



HAL
open science

Perceptimatic: A human speech perception benchmark for unsupervised subword modelling

Juliette Millet, Ewan Dunbar

► To cite this version:

Juliette Millet, Ewan Dunbar. Perceptimatic: A human speech perception benchmark for unsupervised subword modelling. Interspeech 2020 - 21st Annual Conference of the International Speech Communication Association, Oct 2020, Shanghai / Virtual, China. hal-03087252

HAL Id: hal-03087252

<https://hal.science/hal-03087252v1>

Submitted on 23 Dec 2020

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Perceptimatic: A human speech perception benchmark for unsupervised subword modelling

Juliette Millet^{1,2,3}, Ewan Dunbar^{1,2}

¹ CoML, ENS/CNRS/EHESS/INRIA/PSL Research University, Paris, France

² Université de Paris, LLF, CNRS, Paris, France

³CRI, Département Frontières du Vivant et de l’Apprendre, IIFR, Université de Paris

juliette.millet@cri-paris.org

Abstract

In this paper, we present a data set and methods to compare speech processing models and human behaviour on a phone discrimination task. We provide Perceptimatic, an open data set which consists of French and English speech stimuli, as well as the results of 91 English- and 93 French-speaking listeners. The stimuli test a wide range of French and English contrasts, and are extracted directly from corpora of natural running read speech, used for the 2017 Zero Resource Speech Challenge. We provide a method to compare humans’ perceptual space with models’ representational space, and we apply it to models previously submitted to the Challenge. We show that, unlike unsupervised models and supervised multilingual models, a standard supervised monolingual HMM–GMM phone recognition system, while good at discriminating phones, yields a representational space very different from that of human native listeners.

Index Terms: evaluation, unsupervised, speech recognition

1. Introduction

Not all errors are equal. If your task is to identify animals, mistaking a chimpanzee for a marmoset is less important than mistaking a chimpanzee for a hippo. This is also true when evaluating models that deal with human speech. Speech sounds differ from each other in different ways, and mistaking one sound for another in a given language is less important if the two sounds are also highly confusable for human native speakers.

Several works have made detailed comparisons between human speech perception experiments and automatic speech recognition (ASR) and related systems. In [1], neural ASR word transcription errors were compared qualitatively with errors made by human annotators, while [2] compared HMM phoneme confusions with the results of a human phoneme labelling experiment. The finer-grained study in [3] compared Japanese speakers’ perceptual boundaries for the (allophonic) [s]/[ç] contrast with the behaviour of GMM phoneme classifiers, while [4] compared phonetic adaptation by Dutch listeners to artificially “accented” productions of the [r]/[l] to acoustic models adapted to the experimental stimuli. Several studies also investigate whether ASR and unsupervised speech representation learning accords with established phenomena in human speech perception at a qualitative level [5, 6, 7, 8].

All of these comparisons focus either on one language, one model, or a handful of sounds. Most deal with models that predict phoneme labels, using methods that do not work for unsupervised models. In this paper, we present a large-scale phone discrimination task that can be performed by a wide range of models, and make a quantitative comparison with human listeners performing the same task.

Our contribution is twofold. First, we provide **Perceptimatic**, a data set containing stimuli and human results for a speech perception task in English and French (an extension of the English-only data described in [9]). Second, we use this data to evaluate the models submitted to the 2017 Zero Resource Speech Challenge (ZeroSpeech: see below).

We base our data set on test data used in the evaluation of the ZeroSpeech Challenge [10, 11, 12]. This machine learning challenge series aims to find methods for autonomously acquiring language, focusing on unsupervised learning for speech-related tasks. We focus on the French and English unsupervised subword modeling task in the 2017 edition. This task proposes to build systems that autonomously learn phoneme-like representations of speech from raw, continuous speech recordings. It is evaluated by a global phone discriminability measure applied to the learned representations. We provide a clean and manually verified subset of the 2017 test stimuli, along with experimental results from French- and English-speaking listeners on the phone discrimination task. To our knowledge, **Perceptimatic** is the first freely available data set of stimuli taken from natural, running French and English speech accompanied with human responses in a phone discrimination task, and as such, is the first human data comparable to the automatic phone discriminability evaluation used for unsupervised subword modelling.¹

We use two methods to apply these results to the models submitted to ZeroSpeech 2017: we compare models on how well they predict human behaviour in a speech discrimination task using the method in [13], and, in addition, we re-weight the results of the original ZeroSpeech 2017 phone discriminability scores to take account of native listeners’ relative discriminability of different sounds. We show that, with this reweighting, one of the submitted systems performs better than the challenge’s supervised ASR reference system. We also show that this supervised ASR system, while performing well on the phone discriminability evaluation, has a representational space that is very different from that of human native listeners, and very different from the systems submitted to the ZeroSpeech challenge.

2. Human and machine discrimination

2.1. Machine ABX discrimination

A machine ABX phone discrimination test is a binary decision task in which two speech stimuli, A and B, which only differ in their centre phone (for example, [serk]–[soøk] or [zfa]–[zpa]), are presented. A third stimulus, X, must be identified as being

¹All code used for the experiments, as well as stimuli and human results are available online at: <https://github.com/JAMJU/interspeech-2020-perceptimatic>, an extension of <https://github.com/JAMJU/Cogsci2020-Perceptimatic-English>

more similar to either A or B. This stimulus also shares the same first and last phones, and has the same centre phone as either A or B. R_A , R_B and R_X , the representations of respectively A, B, and X, are extracted from the models to be evaluated, and distances $d(R_A, R_X)$ and $d(R_B, R_X)$ are computed. Machine ABX phone discrimination tasks are often used for evaluating unsupervised speech models. They are generic, and are also appropriate for evaluating supervised models.

Since stimuli are not all of the same duration, distance functions based on dynamic time warping are typically used (notably in the ZeroSpeech Challenges). Dynamic time warping takes two sequences C and D as input (in our case either R_A and R_X or R_B and R_X), as well as a function γ for comparing pairs of sequence elements. It aligns C and D by matching the elements of one to the other so as to minimize the sum of $\gamma(c, d)$ for all matched elements (c, d) . Each element of C must be matched with at least one element of D , and alignments must respect temporal order. We calculate distances between stimuli $C = c_1, c_2, \dots, c_p$ and $D = d_1, d_2, \dots, d_q$ as:

$$d(C, D) = \frac{\sum_{c_i, d_j \text{ are matched}} \gamma(c_i, d_j)}{\max(p, q)} \quad (1)$$

As in ZeroSpeech 2017, we take γ as either the arc cosine of the normalized dot product or the symmetrised KL-divergence.

Once $d(R_A, R_X)$ and $d(R_B, R_X)$ are obtained using one of these methods, we compute $\delta = d(R_B, R_X) - d(R_A, R_X)$ if A and X belong to the same category, $\delta = d(R_A, R_X) - d(R_B, R_X)$ if it matches B. If $\delta > 0$, then the model is considered to be correct, otherwise, it is considered to be wrong. We perform this operation on many ABX triplets for a given pair of phones. The percent accuracy gives a measure of the model’s discriminability of the two phone categories.

The ZeroSpeech 2017 evaluation is an ABX phone discrimination task in which the A, B, and X stimuli are extracts taken from running speech. Models generate representations of audio files containing stimuli plus a surrounding context, from which their representations of the stimuli are extracted. In the current paper, we focus on the condition in which models are given one-second audio files, and in which X is uttered by a different speaker than A and B (across speaker). The ZeroSpeech 2017 evaluation measure was the average accuracy over all phone pairs.² Instead of relying only on this accuracy, we propose to evaluate models with respect to a human reference data set.

2.2. Human ABX discrimination

An ABX phone discrimination task for human participants is as follows: participants hear three stimuli in sequence (a triplet) and are asked to identify which one of the first two stimuli (A or B) is more similar to the third (X). As in the machine ABX task, the stimuli are such that one of the two responses is always considered correct (here, the centre phone of X matches either that of A or of B). We use the same stimuli as for the machine ABX task, and obtain multiple human responses for each triplet, which we code as correct or incorrect. We take two approaches to dealing with this data. First, we combine data across participants to compute an item-level accuracy for

²In [11], this average was computed in several steps by averaging accuracies the across contexts (flanking phones), then across speakers, and then across all centre phones. In our case, the test set on which we evaluate the models is almost exactly balanced for context, speakers, and centre phones. Thus, when we calculate accuracies, we simply take the global accuracy over all triplet items.

each triplet item. These item-level accuracies can be averaged into accuracies for individual phone contrasts, comparable with those calculated in the machine ABX task. Alternatively, we can drop the notion of category discriminability, and compare models’ gradient δ values with the probability that listeners give the correct answer, at the level of individual triplet items.

2.3. Comparing model and human performance

We relate human and model results in two ways. First, in order to have a phone discrimination score that takes into account differences in how “hard” each triplet item is, we weight models’ decisions by human accuracies as follows:

$$RewAcc_{ABX} = \frac{\sum_{t \in test} \mathbb{1}_{\delta_M(t) > 0} \times Hum(t)}{\sum_{t \in test} Hum(t)} \quad (2)$$

where $\delta_M(t)$ is the δ value for a triplet item t given by a model M , and with $Hum(t)$ is the percent accuracy for human listeners for the triplet item t (for French items, the accuracy is computed over French listeners, and for English items, the accuracy is computed over English listeners). This score gives more importance to triplets that are “easy” for human native listeners, and thus more important to discriminate correctly. It reduces the impact of triplets that are “hard” for humans, either because A, B and X are perceived as very similar, or because X is perceived to be more similar to the “wrong” answer.

Second, we compare models’ representational spaces with humans’ perceptual space in a detailed way. To do this, we evaluate models’ ability to predict individual human responses. We evaluate use the δ values as predictors in a binary regression, predicting whether a human listener will have the correct response or not on a given trial. Each model we evaluate is trained separately on French and on English, yielding two δ values.³ For each model, we fit an overparameterized probit regression with two zero-one language indicators as bias predictors, one which is 1 for French observations, and another which is 1 for English observations. We then construct French- and English-only δ predictors by multiplying the two δ values by the two indicator variables.⁴ We calculate the log-likelihood. Models with better (higher) log-likelihoods have representational spaces more similar to humans’, in the sense that relative distances in the model predict human discriminability.

3. Perceptimatic data set construction

The stimuli are taken from the French and English ZeroSpeech 2017 test stimuli. These stimuli were originally generated using forced alignment on the LibriVox audio book collection, and identifying all sequences of three phones. There are several problems with these materials. First, the number of triplets is too large to feasibly test on a large number of subjects. Second, there are labelling errors, as well as phones which may not be part of the language variety of certain listeners. Third, the phonetic similarity of the centre phone between X and the correct answer is often doubtful because of contextual variability. Fourth, the flanking phones are sometimes very different for A, B and X, due to contextual variability. Finally, the phone boundaries are not precise enough. We eliminated incorrectly labelled

³For models trained on other languages, we use the same δ values for French and English.

⁴We also add as a predictor a binary variable indicating whether the correct answer was presented first or second, and the trial’s position in the experiment, as well as predictors for participant.

phones, and then, as native English and French listeners, selected and realigned a subset of triplets by hand to minimize these issues. In total, the cleaned subset consists of 5202 triplets (2214 from English), making 461 distinct centre phone contrasts (212 English, 249 French), in a total of 201 distinct contexts (118 English, 83 French), with most phone comparisons appearing in three contexts each (a total of 47 English contrasts appear in either one, two, or four contexts). The speakers used (15 English, 18 French) have, in our assessment, pronunciations close to standard American English/Metropolitan French.⁵

The data set includes 91 participants located in the United States reporting English as the sole language to which they were primarily exposed up to the age of eight, and 93 located in France, attesting the same for French. They were recruited online on Amazon Mechanical Turk (all English speakers, and 57 French speakers) and in person, and all performed the test online with LMEDS software [14].⁶ They were paid for participation. All participants were tested both on English and French stimuli. Here we only report results on English contrasts for English listeners and on French contrasts for French listeners.

For testing, triplets were counterbalanced into lists of 190 per participant.⁷ Each triplet was tested three times, so that most contrasts are tested at least 36 times. Participants respond as to which of the two reference stimuli the probe corresponded to on a six-point scale, ranging from *first for sure* to *second for sure*, with two intermediate degrees of certainty for each. The data set includes both these responses and a binarized version. Except in Figure 2, we report only the binarized responses.

4. Experiments

We evaluate the models submitted to the 2017 ZeroSpeech challenge. We compare them with the topline representation used in the challenge, posteriorgrams from a supervised HMM-GMM phone recognition system with a bigram phone-level language model, trained with a Kaldi recipe [16]. We also compare them with the baseline representation, mel filterbank cepstral coefficients (MFCCs) (thirteen first coefficients with Δ and $\Delta\Delta$, with mean-variance normalization over a moving 300 millisecond window). We add two models that have been shown to have representational spaces similar to human perceptual space [9]: multilingual bottleneck features (**Bot**: [17]) and a Dirichlet process Gaussian mixture model (**DP**: [18]). **Bot** representations are from a bottleneck layer of a model trained to label phoneme states in seventeen phonetically diverse languages (107 hours training data in total; French and English are not included).⁸ We use the Shennong package⁹ to extract these representations.

⁵The full set of English centre phones included in at least one item is [æ a b d ð eɪ e f g h i ɪ k l m n ŋ oʊ p ɪ s ʃ t tʃ u ʊ v ʌ w z]. The full set of French phones is [a ɑ̃ b d e ɛ ɛ̃ f g j k l m n ɲ o ɔ̃ ɔ̃ p ʁ s ʃ t u v w y z ʒ]. For the full list of pairs and contexts, see the online repository.

⁶[15] compared data from an in-lab speech perception experiment with a Mechanical Turk replication and found a close correspondence. 48 additional participants (15 US, 33 France) were tested but did not meet the language background requirements, and 115 participants (65 US, 50 France) were rejected for failing at least three out of twelve catch trials or not finishing the task. The catch trials consisted of additional, highly distinct ABX stimuli, including several which required participants to distinguish *cat* from *dog* for English speakers or *caillou* from *hibou* for French speakers.

⁷No participant is tested twice on the same phone pair, and the combination of speakers is not predictive of the right answer.

⁸If the same sound belongs to different inventories, it is treated as distinct, for a total of 1032 possible phonemes.

⁹<https://github.com/bootphon/shennong>

DP is a Gaussian mixture model over individual frames, from which we draw posteriorgrams. We use pretrained models from [13], trained on 34 hours of English speech and on 33 hours and 42 minutes of French speech, both from LibriVox.

For each ZeroSpeech 2017 model, we use the English and French features as they were submitted to the challenge, in the one second condition. All the models are presented in [11], and we evaluate the systems **[H]** [19], **[P1]**, **[P2]**, [20] **[A3]** [21], **[Y1]**, **[Y2]** [22], **[S1]**, **[S2]** [23], **[C1]** and **[C2]** [24]. Broadly the systems follow four types of strategies. The first (**[H]**, **[P1]** and **[P2]**) consists in performing a bottom-up frame-level clustering. The second (**[C1]**, **[C2]** and **[A3]**) is to construct language-independent embeddings by training neural networks in an unsupervised way. The third (**[Y1]** and **[Y2]**) is to use spoken term discovery to improve the acoustic features. **[Y3]** is a supervised version of these systems, using transcribed pairs from the Switchboard data set. Finally, **[S1]** and **[S2]** use supervised pre-training on out-of-domain languages (only Japanese for **[S1]** and ten different languages, including French and English, for **[S2]**). Except for **[S1]**, **[S2]** and **[Y3]**, all the models are trained only on the 2017 challenge data. **[H]** and **[P1]** use only monolingual data, whereas the rest of the models, while trained separately for each language, make use of the complete set of languages provided with the ZeroSpeech challenge during training. We followed the original choices of the authors concerning γ (cosine distance for all, except for **[H]**, which uses symmetrised KL-divergence). We use the cosine distance for all reference models, except for **DP** and the topline, for which we use the symmetrised KL-divergence.

5. Results

Original ZeroSpeech 2017 ABX accuracies are given on the first row of Table 1. The original 2017 ZeroSpeech test set contains errors, partially corrected by our modifications (see Section 3). ABX accuracies computed only on the Perceptimatic triplets are on the row **Perceptimatic**. We observe that the English and French scores become more similar (French scores increase and English scores decrease). Future submissions to ZeroSpeech 2017 should use our data set for a better evaluation of their model. We also re-weight the results of each model by humans’ results (see row **Perceptually Weighted**) using (2). We observe that re-weighting ABX accuracies by giving more importance to contrasts that are “easy” improves all the scores (see Table 1). This implies that globally, models already have more trouble with “hard” contrasts than with “easy” ones.

We then evaluate models based on how well they predict human behaviour in the experiment. Each “model” we evaluate is in fact two trained models (one French, one English), which each make a gradient prediction of discriminability, δ for each triplet item. We use these two δ values as predictors for the response accuracy of the respective human native listeners (the French δ predicts French stimuli/listeners, and the English δ English stimuli/listeners) on each item. We then calculate the overall log-likelihood of a binary-response regression with respect to the experimental data. (See 2.3 above.) This yields a measure of how well each model predicts human behaviour. To generate confidence intervals, we resample (N=13655: for each triplet item, we draw exactly three human responses without replacement) and re-fit the regressions.¹⁰

Figure 1 plots the phone discrimination accuracy

¹⁰Confidence intervals for the differences in log-likelihood between each pair of models can be found on the online repository.

Table 1: 2017 ZeroSpeech challenge ABX % accuracies (one-second across-speaker conditions), ABX % accuracies on Perceptimatic, and ABX % accuracy reweighted by human discriminability. Higher scores are better.

		topline	S2	Bot	H	S1	DP	C2	A3	Y3	C1	Y2	Y1	P2	P1	MFCC
Zerospeech 2017	French	89.4	88.8	86.9	86.4	86.3	83.7	83	82.8	82.3	82.4	81.3	81.1	79.9	79.7	74.9
	English	91.4	92.1	90.7	89.9	89.9	88.8	87.3	86.8	86.4	86.2	86	85.8	82.4	82.4	73.4
Perceptimatic	French	92.8	92.3	88.5	88.6	89.6	86.5	85.8	84.4	84.5	84.7	83.3	83.4	81.7	81.9	78.3
	English	91.8	92	89.1	90.6	88.5	88.7	86.9	82.8	84.9	84.5	83.2	82.7	82.5	81.8	78.2
Perceptually Weighted	French	93.4	94	90.1	90.3	91.3	88.5	86.5	86.6	86.3	87.4	85.1	85.2	83.2	83.2	79.9
	English	92.8	94.1	91.2	93.0	91.1	91	87.6	85.7	87.1	89	85.7	85.1	84.8	84.3	80.6

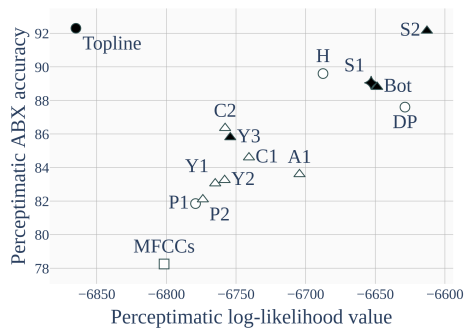


Figure 1: Perceptimatic ABX accuracy (mean of French and English: higher is better) versus log-likelihood. Higher log-likelihood values indicate representational spaces more similar to humans’ perceptual space. White: unsupervised; black: supervised. Triangle: multilingual; circles: monolingual on evaluated language; diamond: monolingual on another language, square: MFCCs.

(Perceptimatic) against the mean resampled log-likelihoods. We observe that having a representational space similar to humans’ perceptual space is correlated with having a good ABX accuracy. However, we notice an obvious outlier: the (topline) supervised ASR model provided as a reference system for ZeroSpeech 2017. This model is unexpectedly bad at predicting human results. This result is consistent with results obtained for neural ASR models trained exclusively on the native language of human listeners [9]. Unsupervised models (such as **DP**) and multilingual models (such as **[S1]**, **[S2]**, and **Bot**) seem to be more predictive of human perception. It is therefore possible to discriminate phones well while making phoneme confusions very different from those of human beings. We discuss some possible explanations for this below.

Conversely, the fact that **DP** obtains the second-highest log-likelihood (not significantly different from **[S2]**) is surprising, given its lower discrimination accuracy. **[H]**, for example, built on **DP** to generate a more speaker-invariant representation (it uses DPGMM clusters as the basis for supervised talker normalization techniques). While these improvements improve phone discriminability, they yield a less human-like representation.

The best model, both in terms of perceptually-reweighted accuracy, and in terms of predicting human behaviour, is **[S2]**. Similarly, **Bot** also performs well. These two models train on multilingual supervised phone recognition objectives. **[S2]** is then further adapted to either English or French, while **Bot** is fixed (and thus has the same δ for the two stimuli languages). That these systems perform so well suggests that human behaviour in this task is driven mainly by universal, general speech




Figure 2: Average of French listeners’ non-binarized results (higher: correct, and more certain of it) against average δ from (left) supervised ASR (right) multilingual system **[S2]**. Each point is a phone pair. Measures are normalized by dividing by standard deviation over the entire data set.

perception mechanisms. With respect to this behaviour, monolingual supervised ASR may be over-training to the phone identification objective. Figure 2 plots French listeners’ continuous discriminability scores against the French δ for supervised ASR and for **[S2]**, normalized to be comparable. The ASR δ values are all shifted to the right, indicating better discriminability, but at the expense of a poorer correlation with human accuracy.

6. Overview

We have presented Perceptimatic, a data set based on the 2017 ZeroSpeech challenge evaluation data. We used it to evaluate the models submitted to the 2017 ZeroSpeech challenge, and showed that, surprisingly, the supervised ASR reference model has a representational space very different from the perceptual space of human listeners. On the other hand, unsupervised models, and supervised models trained on other languages than the one being tested, seem to have more human-like representations. We leave open at least two possible explanations for this result. One is that human listening is simply not tuned to optimize phone classification in the native language, unlike supervised ASR. Another possibility rests on the observation that extracts of phones drawn from running speech, while standard for evaluating unsupervised speech representation learning, are very different from the types of clean speech stimuli used in typical human speech perception experiments. For humans, discriminating naturalistic examples of phones may tap into a more acoustic mode of listening; some other type of listening test might give results more comparable to a supervised ASR.

7. Acknowledgements

This research was supported by the École Doctorale Frontières du Vivant (FdV) – Programme Bettencourt, and by grants ANR-17-CE28-0009 (GEOMPHON), ANR-11-IDFI-023 (IIFR), ANR-11-IDEX-0005 (USPC), ANR-10-LABX-0083 (EFL), and ANR-17-EURE-0017.

8. References

- [1] A. Stolcke and J. Droppo, “Comparing human and machine errors in conversational speech transcription,” in *INTERSPEECH*, 2017.
- [2] B. T. Meyer, M. Wächter, T. Brand, and B. Kollmeier, “Phoneme confusions in human and automatic speech recognition,” in *Eighth Annual Conference of the International Speech Communication Association*, 2007.
- [3] G. Pintér and H. Watanabe, “Do GMM phoneme classifiers perceive synthetic sibilants as humans do?” in *INTERSPEECH*, 2016, pp. 1363–1367.
- [4] O. Scharenborg, S. Tiesmeyer, M. Hasegawa-Johnson, and N. Dehak, “Visualizing phoneme category adaptation in deep neural networks,” in *INTERSPEECH*, 2018, pp. 1482–1486.
- [5] N. H. Feldman and T. L. Griffiths, “A rational account of the perceptual magnet effect,” in *Proceedings of the Annual Meeting of the Cognitive Science Society*, vol. 29, no. 29, 2007.
- [6] T. Schatz, N. Feldman, S. Goldwater, X. N. Cao, and E. Dupoux, “Early phonetic learning without phonetic categories: Insights from machine learning,” *Proceedings of the National Academy of Sciences*, To appear.
- [7] T. Schatz, F. Bach, and E. Dupoux, “ASR systems as models of phonetic category perception in adults,” in *Proceedings of the 39th Annual CogSci Meeting*, 2017.
- [8] T. Schatz and N. Feldman, “Neural network vs. HMM speech recognition systems as models of human cross-linguistic phonetic perception,” in *Proceedings of the Conference on Cognitive Computational Neuroscience*, 2018, pp. 1–4.
- [9] J. Millet and E. Dunbar, “The Perceptimatic English Benchmark for Speech Perception Models,” in *CogSci Conference 2020*, 2020.
- [10] M. Versteegh, R. Thiollière, T. Schatz, X. N. Cao, X. Anguera, A. Jansen, and E. Dupoux, “The Zero Resource Speech Challenge 2015,” in *INTERSPEECH-16*, 2015.
- [11] E. Dunbar, X. N. Cao, J. Benjumea, J. Karadayi, M. Bernard, L. Besacier, X. Anguera, and E. Dupoux, “The Zero Resource Speech Challenge 2017,” in *2017 IEEE Workshop on Automatic Speech Recognition and Understanding (ASRU)*. IEEE, 2017, pp. 323–330.
- [12] E. Dunbar, R. Algayres, J. Karadayi, M. Bernard, J. Benjumea, X.-N. Cao, L. Miskic, C. Dugrain, L. Ondel, A. W. Black *et al.*, “The zero resource speech challenge 2019: TTS without T,” in *INTERSPEECH*, 2019.
- [13] J. Millet, N. Jurov, and E. Dunbar, “Comparing unsupervised speech learning directly to human performance in speech perception,” in *CogSci Conference 2019*, 2019.
- [14] T. Mahr, “LMEDS: Language markup and experimental design software,” 2016.
- [15] D. Kleinschmidt and T. F. Jaeger, “Robust speech perception: Recognize the familiar, generalize to the similar, and adapt to the novel,” *Psychological Review*, vol. 122, no. 2, pp. 148–203, 2015.
- [16] D. Povey, A. Ghoshal, G. Boulianne, L. Burget, O. Glembek, N. Goel, M. Hannemann, P. Motlicek, Y. Qian, P. Schwarz *et al.*, “The Kaldi speech recognition toolkit,” in *IEEE 2011 Workshop on Automatic Speech Recognition and Understanding (ASRU)*, no. EPFL-CONF-192584. IEEE Signal Processing Society, 2011.
- [17] R. Fer, P. Matejka, F. Grezl, O. Plchot, K. Vesely, and J. H. Cernocky, “Multilingually trained bottleneck features in spoken language recognition,” *Computer Speech and Language*, vol. 46, no. Supplement C, pp. 252 – 267, 2017.
- [18] H. Chen, C.-C. Leung, L. Xie, B. Ma, and H. Li, “Parallel inference of Dirichlet process Gaussian mixture models for unsupervised acoustic modeling: A feasibility study,” in *INTERSPEECH-16*, 2015.
- [19] M. Heck, S. Sakti, and S. Nakamura, “Feature optimized dpgmm clustering for unsupervised subword modeling: A contribution to zerospeech 2017,” in *2017 IEEE Automatic Speech Recognition and Understanding Workshop (ASRU)*. IEEE, 2017, pp. 740–746.
- [20] C. Manenti, T. Pellegrini, and J. Pinquier, “Unsupervised speech unit discovery using k-means and neural networks,” in *International Conference on Statistical Language and Speech Processing*. Springer, 2017, pp. 169–180.
- [21] T. Ansari, R. Kumar, S. Singh, and S. Ganapathy, “Deep learning methods for unsupervised acoustic modeling—leap submission to zerospeech challenge 2017,” in *2017 IEEE Automatic Speech Recognition and Understanding Workshop (ASRU)*. IEEE, 2017, pp. 754–761.
- [22] Y. Yuan, C.-C. Leung, L. Xie, H. Chen, B. Ma, and H. Li, “Extracting bottleneck features and word-like pairs from untranscribed speech for feature representation,” in *2017 IEEE Automatic Speech Recognition and Understanding Workshop (ASRU)*. IEEE, 2017, pp. 734–739.
- [23] H. Shibata, T. Kato, T. Shinozaki, and S. Watanabet, “Composite embedding systems for zerospeech2017 track1,” in *2017 IEEE Automatic Speech Recognition and Understanding Workshop (ASRU)*. IEEE, 2017, pp. 747–753.
- [24] H. Chen, C.-C. Leung, L. Xie, B. Ma, and H. Li, “Multilingual bottle-neck feature learning from untranscribed speech,” in *2017 IEEE Automatic Speech Recognition and Understanding Workshop (ASRU)*. IEEE, 2017, pp. 727–733.