4	60.4	4.0	67.6	7.5	3.9	6.0	4.4	3.1	2.7	2.8	25.2	31.8	27.7
S1 – German																
Baseline	2973	3315	30.5	3.0	54.8	94.6	64.9	5.5	0.3	0.6	4.0	0.1	0.2	28.2	1.2	2.3
Topline	7664	2588808	0.0	100.0	100.0	100.0	100.0	28.4	29.5	29.0	42.3	62.7	50.5	70.0	98.3	81.8
K	28675	4258731	66.4	100.0	11.8	5.9	7.9	5.7	11.2	7.5	10.3	14.3	12.0	42.6	56.5	48.6
G1	60648	582009	59.9	71.8	5.7	63.8	10.4	3.2	6.8	4.3	2.4	3.1	2.7	20.6	37.2	26.6
G2	60489	588162	59.9	71.8	5.8	64.0	10.7	3.2	6.8	4.3	2.4	3.1	2.7	20.6	37.2	26.5
S2 – Wolof																
Baseline	462	545	33.5	3.2	39.1	72.1	32.8	2.3	0.1	0.2	1.6	0.0	0.1	25.3	1.0	2.0
Topline	2472	602339	0.0	100.0	100.0	100.0	100.0	43.0	47.6	45.2	55.6	65.6	60.2	81.3	93.2	86.9
K	3593	439321	72.2	100.0	4.7	5.4	5.0	4.6	10.0	6.3	4.9	5.2	5.0	42.4	44.3	43.3
G1	5468	37191	56.8	47.8	7.6	43.3	12.8	5.9	8.2	6.9	6.1	4.2	4.9	29.6	29.6	29.6
G2	5460	37273	56.8	47.8	7.7	44.0	13.0	6.0	8.3	7.0	6.1	4.2	4.9	29.7	29.6	29.6

gets for a multitask neural network, where the loss function is the mean of the five cross-entropy losses. A linear bottleneck layer in the network is the final representation. System [C1] is trained on MFCCs; [C2] transforms the input features using unsupervised linear VTLN. The three systems in **Ansari et al.** [25] combine two sets of features, also trained on all five languages. The first set of features is always a high-dimensional hidden layer from an autoencoder trained on the MFCC frames. The second set of features is a hidden layer from a network trained to predict unsupervised frame labels. These labels are obtained by training a Gaussian mixture model on the speech frames. In system [A1], the cluster labels are used as targets for the DNN trained with the MFCCs as input. System [A2] instead uses time-smoothed frame labels as decoded by a large Gaussian-mixture-HMM (one state for each of the initial cluster labels), in which states have a high self-transition probability. System [A3] does the same but trains the DNN with the autoencoder features as inputs. In systems [A1] and [A2], the DNN features are combined with the autoencoder features in the dissimilarity function for the ABX task: the dissimilarity is a weighted average between the cosine DTW distance using the autoencoder features, and the cosine DTW distance using the other set of features (the weights are optimized on the development languages and not reported). An additional system is described in [26], system [A4], which consists of the DNN features from system [A2].

The third strategy is to use spoken term discovery to improve the acoustic features, following [27, 28]. The system of **Yuan et al.** [29] generates bottle-neck features in the same way as [24], but then applies the STD system of [30] to discover pairs of matched acoustic motifs, which are then used to further train the features. Specifically, following [28], a stacked autoencoder is trained on DTW-aligned frame pairs from instances of motifs paired by the STD system. System

[Y1] uses pairs obtained by applying the STD system (only) to the English training data. System [Y2] (not reported in [24]) obtained pairs from all corpora of all five languages. System [YS] is a supervised comparison, using transcribed pairs from the Switchboard corpus rather than STD.

The strategy of **Shibata et al.** [31] is to use supervised training on out-of-domain languages. In [S1], the features are bottleneck features from a neural network acoustic model trained supervised on Japanese as part of an HMM. In [S2], an end-to-end convolutional network plus bidirectional LSTM is trained using corpora from ten languages, and features from a hidden layer are combined with the bottleneck Japanese features. This second scheme can only be considered out of domain in a weak sense: new corpora are used, but the ten additional languages include English, Mandarin, and German.

## 5.2. Track 2

The baseline of track two was computed using [30], which does pair-matching using locally sensitive hashing applied to PLP features and then groups pairs using graph clustering. The parameters of the STD stayed the same across all languages, except that the DTW threshold was increased for Mandarin (to 0.90, from 0.88) in order to obtain a NED value similar to that of other languages. The topline was an exhaustive-parsing word segmentation model based on the textual transcriptions (a unigram grammar trained in the adaptor grammar framework: [32]).

System [K], from **Kamper et al.** [33] uses  $k$ -means to discover recurring acoustic patterns, jointly optimized with an exhaustive segmentation. The variable-length acoustic chunks of speech yielded by a segmentation are reduced to fixed dimension by down-sampling in the time dimension for clustering by  $k$ -means. The systems of **García-Granada**

**et al.** [34] use a supervised ASR system for Hungarian to decode the speech. Chunks of (transcribed) speech match if they have the same transcription. These matches are then filtered: a representation of the speech is obtained by training an autoencoder, and only pairs with DTW sufficiently low in this representational space (below a threshold) are retained (system [G1]). In system [G2], pairs are tolerated that have a one-phone difference in their transcriptions.

## 6. RESULTS

The results of track one are summarized in Table 2. The clear winner among the unsupervised models is [H]. While the paper focuses on the role of learned feature transformations—and demonstrates that they improve the evaluation—these are unlikely to be solely responsible for the very good performance of the system. The baseline system described in the paper, with no transformations applied, consisting of posteriorgrams from a frame-wise DPGMM clustering, still performs extremely well. Although it was not evaluated on the surprise languages, its average ABX score on the development languages (9.7) is still better than that of the second-ranked system [C2] (10.7), from **Chen et al.** The systems presented by **Heck et al.** all have in common monolingual training, direct evaluation of GMM posteriorgrams, and the fact that they use PLP features as input, rather than MFCCs.

The systems of **Chen et al.** (as well as **Yuan et al.**) use an initial stage of monolingual training, followed by multilingual training. Among the other systems in [24] (not submitted for final evaluation) are some with fully monolingual training, which serve to partly evaluate the effect of multilingual training: monolingual training fares better, within language. However, as the numerical scores are not presented, it is not possible to determine whether the use of monolingual training could account for the good scores of **Heck et al.** What is clear is that multilingual training still works reasonably well—all the remaining systems, except for the lower performing system [P], use multilingual training, and show substantial improvements over the baseline. The setting for these systems corresponds to the situation of an infant raised in a multilingual environment.

The STD strategy for improving subword features fares relatively poorly; a limitation of the **Yuan et al.** systems is that they use an architecture (autoencoders) that was previously reported to be suboptimal for improving acoustic features [28]. A Siamese network [27] might improve the result.

The results of track two are in Table 3. As in the previous challenge, the baseline provided by the system of [30] has a high NED (very good matching), but has a low coverage (few matches). Given that system [K] demands an exhaustive parse, it is guaranteed to have full coverage, which risks poorer matches. It is thus notable that system [K] has a NED on the development languages comparable to the partly supervised systems [G1] and [G2]. On the other hand, grouping

by supervised transcription clearly works better than  $k$ -means clustering; while the groups have low precision (many erroneously grouped items) for both kinds of systems, [G1]/[G2] have much higher grouping recall (more of the true group members among the items grouped together) than system [K]. The quality of segmentation, in terms of tokens and of boundaries, is, on the other hand, much better in system [K], presumably as a result of increased coverage.

## 7. CONCLUSION

The results of the 2017 Zero Resource Speech Challenge build very clearly on the previous challenge. The strategies most successful in the previous challenge for track one, bottom-up clustering and training based on pairs of matching words discovered using spoken term discovery—are built upon, and a new strategy, multilingual training, is introduced. This strategy, although not completely foreseen by the organizers, would be similar to the situation of a multilingual infant. This new approach is used in conjunction with additional tools, including further training using neural networks and the application of time series models. Interestingly, simple bottom-up clustering still seems to work best in the unsupervised case. The fact that supervised labels trained fully out of domain seem to yield good results in both track one and track two (on Japanese, in the case of [S1], and on Hungarian in the case of [G1] and [G2]) indicates that the approach of having systems with strong prior knowledge about human speech is still a benchmark to be beaten in the unsupervised case. (The excellent results of [S2] may yet be partly attributable to the use of supervised training in German, which was one of the surprise languages.) Further work is needed to systematically evaluate the effect of each of these different manipulations, however, and to directly compare these new approaches with those from the previous challenge.

Potential future lines of inquiry building on the new track two systems might include exploring using these systems to feed subword modelling, especially in combination with some of the new subword modelling strategies used here, in light of the successful use of STD to improve subword models in the previous challenge, and the reasonable improvement over baseline shown in this challenge. In the other direction, the major unsupervised improvements to the speech features demonstrated in track one suggest strong room for improvement in unsupervised spoken term discovery systems like [K].

The 2017 ZeroSpeech Challenge remains open for submissions exploring these, or any other new strategies for discovering phoneme- and word-like units without access to linguistic resources.

## 8. REFERENCES

- [1] Awni Hannun, Carl Case, Jared Casper, Bryan Catanzaro, Greg Diamos, Erich Elsen, Ryan Prenger, Sanjeev Satheesh, Shubho Sengupta, Adam Coates, et al., “Deep speech: Scaling up end-to-end speech recognition,” *arXiv preprint arXiv:1412.5567*, 2014.
- [2] Yajie Miao, Mohammad Gowayyed, and Florian Metze, “Eesen: End-to-end speech recognition using deep rnn models and wfst-based decoding,” in *Automatic Speech Recognition and Understanding (ASRU), 2015 IEEE Workshop on*. IEEE, 2015, pp. 167–174.
- [3] James Glass, “Towards unsupervised speech processing,” in *Information Science, Signal Processing and their Applications (ISSPA), 2012 11th International Conference on*. 2012, pp. 1–4, IEEE.
- [4] Aren Jansen, Emmanuel Dupoux, Sharon Goldwater, Mark Johnson, Sanjeev Khudanpur, Kenneth Church, Naomi Feldman, Hynek Hermansky, Florian Metze, Richard Rose, Michael Seltzer, Pascal Clark, Ian McGraw, Balakrishnan Varadarajan, Erin Bennett, Benjamin Borschinger, Justin Chiu, Ewan Dunbar, Abdellah Fourtassi, David Harwath, Chia-ying Lee, Keith Levin, Atta Norouzian, Vijayaditya Peddinti, Rachael Richardson, Thomas Schatz, and Samuel Thomas, “A summary of the 2012 JH CLSP Workshop on zero resource speech technologies and models of early language acquisition,” in *Proceedings of ICASSP 2013*, 2013.
- [5] Gilles Adda, Sebastian Stücker, Martine Adda-Decker, Odette Ambourou, Laurent Besacier, David Blachon, Hélène Bonneau-Maynard, Pierre Godard, Fatima Hamlaoui, Dmitri Idiatov, Guy-Noël Kouarata, Lori Lamel, Emmanuel-Moselly Makasso, Annie Rialland, Mark Van de Velde, François Yvon, and Sabine Zerbian, “Breaking the unwritten language barrier: The Bulb project,” in *Proceedings of SLTU (Spoken Language Technologies for Under-Resourced Languages)*, Yogyakarta, Indonesia, 2016.
- [6] Long Duong, Antonios Anastasopoulos, David Chiang, Steven Bird, and Trevor Cohn, “An attentional model for speech translation without transcription,” in *Proceedings of NAACL-HLT*, 2016, pp. 949–959.
- [7] Alexandre Bérard, Olivier Pietquin, Christophe Servan, and Laurent Besacier, “Listen and translate: A proof of concept for end-to-end speech-to-text translation,” in *NIPS Workshop on end-to-end learning for speech and audio processing*, 2016.
- [8] Laurent Besacier, Bowen Zhou, and Yuqing Gao, “Towards speech translation of non written languages,” in *Spoken Language Technology Workshop, 2006. IEEE*, 2006, pp. 222–225.
- [9] Emmanuel Dupoux, “Cognitive science in the era of artificial intelligence: A roadmap for reverse-engineering the infant language-learner,” *arXiv preprint arXiv:1607.08723*, 2016.
- [10] Maarten Versteegh, Roland Thiollie, Thomas Schatz, Xuan Nga Cao, Xavier Anguera, Aren Jansen, and Emmanuel Dupoux, “The zero resource speech challenge 2015,” in *Proc. of Interspeech*, 2015.
- [11] Maarten Versteegh, Xavier Anguera, Aren Jansen, and Emmanuel Dupoux, “The zero resource speech challenge 2015: Proposed approaches and results,” *Procedia Computer Science: Proceedings of SLTU 2016*, vol. 81, pp. 67–72, 2016.
- [12] Leigh Lisker and Arthur S Abramson, “A cross-language study of voicing in initial stops: Acoustical measurements,” *Word*, vol. 20, no. 3, pp. 384–422, 1964.
- [13] Janet F Werker and Richard C Tees, “Cross-language speech perception: Evidence for perceptual reorganization during the first year of life,” *Infant behavior and development*, vol. 7, no. 1, pp. 49–63, 1984.
- [14] Patricia K Kuhl, “Human adults and human infants show a perceptual magnet effect for the prototypes of speech categories, monkeys do not,” *Attention, Perception, & Psychophysics*, vol. 50, no. 2, pp. 93–107, 1991.
- [15] Daniel Povey, Arnab Ghoshal, Gilles Boulianne, Lukas Burget, Ondrej Glembek, Nagendra Goel, Mirko Hanemann, Petr Motlicek, Yanmin Qian, Petr Schwarz, Jan Silovsky, Georg Stemmer, and Karel Vesely, “The kald speech recognition toolkit,” in *IEEE 2011 Workshop on Automatic Speech Recognition and Understanding*, 2011.
- [16] Dong Wang, Xuewei Zhang, and Zhiyong Zhang, “Thchs-30: A free chinese speech corpus,” *arXiv preprint arXiv:1512.01882*, 2015.
- [17] Elodie Gauthier, Laurent Besacier, Sylvie Voisin, Michael Melese, and Uriel Pascal Elingui, “Collecting Resources in Sub-Saharan African Languages for Automatic Speech Recognition: a Case Study of Wolof,” *LREC*, 2016.
- [18] Thomas Schatz, Vijayaditya Peddinti, Francis Bach, Aren Jansen, Hynek Hermansky, and Emmanuel Dupoux, “Evaluating speech features with the Minimal-Pair ABX task (I): Analysis of the classical MFC/PLP pipeline,” in *INTERSPEECH*, 2013.

- [19] Thomas Schatz, Vijayaditya Peddinti, Xuan-Nga Cao, Francis Bach, Hynek Hermansky, and Emmanuel Dupoux, "Evaluating speech features with the Minimal-Pair ABX task (II): Resistance to noise," in *INTER-SPEECH*, 2014.
- [20] Hongjie Chen, Cheung-Chi Leung, Lei Xie, Bin Ma, and Haizhou Li, "Parallel inference of Dirichlet process Gaussian mixture models for unsupervised acoustic modeling: A feasibility study," in *INTERSPEECH*, 2015.
- [21] Michael Heck, Sakriani Sakti, and Satoshi Nakamura, "Feature optimized dpghmm clustering for unsupervised subword modeling: A contribution to zerospeech 2017," *Submitted to ASRU 2017*.
- [22] Jason Chang and John W Fisher III, "Parallel sampling of DP mixture models using sub-cluster splits," in *Advances in Neural Information Processing Systems*, 2013, pp. 620–628.
- [23] Thomas Pellegrini, Céline Manenti, and Julien Pinquier, "Unsupervised discovery of sub-lexical units in speech based on ZCA and k-means," *Submitted to ASRU 2017*.
- [24] Hongjie Chen, Cheung-Chi Leung, Lei Xie, Bin Ma, and Haizhou Li, "Multilingual bottle-neck feature learning from untranscribed speech," *Submitted to ASRU 2017*.
- [25] T K Ansari, Sonali Singh, Rajath Kumar, and Sriram Ganapathy, "Deep learning methods for unsupervised acoustic modeling: LEAP submission to ZeroSpeech challenge 2017," in *Submitted to ASRU 2017*.
- [26] T K Ansari, Rajath Kumar, Sonali Singh, Sriram Ganapathy, and Susheela Devi, "Unsupervised HMM posteriors for language independent acoustic modeling in zero resource conditions," in *Submitted to ASRU 2017*.
- [27] Roland Thiollere, Ewan Dunbar, Gabriel Synnaeve, Maarten Versteegh, and Emmanuel Dupoux, "A hybrid dynamic time warping-deep neural network architecture for unsupervised acoustic modeling," in *INTER-SPEECH*, 2015, pp. 3179–3183.
- [28] Daniel Renshaw, Herman Kamper, Aren Jansen, and Sharon Goldwater, "A comparison of neural network methods for unsupervised representation learning on the zero resource speech challenge," in *Sixteenth Annual Conference of the International Speech Communication Association*, 2015.
- [29] Yougen Yuan, Cheung-Chi Leung, Lei Xie, Hongjie Chen, Bin Ma, and Haizhou Li, "Extracting bottleneck features and word-like pairs from untranscribed speech for feature representations," *Submitted to ASRU 2017*.
- [30] Aren Jansen and Benjamin Van Durme, "Efficient spoken term discovery using randomized algorithms," in *Automatic Speech Recognition and Understanding (ASRU), 2011 IEEE Workshop on*. IEEE, 2011, pp. 401–406.
- [31] Hayato Shibata, Taku Kato, Takahiro Shinozaki, and Shinji Watanabe, "Composite embedding systems for ZeroSpeech2017 track 1," *Submitted to ASRU 2017*.
- [32] Sharon Goldwater, Thomas L Griffiths, and Mark Johnson, "A bayesian framework for word segmentation: Exploring the effects of context," *Cognition*, vol. 112, no. 1, pp. 21–54, 2009.
- [33] Herman Kamper, Karen Livescu, and Sharon Goldwater, "An embedded segmental k-means model for unsupervised segmentation and clustering of speech," *Submitted to ASRU 2017*.
- [34] F García-Granada, E Sanchis, M.J. Castro-Bleda, J.A. González, and L.F. Hurtado, "ZeroSpeech2017 ELIRF-UPV system," *Submitted to ASRU 2017*.