



HAL
open science

Self-supervised language learning from raw audio: Lessons from the Zero Resource Speech Challenge

Ewan Dunbar, Nicolas Hamilakis, Emmanuel Dupoux

► **To cite this version:**

Ewan Dunbar, Nicolas Hamilakis, Emmanuel Dupoux. Self-supervised language learning from raw audio: Lessons from the Zero Resource Speech Challenge. *IEEE Journal of Selected Topics in Signal Processing*, In press, 10.1109/jstsp.2022.3206084 . hal-03789716

HAL Id: hal-03789716

<https://hal.science/hal-03789716>

Submitted on 27 Sep 2022

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.

Self-supervised language learning from raw audio: Lessons from the Zero Resource Speech Challenge

Ewan Dunbar, Nicolas Hamilakis, and Emmanuel Dupoux

Abstract—Recent progress in self-supervised or unsupervised machine learning has opened the possibility of building a full speech processing system from raw audio without using any textual representations or expert labels such as phonemes, dictionaries or parse trees. The contribution of the Zero Resource Speech Challenge series since 2015 has been to break down this long-term objective into four well-defined tasks—Acoustic Unit Discovery, Spoken Term Discovery, Discrete Resynthesis, and Spoken Language Modeling—and introduce associated metrics and benchmarks enabling model comparison and cumulative progress. We present an overview of the six editions of this challenge series since 2015, discuss the lessons learned, and outline the areas which need more work or give puzzling results.

Index Terms—textless speech processing, unsupervised and self-supervised learning, representation learning.

I. INTRODUCTION

FOR a long time, language technology has been developed principally using large quantities of textual resources. This makes sense, since, as far as technological applications are concerned, language has primarily been used in written form. When it comes to dealing with spoken language, however, this has given rise to a division of labor between, on the

Ewan Dunbar is with the Department of French, University of Toronto, Toronto, ON, M5S 1A1, Canada (e-mail: ewan.dunbar@utoronto.ca). Nicolas Hamilakis is with the École Normale Supérieure, 75005 Paris, France, and with Inria, 75013 Paris, France (e-mail: nick.hamilakis562@gmail.com).

Emmanuel Dupoux is with the École des Hautes Études en Sciences Sociales, 75006 Paris, France, and with Meta AI, 75009 Paris, France (e-mail: emmanuel.dupoux@gmail.com).

one hand, speech components which aim at converting speech to text or text to speech (ASR, automatic speech recognition, and TTS, text-to-speech synthesis), and, on the other hand, components that perform a variety of language tasks based on text (language understanding, dialogue, language generation). As a result, even speech-first applications like speech-to-speech translation or speech assistants like Alexa or Siri are cobbled together in a Frankensteinian fashion, with some components trained on text and others trained on speech (see Figure 1a)—and with all the speech components trained using large amounts of supervision (textual transcription) so that they can communicate with the text-based components. But is this a necessity? Could we build spoken-language based applications directly from the audio stream without using any text?

Preschoolers around the world demonstrate clearly that it is possible to learn to naturally interact in language without knowing how to read or write [1], [2]. Written language is, in a way, a tool for archiving spoken or signed language. Many languages have no writing system at all, and many other language communities do not use the written form of their language much (reportedly, more than half the world's languages do not have a stable or widely used writing system).

Reverse engineering the feat of learning a spoken language from speech input only is a fascinating scientific question. The Zero Resource Speech Challenge (ZRC) focuses on spoken languages. (It is an important and pressing question how zero-resource technology could be applied to working with signed languages.) For spoken language technology, advancing this question would unlock a number of novel applications.

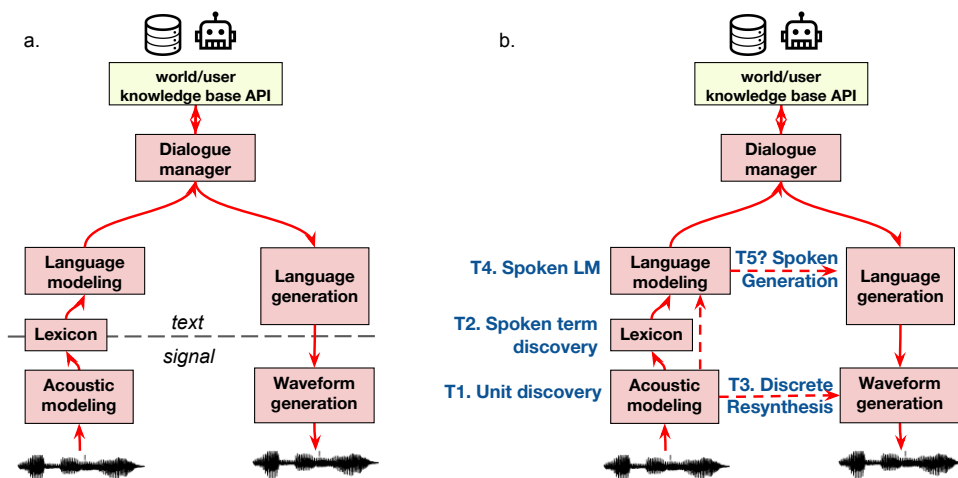


Fig. 1. a. Traditional pipeline for a spoken assistant based on textual resources b. Pipeline and tasks for the zero resource challenge series.

TABLE I
SUMMARY OF THE METRICS USED IN THE ZERO RESOURCE CHALLENGE SERIES FOR EACH TASK.

Task	Ling. Level	Model outputs examined	Metric	Example
T1. Unit Discovery	Phonetic	sequences of frame embeddings	ABX: $d(a, x) < d(b, x)?$ $a \in A, b \in B, x \neq a \in A$	$a = \text{'fly'}_1, b = \text{'fry'}_1, x = \text{'fly'}_2$
	Lexical (matching)	pairs of segments	NED: $ED(a, b)/\max(a , b)$ COV: fraction of corpus covered	$a = \text{'rose'}, b = \text{'prose'}$
T2. Spoken Term Discovery	Lexical (clustering)	cluster memberships of segments	Grouping F-score Type F-score	$c_1 = \{\text{'rose'}_1, \text{'prose'}_1, \text{'ose'}_2\}$ $c_2 = \{\text{'rose'}_2, \text{'rose'}_3\}$
	Lexical (segmentation)	list of time-stamps	Token F-score: as defined in text Boundary F-score: segmentation	'Ho ware youto day'
T3. Unsup. Discrete Resynthesis ("TTS without T")	Phonetic	series of discrete units waveforms	Bitrate: $\frac{n}{D(U)} \sum p(u_i) \log(p(u_i))$ MOS: human evaluation	$U = u_1, \dots, u_n$ $U = \text{'How are you today'}$
	Lexical (frequency)	(pseudo) probabilities	spot-the-word: $\hat{p}(a) > \hat{p}(b)?$	$a = \text{'brick'}, b = \text{'*blick'}$
T4. Spoken LM	Lexical (semantics)	word embeddings	similarity: $d(a, b) \propto d_h(a, b)?$	$a = \text{'abduct'}, b = \text{'kidnap'}$
	Syntax	(pseudo) probabilities	accept. judgment: $\hat{p}(a) > \hat{p}(b)?$	$a = \text{'dogs sleep'}, b = \text{'*dogs sleeps'}$

Note. d is a dissimilarity measure between word or frame embeddings, d_h is a human dissimilarity judgment, ED the edit distance over the phonetic transcriptions of segments (stretches of speech between two time-stamps), $D(U)$ the total duration of dataset U (in sec), p the probability of a given discrete unit u_i in U , \hat{p} a pseudo-probability computed by the model over an input sequence (word or sentence). * indicates ungrammaticality

For one thing, it would allow for applications that address the needs of languages that are entirely or mostly unwritten. Even in languages with large amounts of textual resources, learning language representations from raw audio would help capture the dimensions of language that are typically poorly represented in text (prosody, emotional and non-verbal vocalizations, and so on). Beyond helping to model these aspects of language, capturing unwritten oral expression could improve the ability of machine learning systems to deal with spontaneous speech, thereby unlocking the rich syntax and vocabulary of oral registers, which very often differ greatly from the written form. This would foster more naturalistic human/machine interactions.

While some research has focused on how self-supervised modelling can improve existing supervised speech tasks (for example, the SUPERB benchmark: [3]), the Zero Resource Challenge series assesses progress toward spoken language systems that function without any textual supervision at any point. Building a text-free dialogue system using only raw speech is a difficult machine learning challenge. It requires us to rethink the spoken language processing stack from the ground up. The ZRC series has been designed to address two interlocking research problems: the task problem and the evaluation problem.

The *task problem* is to break down the ill-defined objective of “learning to process spoken language without text” into a series of well-defined sub-problems. The ZRC series follows the general architecture of a spoken digital assistant to define the learning problem implied by the training of each component—the acoustic model, a lexicon, language models, waveform generation, and so on. But instead of using phonemes, characters or words as an intermediate representation, the components are allowed to develop their own latent representations. For instance, instead of outputting phonemes or characters, the acoustic model is assumed to output *acoustic units* which may or may not be discrete. Such an architecture (see Figure 1b) naturally gives rise to the following four

tasks: (Task 1) Acoustic Unit Discovery; (Task 2) Spoken Term Discovery; (Task 3) Unsupervised Discrete Resynthesis; (Task 4) Spoken Language Modeling. These are the textless counterparts of well-known tasks: (Task 1) phone recognition; (Task 2) word recognition (i.e., ASR); (Task 3) TTS; (Task 4) Language Modeling. We will review these tasks in turn in the following sections.

The *evaluation problem* is to define metrics that enable model comparison and cumulative progress for tasks that are defined only through a self-supervised objective. For instance, ASR systems can readily be measured through phone or word error rates. But their self-supervised counterparts, Acoustic Unit Discovery systems, do not aim at recovering phonemes, but a latent representation. How can we evaluate these systems? Interest in some of the above mentioned tasks predates the ZRC series (see for instance [4]–[9], for Task 1), but since each of the published papers used its own metrics, it was difficult (and still is) to compare systems and measure progress. The general strategy of the ZRC series is to develop zero-shot probe tasks that are inspired by human psycholinguistics, and which require no model retraining. The reasoning is that, since the aim is to probe for latent representations at various levels of a self-supervised system, it is best to not train any classifier on top of it. Otherwise, it would be unclear whether the performance obtained reflects the system under observation or simply the classifier. For each task, zero-shot metrics were developed that probe for the different levels of linguistic knowledge that the self-supervised system is supposed to have learned. They only require the extraction of information readily available in the system (for example, embeddings, pseudo-probabilities), and are computed by a separate fixed module which is identical across systems. The evaluation metrics that go with the tasks are listed in Table I and will be presented in more detail in the following sections.

In this paper, we provide a comprehensive overview of the results obtained across the different tasks and metrics of the

TABLE II
SUMMARY OF TASKS AND BENCHMARKS IN THE ZERO RESOURCE CHALLENGE SERIES.

Chall.	Tasks	Benchmark
2015 [10]	T1, T2	ABX-15, TDE-15
2017 [11]	T1, T2	ABX-17, TDE-17
2019 [12]	T3	TTS0-19
2020 [13]	T1,T2,T3	<i>reboot of 2017 & 2019</i>
2021a [14]	T1,T4	ABX-LS, sLM-21 (sWUGGY,
2021b	T1,T4	sBLIMP, sSIMI)

ZRC series since 2015, and we discuss the lessons learned and outline the areas that need more work or give puzzling results.

II. PAST AND PRESENT

Six editions of the Zero Resource Challenge have been proposed over the years as events in different venues (Interspeech, ASRU, NeurIPS) and are summarized in Table II. Each edition has explored a different combination of tasks and introduced different datasets. Overall, the six editions have received a total of 115 submissions from 29 teams. In addition, several papers have been published using some of the Zero Resource benchmark metrics, which we also include in our review.

Table III gives the complete list of submitted systems to all four tasks, and the abbreviations we use for them, including citations for published systems and model types as explained in the upcoming sections.

A. Task 1: Acoustic Unit Discovery

The goal of acoustic unit discovery is to learn representations (embeddings) of speech sounds that retain linguistically relevant information and discard acoustic information which is irrelevant or secondary to recovering the linguistic content, like speaker voice type or recording conditions (additive noise, reverberation, etc). In text-based systems, such representations are typically phonemes (as defined by a pronunciation dictionary) or characters. Here, we let the representations be latent, which means that they may take any form (dense vectors for each frame, probabilistic codes, discrete codes, etc). This poses an evaluation problem. The ZRC series takes the view that, while discovered units may not necessarily take the shape of straightforwardly linguistically interpretable entities like phonemes or phonetic features, they should at least maintain the same key linguistic function: *linguistic contrast*.

1) *Evaluation metrics*: Phoneme categories are typically defined as the smallest element of speech that can induce a difference in meaning between words. In English, for instance, the phonemic contrast between /r/ and /l/ is demonstrated by the fact that “fly” and “fry” are distinct words. Two instances of “fly” would remain instances of the word “fly” even in spite of variations in speaker or recording condition, and an instance of “fry” is (for speakers with a standard pronunciation) never an instance of “fly.” The same goes for possible, as opposed to actual, words: “pla” and “pra” would not be the same word, if they were words.

The minimal pair **ABX** task [72], [73] is inspired by match-to-sample tasks used in human psychophysics and measures discriminability between two sound categories. We define Δ ,

TABLE III
LIST OF CONTRIBUTED SYSTEMS FOR EACH BENCHMARKS

System	Date	Paper	Author affiliation	Model type
<i>ABX-15 (Task 1)</i>				
Bad15a-c	2015	[15]	LCSL IIT-MIT	AutoEnc
Bal15a-e	2015	[16]	CMU	Other
Che15a-d	2015	[17]	ASTAR	GMM
Fah15	2015	[18]	KTH	Siamese
Hec16a-d	2016	[19]	NAIST	GMM
Ren15a-d	2015	[20]	Edinburgh	AutoEnc
Sri15a,b	2015	[21]	IIT Hyderabad	AutoEnc
Thi15	2015	[22]	CoML	Siamese
Zeg15	2015	[23]	CoML	Siamese
<i>ABX-17 (Task 1)</i>				
Ans17a-d	2017	[24]	LEAP-IIS	GMM/AutoEnc
Che17a,b	2017	[25]	NU Singapore	GMM
Hec17a,b	2017	[26]	NAIST	GMM
Pel17a,b	2017	[27]	Toulouse	Other
Ras17	2017	Unpub	Aalto	Other
Shi17a,b	2017	[28]	Tokyo IT	Other
Yua17a-c	2017	[29]	NU Singapore	AutoEnc
Cho19	2019	[30]	Wroclaw/DMind	AutoEnc
Kha20	2020	[31]	Meta	Predict.
<i>TDE-15 (Task 2)</i>				
BAS2	2015	[32]	JHU	Match-first
Kam15a	2015	[33]	Stellenbosch	Match-first
Lyz15a-c	2015	[34]	JHU	Match-first
Ras15a-c	2015	[35]	Aalto	Match-first
<i>TDE-17 (Task 2)</i>				
BAS2	2017	[32]	JHU	Match-first
Kam22	2022	[36]	Stellenbosch	Segment-first
Alg22	2022	[37]	CoML	Segment-first
Kam17	2017	[38]	Stellenbosch	Segment-first
Bha20a,b	2020	[39]	JHU	Segment-first
Gar17a,b	2017	Unpub	UP València	Match-first
Ras20a,b	2020	[40]	Aalto	Match-first
Ras17	2017	[41]	Aalto	Match-first
Iwa21	2021	[42]	Tokyo IT	Segment-first
<i>TTS0-19 (Task 3)</i>				
BAS3	2019	[43]	Brno	AutoEnc/Low br
Tja19a,b	2019	[44]	NAIST	AutoEnc/High br
Hor19a,b	2019	Unpub	Horizon	AutoEnc/High br
Pan19a,b	2019	[45]	IIT Madras	Other/Low br
Kam19a,b	2019	[46]	Stellenbosch	AutoEnc/High br
Fen19a,b	2019	[47]	CUHK	AutoEnc/High br
Ral19	2019	Unpub	CMU	GMM/Low br
Cho19a,b	2019	Unpub	Korea U	AutoEnv/High br
Liu19a,b	2019	[48]	Nat Taiwan U	AutoEnc/Low br
Kum19a,b	2019	[49]	IIT Hyderabad	AutoEnc/Low br
Yus19	2019	[50]	Boğaziçi	Other/Low br
Gök19	2019	[50]	Boğaziçi	AutoEnc/Low br
Gün20	2020	[51]	Boğaziçi	AutoEnc/Low br
Yus20a-c	2020	[52]	Brno	Other/Low br
Hou20a,b	2020	Unpub	Tokyo IT	AutoEnc/High br
Mor20	2020	[53]	Kyoto	AutoEnc/High br
Kum20a,b	2020	[54]	IIT Madras	Other/Low br
Che20a,b	2020	[55]	Sheffield	AutoEnc/High br
Lum20a,b	2020	[56]	Nagoya	AutoEnc/High br
Nie20a,b	2020	[57]	Stellenbosch	Other/High br
Tja20a,b	2020	[58]	NAIST	AutoEnc/High br
<i>ABX-LS, sLM-21 (Task1, Task4)</i>				
BAS4	2021	[59]	CoML	Predict.+BERT
BAS4vg	2021	[60]	CoML	Predict.+BERT
Nie21	2021	[61]	Stellenbosch	Predict.+LSTM
Cho21a,b	2021	[62]	Wroclaw	Predict./Other
Liu21	2021	Unpub	PLA IEU	Predict.+LSTM
Mae21a-d	2021	[63]	Yahoo/CMU	Predict.+BERT
Pen21	2021	[64]	UT Austin	Predict.+BERT
Lee21a,b	2021	Unpub	Seoul	Other+BERT
Gan21	2021	Unpub	LEAP-IIS	Predict.+LSTM
Ngu21a-d	2021	[65]	CoML	Predict.+BERT
Łań21	2021	Unpub	Wroclaw	Predict.+LSTM
Bha21a	2021	[66]	JHU	Predict.
Gao21a-c	2021	Unpub	TU Delft	Predict.+BERT
Qia22	2022	[67]	MIT-IBM	Predict.+Transf
Cue22	2022	[68]	Wroclaw	Predict.

Note. br: bitrate. See model type details in Section II.

TABLE IV
CHARACTERISTICS OF THE ZRC BENCHMARK DATASETS. *FOR TDE-17, TEST SET IS SAME AS TRAIN SET.

Benchmark	Language (source dataset)	Type	(Default) Train Set	Test Set
ABX-15, TDE-15	English (Buckeye)	conversations	Same as test	5h, 12 spk
	Xitsonga (NCHLT)	TIMIT-like		2h30, 24 spk
	English (Librivox)	audiobook	45h	27h, 9 spk
ABX-17, TDE-17*	French (Librivox)	audiobook	24h	17h, 10 spk
	Mandarin (THCHS-30)	read speech	2h30	25h, 4 spk
	German[L1] (Librivox)	audiobook	25h	11h, 10 spk
	Wolof[L2] ([69])	TIMIT-like	10h	5h50, 4 spk
ABX-LS	English (LibriSpeech)	audiobook	960h or 100h	dev/test x clean/other: 5h each 40 spk for clean, 33 spk for other
TTS0-19	English (Librispeech)	audiobooks	Voice: 4h40, 2 spk; Unit: 15h, 100 spk	28min, 24spk
	Indonesian ([70], [71])	TIMIT-like	Voice: 1h30, 1 spk; Unit: 15h, 112 spk	29 min, 15 spk
sLM-21 (sWUGGY, sBLIMP, sSIMI)	English (LibriSpeech)	audiobook	960h	Synthetic words or short sentences or extracted from LibriSpeech

the ABX-discriminability of category A from category B , as the probability that tokens a and x from A are further apart than token b from B and x are, according to a dissimilarity function d , see Equation 1:

$$\Delta(A, B) = \frac{1}{|A|(|A|-1)|B|} \sum_{a \in A} \sum_{b \in B} \sum_{x \in A \setminus a} (\mathbb{1}_{d(a,x) < d(b,x)} + \frac{1}{2} \mathbb{1}_{d(a,x) = d(b,x)}) \quad (1)$$

where $\mathbb{1}$ is the indicator function and $|A|$ the number of tokens in category A . The discriminability score is symmetrized by averaging $\Delta(A, B)$ and $\Delta(B, A)$.

Tokens are representation of speech sequences as output by the model under evaluation. In general, they will be sequences of frame embeddings, and dynamic time warping is used to realign them. Frame-level dissimilarities are averaged along the realignment path to obtain d . Most submissions used angular dissimilarity (arccos of the normalized dot product of the frame embeddings), while others used the KL divergence when the frame representations were posterior probabilities.

For the categories A and B we use triphone sequences that only vary in the middle phoneme (like “fly” and “fry”), as extracted from longer utterances in our test set. Thus, participants apply their trained models to these audio files, and output sequences of representations for the entire file. These representations must be time-stamped, so that the ABX evaluation software can identify the beginning and end of each three-phone token in the output. This was done in all three **ABX**- benchmarks; only the ABX evaluation that was contained in the **TTS0-19** evaluation package did otherwise, using small audio files containing only the three-phone sequence (see Section II-C)¹.

For the *within-speaker* variant of this task, all of the phone triplets belong to the same speaker (e.g., $a = \text{fly}_{T_1}$, $b = \text{fry}_{T_1}$, $x = \text{fly}'_{T_1}$). The error rates 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 $/l-/r/$, average the scores for all attested contexts such as $/f_y/$, $/f_i/$, $/a_o/$, 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* variant, a and b belong to the same speaker, and x to a different one. $a = \text{fly}_{T_1}$, $b = \text{fry}_{T_1}$, $x = \text{fly}_{T_2}$. 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.

2) *Datasets*: As seen in Table IV, the first two benchmarks (ABX-15 and ABX-17) used relatively modest datasets sizes for training (from 2.5h to 45h) over 6 languages. ABX-15 used the training set as test set, while the ABX-17 introduced a separate test set with new speakers, to test for the ability of the learned representations to generalize to new speakers. The more recent benchmark (ABX-LS) uses the full LibriSpeech (960h) as training set, and introduced a separate dev and test set in order to avoid overfitting the model’s hyperparameters to the test set. As more recent models tend to use more and more data, the benchmarks are open to submissions of systems trained on datasets other than the default ones, so long as they contain no labels and are described in detail.

3) *Baselines*: For Task 1, our low end reference score (“baseline”) is calculated by computing the distances on spectral representations (MFCC). Good representations that highlight linguistically relevant differences, and thus neutralize speaker or channel differences, should at least do better than MFCCs. On the high end, using the gold annotations to generate a frame-by-frame one-hot phonemic representation mechanically leads to a perfect ABX score. To give a fairer high end comparison, we have also often included scores for the output of an off-the-shelf supervised GMM-HMM ASR. We included such a “topline” score in the **ABX-15** and **ABX-17** results, as well as in the unit evaluation component of the **TTS0-19** evaluation, which also contained an ABX test (see below). In the first two cases, the representations we

¹The risk of the time stamping approach is that some models’ representations may warp time in unpredictable ways, resulting in apparent errors due to misalignment with the time-stamps presupposed by the evaluation. On the other hand, providing very short audio clips at test time may cause its own problems if models have been trained with longer sequences.

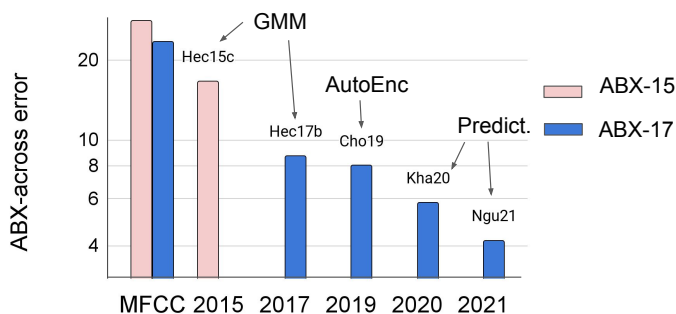


Fig. 2. ZR Task 1 results on English ABX test sets (ABX-15: Conversational speech–Buckeye; ABX-17: Audiobooks–LibriVox). The left two scores are on MFCC representations. The right two scores have been trained on Librispeech 960. See Table III for references.

used were posteriorgrams (that is, rather than a one-hot vector at each frame representing the decoding, we calculate the model’s posterior decoding probability for each frame). This reference score was beaten not long after the 2017 challenge (by **Cho19**), and remains quite poor compared to modern systems (see below). In the **TTS0-19** evaluation, we instead used the hard decoding, rather than probabilities, and observed that this reference score was in fact *very* easy to beat (for one of the two languages, the MFCC scores were already better), because of the numerous errors in decoding made by the off-the-shelf ASR system. More recently, submitted systems have become so good at the ABX task (see Figure 2 and below) that such low-fidelity topline systems have become less relevant.

4) *Results:* Since 2015, several approaches have been taken to Task 1. Most start from the principle that a central characteristic of text (or phonemic transcription) is that it provides a highly *compressed latent representation* for speech. For reference, a 16 kHz 16-bit waveform is coded using 256 Kb/sec, which generic audio codecs like Opus or MP3 can compress down to between 32 and 16 Kb/sec (a factor of 8 to 16). In contrast, a phonemic transcription is about 70 b/sec. This represents a compression of more than 200x compared to general audio codecs! Many objective functions proposed for Task 1 have as their primary goal to reduce the amount of information coded.

A simple and remarkably successful version of this idea—inspired by the *universal background models* used in speaker encodings—is to model acoustic frames using a mixture of full-covariance Gaussians (**GMM**). The posteriorgram of the mixture is taken as a new, sparse code for the speech input. In other words, each acoustic frame in the input file is assigned a sparse vector of probabilities, which correspond to the a posteriori probability of each of the discovered Gaussian distributions as the source of the given frame. Since individual frames are very short, and they are clustered as independent observations, this gives rise to a code which classifies speech in terms of short-term acoustic events, typically with several hundred Gaussian clusters discovered. This strategy, supplemented with additional speaker normalization or teacher-student tricks, was able to obtain top scores in the 2015 and 2017 editions (**Che15a-d**, **Hec17a,b**: see Figure 2).

Another type of approach seeking to find a compressed

latent representation uses *autoencoders* (**AutoEnc**), which aim at reconstructing the signal through an information bottleneck, sometimes achieved by using a discrete codebook: **Bad15a-c**, **Cho19**. The codebook + WaveNet autoencoder of **Cho19** obtained better results than previous, mixture model-based systems.

Since 2020, a new generation of *predictive models* (**Predict**) began to appear which have given rise to never-before-seen performance: contrastive predictive coding, or CPC [74], wav2vec 2.0 [75], HuBERT [76]; see [77] for a review. Two salient differences with this new wave of models is how they integrate context, and how they scale by working from the raw audio. The compression based models tended to have a frame-based view of the speech signal, modeling the probability distribution of individual acoustic frames through a compressed latent representation. In contrast, the predictive models aim to reconstruct large missing or masked parts of the signal *conditional* on visible parts of the sequence. For instance, CPC predicts future frames within a 10 to 120 ms window based on past frames, and obtained excellent ABX scores (**Kha20**, **Ngu21**, **BAS4-sm,lg**: around 4.5% across speaker). See a more thorough discussion of CPC in the section about Task 4 below. Wav2vec 2.0 and HuBERT try to reconstruct a masked part of the signal (of the order of 100ms), based on left and right context.

Independently of their predictive objectives, these models are also more sophisticated than previous ones in their encoding of context. Instead of processing signals within small receptive fields (**Bad15a-c**, **Thi15**, **Che15a-d**), the new systems use recurrences, transformers, or both, allowing them to model temporal correlation at longer distances. At the phonological level, language can be viewed as an autocorrective code due to redundancies introduced by phonotactics and lexical regularities. Previous work showed that top-down lexical context [22], [78] or phonotactics [79], [80] can indeed help with the discovery of phonetic units. This may explain why predictive/masked objectives together with large receptive fields help learning the acoustic properties of units jointly with their functional role in the language, yielding more relevant units.

The new models are also large, and, accordingly, are typically trained on large amounts of data (thousands of hours), which is orders of magnitude more than the training sets used in the initial ZRC series. In addition, some of the new models work *directly from waveform* instead of relying on engineered features like MFCC or mel filterbanks. Allowing models to be large, working from the raw audio, and training them on large amounts of data might push them to mimic the evolution of the human auditory system and its adaptations to speech. Indeed, [81], [82] find that wav2vec 2.0, HuBERT, and, particularly, CPC, are good predictors of low-level (sub-phonemic) auditory and speech processing in humans. In addition, it is well known that human perception relies on temporal fine structure not captured by magnitude spectrograms [83], particularly in difficult listening conditions. Models working from waveform

might have an advantage.²

To sum up, predictive models seem to have an edge on compressive models, and enjoy increasing popularity for a variety of downstream applications (see models presented in the SUPERB benchmarks [3], [85]). In the context of the ZRC series, a fair comparison equating architecture and dataset size would be necessary before claiming a definitive win. In addition, new models combining the two ideas like Masked AutoEncoders are emerging and need to be tested [86], [87].

Other interesting ideas have been explored in the ZRC series. Although they have not made it to the top of the leaderboard, they may still have much to contribute, perhaps in combination with other approaches. For example, some systems have attempted to use possible synergies with Task 2, and have used Spoken Term Discovery to obtain pairs of segments that have potentially the same phonemes, and use them in a **Siamese** architecture. Such pairs have also been used as a form of data augmentation with some **AutoEnc** models as well. In addition, most systems have not attempted to model the duration of linguistic units like phonemes or syllables yielding representation of much shorter duration acoustic events (10 ms or so). Yet duration is a principal concern of the HMM-based unit discovery system of **Baseline3** as well as the segmental CPC approach of **Bha21a**. The approach of **Pan19a–b, Kum20a–b** implicitly considers duration by dividing the signal into syllables, which are then further divided into subsyllabic units (see also Task 4 in Section II-D). None of these other approaches has reached state-of-the-art performance yet, but, once again, duration is quite clearly critical to speech perception, and so it seems likely that research will need to examine these ideas further.

B. Task 2: Spoken Term Discovery

Just as the infant learns the words of their language by listening, Spoken Term Discovery seeks to find recurring patterns in the input, proposing a set of boundaries delimiting the start and end of proposed word tokens discovered in the speech, and category labels indicating proposed word types.³ This problem was explored by several papers prior to the ZRC [88]–[92], and served as inspiration for the challenge itself. Although the task of “finding words” seems intuitively simple, it is made up of at least three subproblems which we evaluate separately.

- The *matching* subproblem is to find all pairs of speech fragments that are instances of the same sequence of phonemes. This can be evaluated based on how phonemically similar the discovered fragments are using the gold transcription (normalized edit distance: **NED**) and how much of the corpus they cover (**coverage**).
- The *lexicon discovery* subproblem is to group these fragments into clusters (as opposed to simple pairwise matching). The goal is to find a lexicon of types. A proposed

cluster can be evaluated based on how well the members match on the sequence of phonemes (**Grouping**) and how well the sets match the gold-standard lexicon of word types (**Type F-score**).

- The *word segmentation* subproblem attempts to find onsets and offsets of fragments that are aligned with the word boundaries as defined in the gold-standard text.

1) *Evaluation metrics*: To maximize comparability with text-based word discovery approaches, all of these evaluations are done by forced aligning the test set with its phoneme transcription. Any discovered speech fragment is converted into its transcription (which means taking decisions about phonemes on the left or right edge which may be partially covered: we include a phoneme if the fragment contains more than more than 30 ms of that phoneme or more than 50% of its duration).

The evaluation of spoken term discovery systems as matching systems consists of two scores, **NED** (normalized edit distance) 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 ($ED(a, b)/\max(|a|, |b|)$, where a and b are the two elements of a proposed match). The **coverage** is the fraction of the discoverable part of the corpus that is covered by the set of all 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 n-grams (with n between 3 and 20). This is almost all of the corpus, except for unigram and bigram hapaxes.

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**, **grouping recall**, and **grouping F-score**. These are defined in terms of P_{clus} , the set of all pairs of fragments that are groups in the same cluster, and P_{goldcl} , 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.

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

$$\text{Rec: } \sum_{t \in \text{types}(P_{\text{goldcl}})} w(t, P_{\text{goldcl}}) \frac{|\text{occ}(t, P_{\text{clus}} \cap P_{\text{goldcl}})|}{|\text{occ}(t, P_{\text{goldcl}})|} \quad (3)$$

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**, **type recall**, and **type 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 phonemes long.

To evaluate systems on the word segmentation subproblem, we use the **token** and **boundary F-scores** with respect to the gold text as is usual in text-based word segmentation. The token *F-score* evaluates whether the set of discovered tokens

²We note, however, that, as [84] demonstrates, wav2vec 2.0’s use of temporal fine structure still differs from that of human listeners.

³Note here, that contrary to Task 1, we ask systems to explicitly return linguistically interpretable representations (boundaries and cluster labels) instead of a representation that simply satisfies certain functional properties.

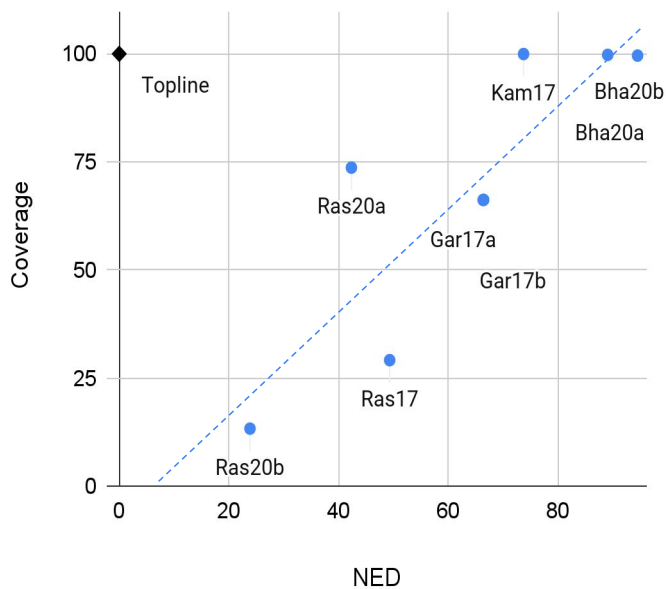


Fig. 3. Task 3 (term discovery): NED vs Coverage Tradeoff. Average results across 5 languages (ZR17 and ZR20).

matches the gold set of tokens, and the boundary F -score evaluates whether the set of boundaries (the delimitations between tokens, in terms of which phones in the gold transcription are separated by a boundary and which are not) corresponds to the gold set of boundaries.

The baseline system we provide is that of [32], which matches acoustic pairs using locally-sensitive hashing and then groups the pairs together using graph clustering. The topline was based on applying the adaptor grammar segmentation algorithm of [93] to the gold textual transcriptions.

2) *Datasets*: The datasets for the 2015 and 2017 Task 2 benchmark (TDE-15 and TDE-17, see Table IV) are coextensive with those of the corresponding Task 1 benchmarks, with one exception: the test set is always same as the training set. This may seem unorthodox from a machine learning point of view, but is quite common in the text segmentation literature, as the models are evaluated on their ability to extract words and boundaries from the training set.

3) *Results*: The Spoken Term Discovery task is still very challenging and has not received the same attention as Acoustic Unit Discovery. One major finding across the three ZRC editions that featured this task is the existence of a **tradeoff** between attempting to find a lot of words and ensuring that the discovered words are accurate. The quality of the set of words that are treated as matches/repetitions by the system, as measured by the normalized edit distance (NED), will necessarily be better if systems do not commit to extracting more dubious word candidates in the first place; however, the more candidates are ignored, the less of the corpus will receive an analysis (lower coverage) and the fewer of the gold word boundaries will be found (leading also to lower boundary F -scores). The tradeoff between term quality and coverage is shown in Figure 3.

Systems that take a “matching first” approach, like **Ras15a-c**, **Ras20a-b**, seek primarily to find recurrent pho-

netic patterns. Boundaries here are merely designations of the edges of these discovered segments. The system that currently does best at balancing NED with coverage and segmentation quality, **Ras20a**, takes a matching-based approach, based only on MFCC inputs. This system begins by doing a low-resolution search for candidate matches by dividing the input utterances into fixed-length down-sampled speech segments. Then, it filters the candidate matches using a higher-resolution matching algorithm based on dynamic time warping.

It might seem surprising that an algorithm that uses MFCC inputs, rather than features learned by acoustic unit discovery, would yield good performance. Nevertheless, [94] demonstrated that ABX error rate may not be the best indicator of downstream lexicon discovery, and our own informal experiments have shown that naively feeding improved acoustic units into a generic matching system (for example, those learned by **Thi15**) can actually make matching quality *worse*.

Other Task 2 systems take segmentation-oriented approaches, putting a priority on discovering boundaries. Building on earlier text-based systems using non-parametric Bayesian models like [93], [95], [96], systems like **Kam15**, **Kam17** jointly optimize an exhaustive segmentation and a dictionary of clustered word embeddings. In a different, bottom-up approach, **Bha20a,b** matches learned segmental acoustic units to construct a full word segmentation. As would be expected, since these systems strive to optimize segmentation measures, they fare rather poorly on matching measures.

Figure 4 focuses on segmentation by itself and displays the Token F -score for each of the submitted systems, compared to a topline unigram adaptor grammar segmentation system trained on the corresponding text (phonemized text without the blank spaces between words). All of the high-coverage segmentation-oriented models are on the left and all of the low NED, matching-first models on the right. The segmentation-oriented models are more likely to do well on this metric, which assesses how many of the true word tokens were correctly segmented. Included here are two new models, **Kam22** and **Alg22**, which do not even attempt to build a lexicon of types. The first one is in the same vein as [97], which posits word boundaries at peaks in surprisal across sequences of learned segmental units, while the second uses a non-parametric Bayesian approach directly on tokens. As can be seen, however, the gap between the best speech based models and the text-based ones is still large.

The reason is likely multi-fold. First, speech is variable, which means that the same word will surface with a variety of acoustic shapes. Even if good invariant quantized acoustic representations are used, the potential (and actual) variability in the different “transcriptions” in terms of these quantized units for the “same” word grows exponentially with word length. This makes it difficult to build a reliable lexicon of word types. Second, speech rate and phone durations are variable in time, with the result that both phoneme duration and word duration can change substantially from occurrence to occurrence, a problem that does not exist in text. Finally, speech is typically coded into frames, which gives it a finer granularity than text (e.g., 10 ms frames, whereas phonemes

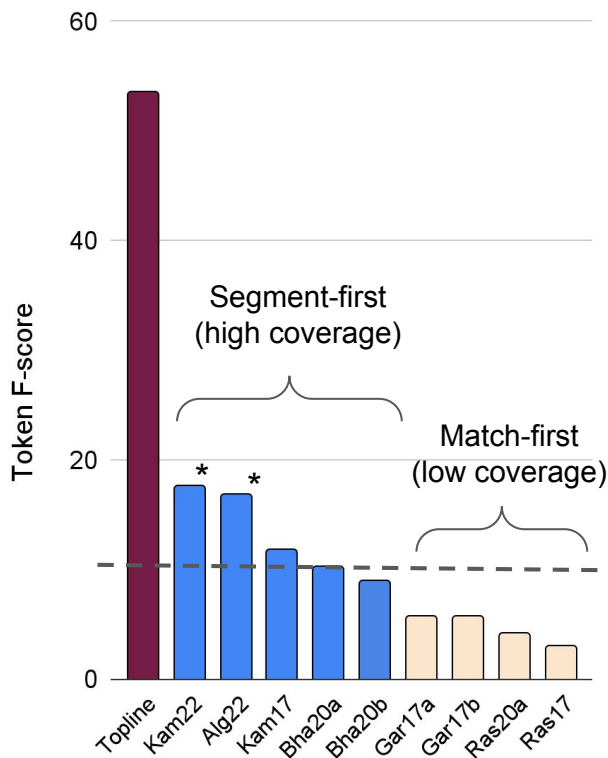


Fig. 4. Task 2 (term discovery): Token F-scores, measuring how many words are correctly segmented, averaged across 5 languages (ZR17 and ZR20 plus two new papers). The topline is a unigram word segmentation adaptor grammar trained on the same amount of text. The dotted line is a baseline consisting in random segmentations every 120ms. Starred (*) models compute the segmentation without even building any lexicon of discrete types.

last on average around 70 ms): a potential word boundary can therefore occur in more places in speech than in text. This increases the number of potential segmentation errors that can be made. This last point is one of the motivations for systems such as **Kam22**, **Bha21a**, **Cue22**, all of which, jointly or sequentially, infer word boundaries hierarchically on the basis of learned acoustic unit boundaries. Any future work will need to address all three of these challenges to achieve better performance on this task.

C. Task 3: Discrete Resynthesis (TTS without T)

Here, we investigate a task which is similar to what infants may do when they repeat a word or a sentence: they encode the signal into some representation, and then reproduce the same content in their own voice. Defined like this, the task is already known as voice cloning or voice transfer, and it can be performed at a rather low level by introducing a target speaker embedding in the decoder part of a simple encoder–decoder architecture. Here, however, we add the constraint that there be a discrete bottleneck between the encoder and decoder, and we measure the bitrate of the encoding. In other words, we ask participants to use discovered acoustic units instead of phonemes, and we push these units to approach the bitrate of phonemic transcription. Prior to the ZRC, [98] demonstrated the feasibility of unsupervised discrete resynthesis. Furthermore, some of the models in Task 1 (**Bad15a–c**, **Cho19**)

already used a similar discrete bottleneck autoencoder architecture, although they did not evaluate the quality of the reconstruction nor the bitrate of the representation. Participants on this task are provided with a *unit dataset* from multiple speakers used to discover discrete units, and a *voice dataset* to train a synthesizer for the target voice taking the units as input. The *test dataset* consists of novel utterances by unseen speakers, which must be resynthesized in the target voice. Participants submit both the acoustic unit representation and the resynthesis for evaluation.

1) *Evaluation metrics*: As for the acoustic units, they are evaluated in terms of their **bitrate**, where each unique embedding value is counted as a single symbol type. For the **bitrate** computation, each vector is processed as a character string. A dictionary of the possible values is constructed over the embedding file for the submitted test set. We thus assume that the entire test set corresponds to a sequence of vectors U of length n : $U = [u_1, \dots, u_n]$. The bit rate for U is then $B(U) = \frac{n}{D(U)} \sum_{i=1}^n p(u_i) \log_2 p(u_i)$, where $p(u_i)$ is the probability of symbol u_i in U . The numerator is n times the entropy of the symbols, which gives the optimal number of bits needed to transmit the sequence of symbols $s_{1:n}$. To obtain a bitrate, we divide by $D(U)$, the total duration of U in seconds.⁴ Not reported here, Task 3 also included a version of ABX in which the minimal triphones (“fly” versus “fry”) were extracted and presented as small audio clips, in order to not penalize the evaluation of sequence-to-sequence systems that would lack the alignment between units and audio signals. As it happened, no such system has been submitted as yet.

As for the generated waveforms, native speakers of the test languages were recruited online to evaluate the quality of the synthesis in terms of intelligibility, naturalness, and similarity. *Intelligibility* was measured by asking participants to orthographically transcribe the synthesized sentence. Each transcription was compared with the gold transcription using the Levenshtein distance, yielding a Character Error Rate (**CER**). The overall *naturalness* of the synthesis was assessed on a 1 to 5 scale, yielding a Mean Opinion Score (**MOS**).⁵ *Speaker similarity* was assessed using a 1 to 5 scale. Sentences were presented in pairs (target voice, system voice).⁶ Each sentence token was evaluated at least once with each system, as well as the original (gold) recordings.

2) *Datasets*: The 2019 Benchmark for Task 3 (TTS0-19) provides training and test data for two language: English (the dev language) and Indonesian (the test language). For each language, one “Unit” dataset is provided to train unit discovery (around 15h, between, 100 and 120 speakers), and one “Voice” dataset is provided to train speech synthesis in the target voice

⁴A fixed frame rate transcription may have a higher bitrate than a “textual” representation due to the repetition of symbols across frames. For instance, the bitrate of a 5 ms framewise gold phonetic transcription is around 450 bits/sec and that of a “textual” transcription around 60 bits/sec.

⁵The question posed was: *Rate how natural the audio is, between 1 and 5 (1=very unnatural, 3 = neutral, 5=very natural).*

⁶The question posed was: *Rate the similarity between the reference voice and the system voice, between 1 and 5 (1 = very different voices, 3 = neither similar nor different voices, 5 = very similar voices).* Ten additional trials were included, for each participant, in which the reference voice was not the target voice but the source voice.

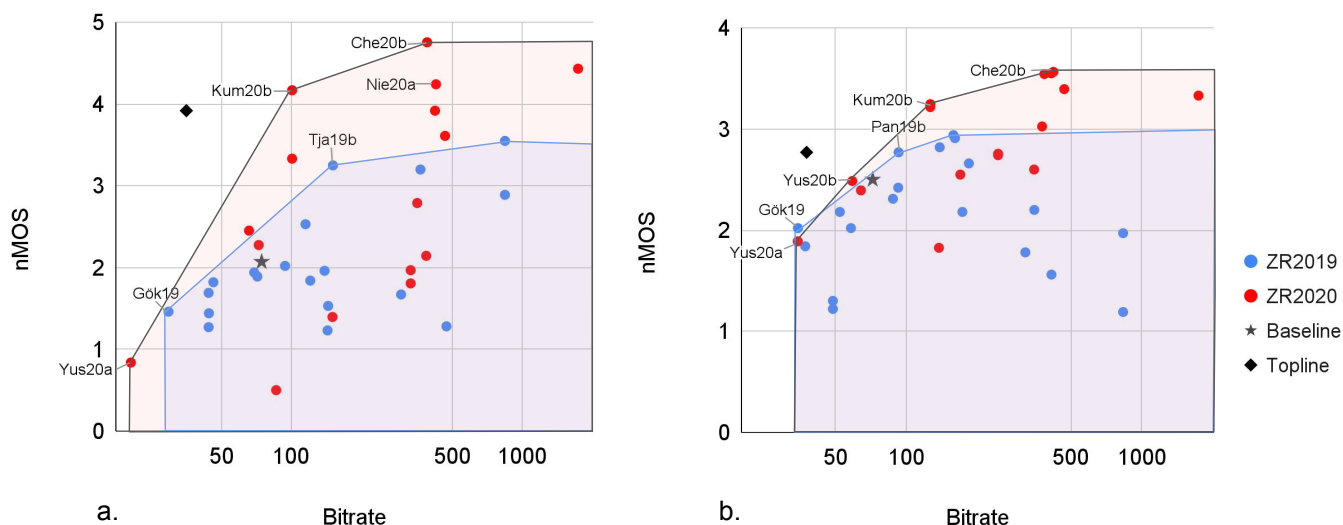


Fig. 5. ZR Task 3 (TTS without T) human evaluation results on the surprise language (Indonesian, panel a) and the dev language (English, panel b) across the two editions of the Challenge (2019 and 2020). The nMOS score is obtained by normalizing the Mean Opinion Score across the two challenges by using the Baseline and Topline as anchor points.

(two voices in English, one in Indonesian, between 1h30 and 2h30 per speaker). The test set consists in new utterances by new speakers (around 30 minutes, 15-24 speakers; see Table IV for detailed numbers). None of these datasets are provided with labels except an anonymous speaker ID.

3) *Baselines*: The baseline system consists of an existing acoustic unit discovery system which discovers GMM-HMM models and clusters them using an unsupervised Bayesian approach (see [43]). We then use decoding from this system (i.e., sequences of unit labels) instead of phonemes, and train an out-of-the-box speech synthesizer (Merlin, with the Ossian front end [99]). For the topline system, we replace the unit discovery with an off-the-shelf GMM-HMM ASR system.

4) *Results*: The performance was overall quite good, with several systems achieving better resynthesis than the text-based topline. As shown in Figure 5, there is a general tradeoff between synthesis quality and bitrate, which held both in the dev language (English) and in the heldout surprise test language (Indonesian). As shown by the black point in the figure (the decoded output of a simple phone recognizer), phonemic transcription is a highly-compressed representation of speech which is excellent for this task (the middling MOS scores are, as for Task 1, attributable to the fact that the out-of-the-box ASR and TTS were not optimized to the task).

Many of the systems that have a **low bitrate** (under 100b/sec) learn a discrete autoencoder on acoustic features (**Kam19a-b**, **Yus19**, **Gök19**, **Liu19a-b**, **Gün20**), generally taking further steps such as filtering or downsampling to reduce the temporal resolution. Taking a slightly different approach, our baseline model, as well as the related **Yus20a-c**, learn latent HMMs as acoustic units, in order to explicitly model duration. On the other hand, **Pan19a-b**, **Kum20a,b** put temporal reduction in an initial step of acoustic segmentation based on syllable-like units. Among these models, **Kum20b**, which presegments and then learns HMM acoustic units, stands out as reaching performance comparable to higher bitrate models

(it admittedly has a somewhat higher bitrate than the other models listed here). Syllable-like presegmentation, as noted above, has also been used productively in Task 2 by **Ras17**, **Alg22**. It is fair to say that syllables have been underutilized in zero resource speech processing, given their promise.

Most of the remaining systems have a **high bitrate** between 100 and 600b/sec. Supervised posteriorgrams are on the upper end of this, and MFCC representations have a bitrate around 1500.⁷ Most of the submitted systems in this range are compression approaches using discrete autoencoders, including the system of **Che20b**, which gives excellent performance. The system of **Nie20a,b** stands out among the others as yielding high quality results. This is the only submitted system which uses a predictive loss based on CPC—although, unlike typical CPC models, it works from spectrogram and is trained on the small (15h) dataset provided for the 2019 edition.

The results of [100] also support the claim that CPC and related approaches are well-adapted to discrete resynthesis. In addition, [100] demonstrated that an automatic evaluation using ASR is strongly predictive of human evaluators' ratings, and that the discrete representations can be used to support learning a language model.

D. Task 4: spoken LM

Spoken language modeling is the task of learning a language model directly from audio. Such a model could be end-to-end, learning directly from speech, or it could take as input discrete or continuous representations from Task 1 or word level representations from Task 2—so long as these input representations are learned without supervision from text or other labels. The task can be understood as the modeling of the probability distribution of spoken utterances in an unknown language.

⁷Note that our bitrates are calculated empirically on a rather small speech corpus. The representational *capacity* of MFCC representations is an order of magnitude higher.

1) *Evaluation metrics*: For Task 4, the evaluation problem is severe. Language models trained from text are typically evaluated by the perplexity over a test corpus, or by fine-tuning on downstream tasks. As discussed above, the ZRC series focuses evaluation on zero-shot tasks that require no training. This excludes a fine-tuning evaluation. As for perplexity, in text-based systems, it is derived from the conditional probability distribution of the next token given a past sequence of tokens. In speech-based systems that use discrete pseudo-text units, the number of such units is a latent variable, making the perplexities difficult to compare across models. The problem becomes worse for systems that do not use discrete representations at all, where the estimation of the conditional probabilities themselves becomes model dependent. The two editions of the ZRC 2021 used a battery of 4 metrics, each one measuring performance at a different linguistic level: acoustic, lexical, syntactic and semantic.

At the level of *acoustics*, we the ABX-LS benchmark as defined in Section II-A, built on top of LibriSpeech dev and test sets (see [101], [102]).

At the *lexical* and *syntactic* levels, instead of computing an average perplexity across a corpus, the ZRC uses a contrastive approach, where a “pseudo-probability” \hat{p} is computed for minimal pairs of utterances—one grammatically legal, the other illegal. The pseudoprobability can be obtained from a language model by decomposing the probability of an utterance U into a product of conditional probabilities of each of its constituent units u_i : $\hat{p}(U) = p(u_1)p(u_2|u_1) \dots p(u_N|u_1 \dots u_{N-1})$, or by computing an average perplexity score or of the loss function over the utterance U . An accuracy Acc is computed by counting how often \hat{p} is higher for the legal than for the illegal utterance:

$$Acc(T) = \frac{1}{|T|} \sum_{(a,b) \in T} (\mathbb{1}_{\hat{p}(a) > \hat{p}(b)} + \frac{1}{2} \mathbb{1}_{\hat{p}(a) = \hat{p}(b)})$$

where T is a test set containing pairs of audio, one legal (a), one illegal (b); the chance level is 0.5. To probe the lexical level, pairs of well matched words versus nonwords (e.g. “brick” versus “blick”) are constructed using the Wuggy non-word generator [103]. The syntactic levels is probed by using pairs of grammatical and non grammatical sentence derived from the BLIMP dataset [104]. All stimuli are synthesized using the Google TTS API resulting in the sWUGGY and sBLIMP test sets, respectively (see below for details).

Finally, ZRC evaluated the *semantic* level by using a similarity probe task used to investigate word embeddings [105]. It correlates the similarity of systems’ representations of words with human similarity judgments. This enables us to measure the extent to which the model is able to extract lexical semantic knowledge. As for the ABX task, participants provide embeddings for input tokens as well as a distance to compute similarity. The Spearman rank correlation is calculated between the dissimilarity scores provided in the submission test set (sSIMI), $d(a, b)$, and the dissimilarity scores given by human judgments, $d_h(a, b)$. The challenge provided by default the cosine distance computed over pooled embeddings (with mean, max or min pooling).

2) *Datasets*: There is only one benchmark (sLM-21) associated with Task 4. The default training set is LibriSpeech 960h, although participants can use other training sets so long as no labels is provided besides speaker ID. The test set is split into sWUGGY, sBLIMP and sSIMI, that evaluate the lexical, syntactic and semantic levels, respectively. These test sets are described in details in [59] and only briefly summarized here.

The sWUGGY dev and test sets consists of 5,000 and 20,000 pairs of words and nonwords respectively, with the existing words being part of the LibriSpeech train vocabulary. There is also an additional OOV-sWUGGY dev and test set consisting of 5,000, and 20,000 pairs respectively, with existing words which do not appear in the LibriSpeech training set. The nonwords are produced with WUGGY [103], which generates, for a given word, a list of candidate nonwords best matched in phonotactics and syllabic structure, which were additionally filtered for pronouncability using G2P, and for having on average the same unigram and bigram phoneme frequencies as words. Waveforms were produced with the Google Speech API.

The sBLIMP dev and test sets are adapted from BLIMP [104], a set of linguistic minimal sentence pairs of matched grammatical and ungrammatical sentences. The dev and test sets contain 6,300 and 63,000 pairs respectively, with no sentence pair overlap. Stimuli were filtered to contain LibriSpeech vocabulary and for natural prosodic contours, and synthesised as above.

The sSIMI dataset was constructed out of 13 existing semantic similarity and relatedness datasets: WordSim-353 [106], WordSim-353-SIM [107], mc-30 [108], rg-65 [109], Rare-Word (or rw) [110], simLex999 [111], simverb-3500 [112], verb-143 [113], YP-130 [106] and the relatedness-based datasets include MEN [114], Wordsim-353-REL [107], mturk-287 [115], and mturk-771 [116]. All scores were normalised on a 0-10 scale, and pairs within a same dataset containing the same words in different order were averaged. Pairs containing a word absent from LibriSpeech train set [117] were discarded. The mturk-771 dataset was set aside as a dev set and the other 12 datasets were used as test sets, after removing overlapping pairs across dev and test sets. Given the unequal size of the test sets, the ZRC Benchmark introduced a weighted average of the Spearman scores which we report here. Two subsets of audio files, one synthetic, and one natural, were created, the latter being obtained by extracting the audio sequences corresponding to each word from LibriSpeech, as in [105]. In this subset, each word can appear in multiple tokens, providing phonetic diversity; duplicated *scores* are averaged in the analysis step. The natural subset is smaller than its synthetic counterpart, as we had to discard pairs from the test and dev sets which were not present in the LibriSpeech test and dev sets respectively. The synthesized subset is composed of 9744 and 705 word pairs for the test and dev sets respectively, and the LibriSpeech subset is composed of 3753 and 309 pairs for the test and dev sets.

3) *Baselines*: Our baseline system, described in [59], is based on first training a contrastive predictive coding (CPC) model. We review the CPC acoustic model [74] here for clarity. Given an input waveform \mathbf{x} , the *encoder* component

TABLE V

LEADERBOARD UPDATE FOR DECEMBER 2021. TEST SET RESULTS ON THE FOUR METRICS FOR ACOUSTIC (ABX), LEXICAL (sWUGGY), SYNTACTIC (sBLIMP) AND SEMANTIC (sSIMI) LEVELS RESPECTIVELY. THE FIRST THREE SYSTEMS ARE FOR REFERENCE AND INCLUDE THE TWO BEST SYSTEMS OF THE FIRST 2021 ZRC EDITION [14]. THE LAST THREE SYSTEMS ARE **VISUALLY GROUNDED (VG)**. THE REMAINING SYSTEMS ONLY USE AUDIO INPUT. THE BEST OVERALL SCORE IN EACH COLUMN AMONG ALL SUBMITTED SYSTEMS IS BOLDDED.

System	Budget	ABX-with.		ABX-across		sWUGGY	sBLIMP	sSIMI	
		clean	other	clean	other			synth.	Libri.
Character-based (RoBERTa)	24576	-	-	-	-	96.25	82.11	23.22	21.08
Best March 2021 Low-budget [Nie21]	60	5.41	8.67	6.89	13.14	72.86	53.59	5.60	2.72
Best March 2021 High-budget [BAS4-1g]	1536	3.28	4.81	4.31	7.92	75.51	56.16	3.19	1.32
Gan21	60	5.04	7.89	7.08	14.01	59.87	52.20	7.47	4.38
Ngu21a	2560	3.49	5.72	4.56	9.19	77.22	55.62	8.43	4.48
Ngu21b	2112	3.80	4.86	4.70	7.70	79.13	58.89	8.62	6.65
Ngu21c	3424	3.03	3.62	3.83	5.63	80.19	59.29	6.8	3.08
Ngu21d	2720	3.49	5.72	4.55	9.20	80.29	59.93	9.43	5.33
Lan21	60	2.85	4.44	3.67	7.33	70.23	51.93	3.09	2.05
Bha21a	60	3.49	5.72	4.51	9.16	50.90	52.10	3.44	1.67
Bha21b	60	3.49	5.73	4.57	9.19	48.55	52.80	1.57	5.2
Gao21a	60	6.71	10.62	8.41	15.16	62.29	53.17	1.67	1.13
Gao21b	60	6.60	10.86	8.51	16.51	68.42	54.02	2.61	1.78
Gao21c	35	7.45	11.76	10.94	19.31	62.29	53.17	5.07	1.2
Pen21 (VG)	468	4.24	5.22	5.08	7.91	75.23	57.40	17.99	12.78
Lee21a (VG)	7	6.50	9.95	9.17	15.46	51.91	50.43	4.51	5.54
Lee21b (VG)	7	6.50	9.95	9.04	15.44	51.91	50.43	8.1	12.21

of the model maps it to a sequence $\mathbf{z} = (z_1, \dots, z_T)$. An autoregressive component then predicts the future, taking z_1, \dots, z_t and outputting a latent representation c_t , which is a representation of the context. Given the context c_t , CPC tries to predict the K next future embeddings $\{z_{t+k}\}_{1 \leq k \leq K}$ by minimizing a contrastive loss:

$$\mathcal{L}_t = -\frac{1}{K} \sum_{k=1}^K \log \left[\frac{\exp(z_{t+k}^\top W_k c_t)}{\sum_{\tilde{z} \in \mathcal{N}_t} \exp(\tilde{z}^\top W_k c_t)} \right] \quad (4)$$

where \mathcal{N}_t is a random subset of negative embedding samples, and W_k is a linear classifier used to predict step k of the future.

Our baseline system then clusters the resulting framewise representations (as independent observations) using k-means, to reduce them to 50 units. The resulting discrete sequences are passed as input to a character-based language model. We experimented with both BERT and LSTM models, and found that large BERT models performed best.

4) *Results:* The first round of submissions was documented in 2021 [14]; the best-performing systems were variants of our baseline system. A second round was opened as a NeurIPS 2021 challenge, including a visually-grounded training option. Briefly, this modified scenario expands the range of data that models can be trained on, to include multi-modal datasets (like speech and image, or speech and video). The rationale is that young children learn in a multimodal, multisensory environment rather than by just listening. Some earlier models of word discovery and representation learning demonstrated the feasibility of such multimodal training [118]–[120]. Following [60], Task 4 was expanded to include “visually-grounded” training. Participants were to indicate the dataset they used. Systems were only tested with speech-only inputs, however, for comparability with non grounded systems. Here, we present for the first time the results of these latest submissions to Task 4 (see Table V).

Similar to the baseline models, the systems of **Gan21**, **Ngu21a,d**, **Bha21a,b**, and **Gao21a-c** take the approach of training acoustic units and then constructing a language model on their outputs. The distinction between *high-budget* systems and *low-budget* systems is made the basis of the number of GPU hours needed to train the language model. **Gao21a-c** apply segmentation and pooling to reduce the temporal resolution of the units, while **Bha21a,b** use Segmental CPC to learn units and segmentation jointly. **Ngu21a,d** are technical improvements on the previous best system **BAS4-1g**. On the other hand, **Ngu21b,c** are HuBERT systems, trained end-to-end on a masked language modelling task.

The systems of **Pen21** and **Lee21a,b** are visually grounded. In the case of these two systems, that means they both start from acoustic units that are trained using parallel speech–image data (picture captions). One difference between the two models is the type of training—**Pen21** trains end-to-end on a masked language modelling objective, while **Lee21a,b** use the pre-trained features as input to a small BERT.

Task 4 is clearly in its very early stages (this in spite of the excellent ABX performance of the units used in systems up to now). However, even at this stage, after only one year’s worth of submissions, spoken language modelling has shown improvement on the spot-the-word task (moving from the best speech-based baseline’s 75% accuracy up to 80%) and on the syntactic judgment task (improving from 56% to 60% accuracy). The approach so far has been simple: high-quality units and a powerful language model. In the baseline models as well as most submissions, these components were trained separately; newer models like HuBERT [76] learn them jointly. The two approaches are currently tied for the top position on the leaderboard (**Ngu21a,d**, CPC units fed to a large BERT model, and **Ngu21b,c**, HuBERT systems). As for capturing word semantics, the Fast-VGS+ system of **Pen21** stands out as a serious competitor. This visually-grounded system takes advantage of spoken image caption data in training.

III. WHAT NEXT?

Over the six editions of the ZRC series, the following lessons can be drawn:

- Great progress has been made in Acoustic Unit Discovery (T1) due to recent breakthroughs in self-supervised representation learning showing good scaling properties in large corpora. The latent units discovered at this stage, though, are not interpretable linguistic units like phonemes but represent shorter-duration acoustic events.
- The Discrete Resynthesis task (T3) obtained excellent results, sometimes surpassing text-based systems in resynthesis quality—at the expense of bitrate, which remains about 4 to 8 times higher than phonemes or text.
- Spoken Word Discovery (T2) still remains disappointingly difficult in all of its three subcomponents (matching, clustering, and segmentation). Presumably, understudied effects of the acoustic and temporal variability inherent to speech still hampers current approaches.
- Spoken Language Modeling (T4) got surprisingly promising results, given that the task is complex, considering the difficulties found in Task 2, and the fact that most systems only worked from Task 1 units with sub-phonemic temporal granularity. There is however room for progress, given the gap between speech-based and text-based language models on syntactic and semantic tests.

Given the large body of results that have accumulated, it may now be useful to reflect on some of the basic assumptions and methods of the ZRC series to determine how to move forward. The assumptions are related to the architecture presented in Figure 1b and its corresponding task decomposition. The methods are related to the particular choice of metrics that were chosen to evaluate each of these tasks. We discuss in particular the role of Acoustic Modeling and the Lexicon.

A. Acoustic Modeling and ABX.

One of the basic assumptions of the ZRC series is the existence of an Acoustic Modeling component that turns speech input into a latent representation, which plays the role of phonemes or text in that it can be directly used as input to other processing components: the Lexicon, Waveform Generation, and possibly even Language Modeling. Methodologically, we proposed the machine ABX task as a metric to gauge the quality of this latent representation. This makes the prediction that there should be a strong positive correlation between ABX scores and the relevant metrics in the other components. Inspection of this correlation across tasks shows that such a correlation exists, but that it is in some cases weak.

- **T1 and T3:** a reanalysis of the 35 ZRC submissions shows a Pearson correlation coefficient of $r = .57$ and $r = .54$ (English and Indonesian, resp.) between ABX and intelligibility as measured by Character Error Rate (CER). The systems differed not only in the encoder but also in the decoder introducing noise in the correlation. A more controlled correlation across 9 systems with matched decoders [100] reported $r = .905$.⁸

⁸There is also a weaker (negative) correlation of ABX and bitrate (-.156 and -.277, in English and Indonesian, resp.) suggesting that a resynthesis objective may push away from a low bitrate phoneme-like representation.

- **T1 and T4:** we also find a reasonably high correlation between ABX and spot-the word ($r = .52$ across the ZRC submitted systems, and $r = .853$ in [100] in a more controlled comparison with matched language models).
- **T1 and T2:** the situation is confusing. Since the beginning, we have seen that these two tasks may require different representations. For instance, unit discovery worked well with MFCC, but word discovery worked better with PLP. Similarly, [94] showed that, across 16 types of word embeddings (supervised or unsupervised), ABX scores correlate only moderately well with two other proxies for word segmentation (frequency estimation: .53; Mean Average Precision: .45).

While the observed level of correlation may be sufficient to still use ABX as a proxy for comparing models before using them for downstream tasks, some caution is necessary, especially for the link between T1 and T2. One possible explanation for this discrepancy is that the assumption of a single acoustic level feeding all downstream tasks is wrong, and that there are instead several different acoustic codes with different properties. Alternatively, it could be that there is a single code, but that the ABX metric is not capturing the linguistic properties of this representation relevant for all the other tasks. While ABX was constructed to measure contrasts in minimal pairs of possible word across changes in speaker, it may not capture well other kinds of invariance (speaking rate, phonetic context) that are crucial for some of the other tasks. Further studies will be needed to sort out this question.

B. The Lexicon, discrete units and interpretability

One reason text-like representations—be they phonemic, alphabetic, or logographic—are fundamental to speech and language processing is that they serve a dual function. On the one hand, they record linguistically important properties of the *form* (what was actually uttered). On the other hand, they support straightforward analysis of the *content* (“meaningful” properties like morphology, syntax, and semantics).

While human listeners are sensitive to detailed, sub-phonemic properties [121], and while various gradations in lexical meaning can be observed [122], the two kinds of variability are not generally correlated. For example, although it is possible to pronounce the noun *sun* with an initial sound that would be intermediate between an /s/ and a /f/—making it sound somewhat more like the adjective *fun*—this gradient change does not evoke a concept of “slightly amusing star,” nor make the word more adjective-like. In other words, text-like representations would thus seem to be necessary for decorrelating form from meaning using an arbitrary mapping between a word’s phonological forms and its semantic or syntactic representations [123].

However, achieving this crucial decorrelation may not necessarily require that the representations be discrete, nor that they correspond to interpretable linguistic units or “words” as defined in a dictionary. First, linguistically, there are at least three distinct notions of ‘words’ (prosodic, syntactic and semantic [124]), which may not may not be aligned with

how dictionaries are constructed.⁹ Second, dictionaries are the result of a long cultural evolution where many design choices have been made that may not be consistent within or across languages. As a result, it could be that the requirement for T2 to provide segmentations and lexicons that are aligned with the written text is too strong. The fact that word-based units like BPE work well for text-based applications does not mean that the equivalent units for speech would align well with word boundaries. We could imagine in the future replacing the T2 metrics by new metrics reflecting the functional role of word segmentation, i.e., that of providing a level of granularity where arbitrary mappings between form and meaning can be learned (along the lines of the sSIMI metric in T4).

C. The future of the Zero Resource Speech Challenge

The submission site, www.zerospeech.com, is now open continuously and allows for running evaluations on all of the past benchmarks. The field of unsupervised representation learning is established enough that it is no longer necessary to channel it through special events. Indeed, self-supervised audio models are such an active domain that there are many relevant new models (for example, WavLM: [125]) which have yet to be evaluated on the ZRC metrics. Existing benchmarks, especially for Tasks 2 and 4, also still have a lot of potential for improvement, without creating more difficult tasks.

One exception to this is shown in Figure 1. Combining Task 3 and Task 4 leads naturally to consider the possibility of *generating* spoken language. A traditional spoken dialogue system will (conditional on some knowledge source) generate text, typically using a neural language model. The text is synthesized into speech. A spoken language model can be made to generate spoken language directly, as demonstrated by [67], [100], [126], [127]. Much as Task 3 is complementary to Task 1—but has slightly different constraints—the task of generating speech from a spoken language model is complementary to Task 4, yielding a potential Task 5. The evaluation of such a potential task may follow [100], replacing human evaluations of intelligibility and meaningfulness of by ASR-based proxy measurements (Phone Error Rate for intelligibility in a Task 3 setting, and continuation BLEU or VERT score for the prompted or unprompted generations).

Another reason for not declaring the ZRC series closed is that there is still a lot to understand on the evaluation side for existing tasks. For Task 1, recent research has shown that many discovered representations may not be speaker invariant [128] and differ from how humans perceive speech sounds [81], [129]. Further, their ability to perform out of domain (noisy environments, accented speech) has not been evaluated [31], [130].

IV. CONCLUSION: TOWARDS TEXTLESS NLP

Research on Acoustic Unit Discovery has led to a wave of new models using unsupervised pre-training to advance ASR. It also opens up the more radical possibility that one may

⁹As an example, “a piglet” could be considered a single unit prosodically, two units semantically, and three units, morphosyntactically (“a+pig+let”).

get rid of text altogether, and proceed with building language processing pipelines directly from raw audio. Up to now, this possibility has done little to change the dominance of text as the basic currency of NLP. One exception is translation, in which the idea of training machine translation directly from speech to speech in an end-to-end fashion has seen substantial uptake [131]–[135]. Other ways of removing text from the processing stack have also been explored, such as for language generation [100]. In general, however, converting speech to text and back remains the first and the last step of current speech-based NLP systems.

The allure of replacing text with low-bitrate unsupervised representations of speech goes beyond bringing NLP to lower-resource languages. Learned pseudo-text promises to be more flexible than traditional orthography: if the transcription system is the result of learning, it can also change to deal with new varieties and accents, and can learn to capture important linguistic information which is captured in a very limited way by typical writing systems, such as prosody. On the other hand, it promises to be more consistent: some writing systems are complicated by arbitrary exceptions, while other languages lack standardized conventions for spelling. Discovered representations could avoid both of these issues. Unlike traditional phonetic transcription, which uses a fixed, universal set of symbols which can in fact have rather different phonetic values across languages, unit discovery allows for a system to be adapted to the language.

The Zero Resource Speech Challenge has spearheaded efforts to build demonstrably useful unit discovery, as well as stimulating progress in applying these representations to more complex tasks. The major advances in building more realistic auditory-like representations have already borne fruit in recognition and synthesis. As we move toward better evaluations, we look ahead to the possibility of truly textless NLP—and a major key to unlocking cognitive models of human language development and speech perception.

ACKNOWLEDGMENT

For E. Dunbar, this research was supported by the Connaught Fund and the Arts and Science Bridging Fund, U. of Toronto, Natural Sciences and Engineering Research Council of Canada (NSERC) RGPIN-2022-04431, and by ANR grant ANR-17-CE28-0009 GEOMPHON. For E. Dupoux in his EHESS role and N. Hamilakis, it was supported by ANR grants ANR-17-EURE-0017 Frontcog, ANR-10-IDEX-0001-02 PSL* and ANR-19-P3IA-0001 PRAIRIE 3IA Institute, and by a Meta AI Research Gift.

REFERENCES

- [1] E. Dupoux, “Cognitive science in the era of artificial intelligence: A roadmap for reverse-engineering the infant language-learner,” *Cognition*, vol. 173, pp. 43–59, 2018.
- [2] E. L. Bavin, Ed., *The Cambridge handbook of child language*. Cambridge: New York: Cambridge University Press, 2009, oCLC: 435767607. [Online]. Available: <http://site.ebrary.com/id/10303044>
- [3] S.-w. Yang, P.-H. Chi, Y.-S. Chuang, C.-I. J. Lai, K. Lakhotia, Y. Y. Lin, A. T. Liu, J. Shi, X. Chang, G.-T. Lin *et al.*, “SUPERB: Speech processing universal performance benchmark,” in *INTERSPEECH*, 2021. [Online]. Available: <https://arxiv.org/abs/2105.01051>

- [4] T. Kohonen, "The 'neural' phonetic typewriter," *Computer*, vol. 21, no. 3, pp. 11–22, 1988.
- [5] B. Varadarajan, S. Khudanpur, and E. Dupoux, "Unsupervised learning of acoustic sub-word units," in *Proc. of ACL: HLT*, 2008, pp. 165–168.
- [6] M. Huijbregts, M. McLaren, and D. van Leeuwen, "Unsupervised acoustic sub-word unit detection for query-by-example spoken term detection," in *ICASSP*, 2011, pp. 4436–4439.
- [7] C.-y. Lee and J. Glass, "A nonparametric Bayesian approach to acoustic model discovery," in *Proceedings of ACL*, 2012, pp. 40–49.
- [8] A. Jansen, E. Dupoux, S. Goldwater, M. Johnson, S. Khudanpur, K. Church, N. Feldman, H. Hermansky, F. Metze, R. Rose, M. Seltzer, P. Clark, I. McGraw, B. Varadarajan, E. Bennett, B. Borschinger, J. Chiu, E. Dunbar, A. Fourtassi, D. Harwath, C.-y. Lee, K. Levin, A. Norouzi, V. Peddinti, R. Richardson, T. Schatz, and S. Thomas, "A summary of the 2012 CLSP Workshop on zero resource speech technologies and models of early language acqu.," in *ICASSP*, 2013.
- [9] L. Badino, C. Canevari, L. Fadiga, and G. Metta, "An Auto-encoder based approach to unsupervised learning of subword units," in *ICASSP*, 2014.
- [10] M. Versteegh, X. Anguera, A. Jansen, and E. Dupoux, "The zero resource speech challenge 2015: Proposed approaches and results," *Procedia Computer Science: SLTU 2016*, vol. 81, pp. 67–72, 2016.
- [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 *IEEE ASRU*, 2017.
- [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] E. Dunbar, J. Karadayi, M. Bernard, X.-N. Cao, R. Algayres, L. Ondel, L. Besacier, S. Sakriani, and E. Dupoux, "The zero resource speech challenge 2020: Discovering discrete subword and word units," in *INTERSPEECH*, 2020.
- [14] E. Dunbar, M. Bernard, N. Hamilakis, T. A. Nguyen, M. de Scyssel, P. Rozé, M. Rivière, E. Kharitonov, and E. Dupoux, "The zero resource speech challenge 2021: Spoken language modelling," in *INTERSPEECH*, 2021.
- [15] L. Badino, A. Mereta, and L. Rosasco, "Discovering discrete subword units with binarized autoencoders and hidden-markov-model encoders," in *INTERSPEECH*, 2015.
- [16] P. Baljekar, S. Sitaram, P. K. Muthukumar, and A. W. Black, "Using articulatory features and inferred phonological segments in zero resource speech processing," in *INTERSPEECH*, 2015.
- [17] 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.
- [18] A. F. Myrman and G. Salvi, "Partitioning of posteriorgrams using siamese models for unsupervised acoustic modelling," in *ISCA Workshop on Grounding Language Understanding*, 2017.
- [19] M. Heck, S. Sakti, and S. Nakamura, "Unsupervised linear discriminant analysis for supporting dpgmm clustering in the zero resource scenario," *Procedia Computer Science*, vol. 81, pp. 73–79, 2016.
- [20] D. Renshaw, H. Kamper, A. Jansen, and S. Goldwater, "A comparison of neural network methods for unsupervised representation learning on the zero resource speech challenge," in *INTERSPEECH*, 2015.
- [21] B. M. L. Srivastava and M. Shrivastava, "Articulatory gesture rich representation learning of phonological units in low resource settings," in *International Conference on Statistical Language and Speech Processing*, 2016, pp. 80–95.
- [22] R. Thiollere, E. Dunbar, G. Synnaeve, M. Versteegh, and E. Dupoux, "A hybrid dynamic time warping-deep neural network architecture for unsupervised acoustic modeling," in *INTERSPEECH*, 2015.
- [23] N. Zeghidour, G. Synnaeve, M. Versteegh, and E. Dupoux, "A deep scattering spectrum—deep siamese network pipeline for unsupervised acoustic modeling," in *ICASSP*, 2016, pp. 4965–4969.
- [24] T. Ansari, R. Kumar, S. Singh, and S. Ganapathy, "Deep learning methods for unsupervised acoustic modeling—leap submission to zrc 2017," in *IEEE ASRU Workshop*, 2017, pp. 754–761.
- [25] H. Chen, C.-C. Leung, L. Xie, B. Ma, and H. Li, "Multilingual bottleneck feature learning from untranscribed speech," in *IEEE ASRU Workshop*, 2017, pp. 727–733.
- [26] M. Heck, S. Sakti, and S. Nakamura, "Feature optimized dpgmm clustering for unsupervised subword modeling: A contribution to zerospeech 2017," in *IEEE ASRU Workshop*, 2017, pp. 740–746.
- [27] T. Pellegrini, C. Manenti, and J. Pinquier, "Technical report the irit-ups system@ zerospeech 2017 track1: unsupervised subword modeling," Tech. rep., IRIT, Université de Toulouse, Tech. Rep., 2017.
- [28] H. Shibata, T. Kato, T. Shinozaki, and S. Watanabe, "Composite embedding systems for zerospeech2017 track1," in *IEEE ASRU Workshop*, 2017, pp. 747–753.
- [29] 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 *IEEE ASRU Workshop*, 2017, pp. 734–739.
- [30] J. Chorowski, R. J. Weiss, S. Bengio, and A. v. d. Oord, "Unsupervised speech representation learning using wavenet autoencoders," *arXiv preprint arXiv:1901.08810*, 2019.
- [31] E. Kharitonov, M. Rivière, G. Synnaeve, L. Wolf, P.-E. Mazaré, M. Douze, and E. Dupoux, "Data augmenting contrastive learning of speech representations in the time domain," in *IEEE SLT Workshop*, 2021, pp. 215–222.
- [32] A. Jansen and B. Van Durme, "Efficient spoken term discovery using randomized algorithms," in *IEEE ASRU Workshop*, 2011, pp. 401–406.
- [33] H. Kamper, K. Livescu, and S. Goldwater, "An embedded segmental k-means model for unsupervised segmentation and clustering of speech," in *IEEE ASRU Workshop*, 2017, pp. 719–726.
- [34] V. Lyzinski, G. Sell, and A. Jansen, "An evaluation of graph clustering methods for unsupervised term discovery," in *INTERSPEECH*, 2015.
- [35] O. Räsänen, G. Doyle, and M. C. Frank, "Unsupervised word discovery from speech using automatic segmentation into syllable-like units," in *INTERSPEECH*, 2015.
- [36] H. Kamper, "Word segmentation on discovered phone units with dynamic programming and self-supervised scoring," *arXiv preprint arXiv:2202.11929*, 2022.
- [37] R. Algayres, T. Ricoul, J. Karadayi, A. Mohammed, B. Sagot, and E. Dupoux, "DP-PARSE: finding word boundaries from raw speech with a token lexicon," *TACL*, 2022.
- [38] H. Kamper, A. Jansen, and S. Goldwater, "A segmental framework for fully-unsupervised large-vocabulary speech recognition," *Computer Speech & Language*, vol. 46, pp. 154–174, 2017.
- [39] S. Bhati, J. Villalba, P. Želasko, and N. Dehak, "Self-expressing autoencoders for unsupervised spoken term discovery," *arXiv preprint arXiv:2007.13033*, 2020.
- [40] O. Räsänen and M. A. C. Blandón, "Unsupervised discovery of recurring speech patterns using probabilistic adaptive metrics," *arXiv preprint arXiv:2008.00731*, 2020.
- [41] S. Seshadri, U. Remes, O. Räsänen *et al.*, "Comparison of non-parametric bayesian mixture models for syllable clustering and zero-resource speech processing," *INTERSPEECH*, 2017.
- [42] Y. Iwamoto and T. Shinozaki, "Unsupervised spoken term discovery using wav2vec 2.0," in *IEEE APSIPA ASC*, 2021, pp. 1082–1086.
- [43] L. Ondel, L. Burget, and J. Černocký, "Variational inference for acoustic unit discovery," ser. *Procedia Computer Science (SLTU)*, vol. 81, 2016, pp. 80–86.
- [44] A. Tjandra, B. Sisman, M. Zhang, S. Sakti, H. Li, and S. Nakamura, "VQVAE unsupervised unit discovery and multi-scale code2spec inverter for ZRC 2019," *INTERSPEECH*, pp. 1118–1122, 2019.
- [45] D. Karthik Pandia and H. A. Murthy, "Zero resource speech synthesis using transcripts derived from perceptual acoustic units," *INTERSPEECH*, pp. 1113–1117, 2019.
- [46] R. Eloff, A. Nortje, B. van Niekerk, A. Govender, L. Nortje, A. Pretorius, E. Van Biljon, E. van der Westhuizen, L. van Staden, and H. Kamper, "Unsup. acoustic unit disc. for speech synthesis using discrete latent-variable neural networks," in *INTERSPEECH*, 2019.
- [47] S. Feng, T. Lee, and Z. Peng, "Combining Adversarial Training and Disentangled Speech Representation for Robust Zero-Resource Subword Modeling," *INTERSPEECH*, pp. 1093–1097, 2019.
- [48] A. T. Liu, P.-c. Hsu, and H.-Y. Lee, "Unsupervised end-to-end learning of discrete linguistic units for voice conversion," *INTERSPEECH*, pp. 1108–1112, 2019. [Online]. Available: <https://arxiv.org/abs/1905.11563>
- [49] S. Nayak, C. S. Kumar, G. Ramesh, S. Bhati, and K. S. R. Murty, "Virtual phone discovery for speech synthesis without text," in *IEEE GlobalSIP*, 2019, pp. 1–5.
- [50] B. Yusuf, A. Gök, B. Gündogdu, O. D. Kose, and M. Saraclar, "Temporally-aware acoustic unit discovery for zerospeech 2019 challenge," in *INTERSPEECH*, 2019, pp. 1098–1102.
- [51] B. Gündogdu, B. Yusuf, M. Yesilbursa, and M. Saraclar, "Vector quantized temporally-aware correspondence sparse autoencoders for zero-resource acoustic unit discovery," in *INTERSPEECH*, 2020, pp. 4846–4850.
- [52] B. Yusuf, L. Ondel, L. Burget, J. Černocký, and M. Saraclar, "A hierarchical subspace model for language-attuned acoustic unit discovery," in *ICASSP*, 2021, pp. 3710–3714.

- [53] T. Morita and H. Koda, "Exploring TTS without T using biologically/psychologically motivated neural network modules (zerospeech 2020)," *arXiv preprint arXiv:2005.05487*, 2020.
- [54] A. Prakash, M. R. Kumar, H. A. Murthy *et al.*, "Exploration of end-to-end synthesizers for zero resource speech challenge 2020," *arXiv preprint arXiv:2009.04983*, 2020.
- [55] M. Chen and T. Hain, "Unsupervised acoustic unit representation learning for voice conversion using wavenet auto-encoders," *arXiv preprint arXiv:2008.06892*, 2020.
- [56] P. L. Tobing, T. Hayashi, Y.-C. Wu, K. Kobayashi, and T. Toda, "Cyclic spectral modeling for unsupervised unit discovery into voice conversion with excitation and waveform modeling," in *INTERSPEECH*, 2020, pp. 4861–4865.
- [57] B. van Niekerk, L. Nortje, and H. Kamper, "Vector-quantized neural networks for acoustic unit discovery in the zerospeech 2020 challenge," *arXiv preprint arXiv:2005.09409*, 2020.
- [58] A. Tjandra, S. Sakti, and S. Nakamura, "Transformer vq-vae for unsupervised unit discovery and speech synthesis: Zerospeech 2020 challenge," *arXiv preprint arXiv:2005.11676*, 2020.
- [59] T. A. Nguyen, M. de Seyssel, P. Rozé, M. Rivière, E. Kharitonov, A. Baevski, E. Dunbar, and E. Dupoux, "The zero resource speech benchmark 2021: Metrics and baselines for unsupervised spoken language modeling," *arXiv preprint arXiv:2011.11588*, 2020.
- [60] A. Alishahi, G. Chrupala, A. Cristia, E. Dupoux, B. Higy, M. Lavechin, O. Räsänen, and C. Yu, "Zr-2021vg: Zero-resource speech challenge, visually-grounded language modelling track," *arXiv preprint arXiv:2107.06546*, 2021.
- [61] B. van Niekerk, L. Nortje, M. Baas, and H. Kamper, "Analyzing speaker information in self-supervised models to improve zero-resource speech processing," *arXiv preprint arXiv:2108.00917*, 2021.
- [62] J. Chorowski, G. Ciesielski, J. Dzikiowski, A. Łańcucki, R. Marxer, M. Opala, P. Pusz, P. Rychlikowski, and M. Stypułkowski, "Information retrieval for zerospeech 2021: The submission by university of wroclaw," *arXiv preprint arXiv:2106.11603*, 2021.
- [63] T. Maekaku, X. Chang, Y. Fujita, L.-W. Chen, S. Watanabe, and A. Rudnicky, "Speech representation learning combining conformer cpc with deep cluster for the zerospeech challenge 2021," *arXiv preprint arXiv:2107.05899*, 2021.
- [64] P. Peng and D. Harwath, "Self-supervised representation learning for speech using visual grounding and masked language modeling," *arXiv preprint arXiv:2202.03543*, 2022.
- [65] T. A. Nguyen, B. Sagot, and E. Dupoux, "Are discrete units necessary for spoken language modeling?" *IEEE J. Sel. Top. in Sig. Proc.*, 2022.
- [66] S. Bhati, J. Villalba, P. Żelasko, L. Moro-Velazquez, and N. Dehak, "Unsupervised speech segmentation and variable rate representation learning using segmental contrastive predictive coding," *arXiv preprint arXiv:2110.02345*, 2021.
- [67] K. Qian, Y. Zhang, H. Gao, J. Ni, C.-I. Lai, D. Cox, M. Hasegawa-johnson, and S. Chang, "Contentvec: An improved self-supervised speech representation by disentangling speakers," in *International Conference on Machine Learning*, 2022.
- [68] S. Cuervo, M. Grabias, J. Chorowski, G. Ciesielski, A. Łańcucki, P. Rychlikowski, and R. Marxer, "Contrastive prediction strategies for unsupervised segmentation and categorization of phonemes and words," *arXiv preprint arXiv:2110.15909*, 2021.
- [69] E. Gauthier, L. Besacier, S. Voisin, M. Melese, and U. P. Elingui, "Collecting Resources in Sub-Saharan African Languages for Automatic Speech Recognition: a Case Study of Wolof," in *LREC*, 2016.
- [70] S. Sakti, E. Kelana, H. Riza, S. Sakai, K. Markov, and S. Nakamura, "Development of Indonesian large vocabulary continuous speech recognition system within A-STAR project," in *Proceedings of the TCAST Workshop*, 2008.
- [71] S. Sakti, R. Maia, S. Sakai, T. Shimizu, and S. Nakamura, "Development of HMM-based Indonesian speech synthesis," in *Proc. Oriental COCOSA*, 2008, pp. 215–219.
- [72] T. Schatz, V. Peddinti, F. Bach, A. Jansen, H. Hermansky, and E. Dupoux, "Evaluating speech features with the Minimal-Pair ABX task (I): Analysis of the classical MFC/PLP pipeline," in *INTERSPEECH*, 2013.
- [73] T. Schatz, V. Peddinti, X.-N. Cao, F. Bach, H. Hermansky, and E. Dupoux, "Evaluating speech features with the Minimal-Pair ABX task (II): Resistance to noise," in *INTERSPEECH*, 2014.
- [74] A. van den Oord, Y. Li, and O. Vinyals, "Representation learning with contrastive predictive coding," *CoRR*, vol. abs/1807.03748, 2018.
- [75] A. Baevski, H. Zhou, A. Mohamed, and M. Auli, "wav2vec 2.0: A framework for self-supervised learning of speech representations," *arXiv preprint arXiv:2006.11477*, 2020.
- [76] W.-N. Hsu, B. Bolte, Y.-H. H. Tsai, K. Lakhotia, R. Salakhutdinov, and A. Mohamed, "HuBERT: Self-supervised speech representation learning by masked prediction of hidden units," *IEEE Trans. on Audio, Speech, and Lang. Proc.*, vol. 29, pp. 3451–3460, 2021.
- [77] L. Borgholt, J. D. Havtorn, J. Edin, L. Maaløe, and C. Igel, "A brief overview of unsupervised neural speech representation learning," 2022.
- [78] N. H. Feldman, T. L. Griffiths, S. Goldwater, and J. L. Morgan, "A role for the developing lexicon in phonetic category acquisition," *Psychological review*, vol. 120, no. 4, pp. 751–778, 2013.
- [79] M. Elsner, S. Goldwater, and J. Eisenstein, "Bootstrapping a unified model of lexical and phonetic acquisition," in *Proceedings of ACL*, 2012, pp. 184–193.
- [80] B. Dillon, E. Dunbar, and W. Idsardi, "A single-stage approach to learning phonological categories: Insights from Inuktitut," *Cognitive Science*, vol. 37, no. 2, pp. 344–377, 2013.
- [81] J. Millet, I. Chitoran, and E. Dunbar, "Predicting non-native speech perception using the perceptual assimilation model and state-of-the-art acoustic models," in *Proceedings of CoNLL*, 2021, pp. 661–673.
- [82] J. Millet and E. Dunbar, "Do self-supervised speech models develop human-like perception biases?" in *Proceedings of ACL 2022*, 2022.
- [83] B. C. Moore, *An introduction to the psychology of hearing*. Brill, 2012.
- [84] L. Weerts, S. Rosen, C. Clopath, and D. F. Goodman, "The psychometrics of automatic speech recognition," *bioRxiv*, 2021.
- [85] H.-S. Tsai, H.-J. Chang, W.-C. Huang, Z. Huang, K. Lakhotia, S.-w. Yang, S. Dong, A. Liu, C.-I. Lai, J. Shi *et al.*, "SUPERB-SG: Enhanced speech processing universal performance benchmark for semantic and generative capabilities," in *Proceedings of ACL*, 2022, pp. 8479–8492.
- [86] P.-Y. Huang, H. Xu, J. Li, A. Baevski, M. Auli, W. Galuba, F. Metzke, and C. Feichtenhofer, "Masked autoencoders that listen," 2022. [Online]. Available: <https://arxiv.org/abs/2207.06405>
- [87] A. Baade, P. Peng, and D. Harwath, "MAE-AST: Masked autoencoding audio spectrogram transformer," in *arxiv:2203.16691*.
- [88] A. S. Park and J. R. Glass, "Unsupervised Pattern Discovery in Speech," *IEEE Transactions on Audio, Speech, and Language Processing*, vol. 16, no. 1, pp. 186–197, 2008.
- [89] Y. Zhang and J. R. Glass, "Towards multi-speaker unsupervised speech pattern discovery," in *ICASSP*, 2010, pp. 4366–4369.
- [90] A. Jansen and B. Van Durme, "Efficient spoken term discovery using randomized algorithms," in *IEEE ASRU*, 2011, pp. 401–406.
- [91] A. Muscariello, G. Gravier, and F. Bimbot, "Unsupervised Motif Acquisition in Speech via Seeded Discovery and Template Matching Combination," *IEEE Transactions on Audio, Speech and Language Processing*, vol. 20, no. 7, pp. 2031–2044, 2012.
- [92] A. Jansen, S. Thomas, and H. Hermansky, "Weak top-down constraints for unsupervised acoustic model training," in *ICASSP*, 2013, pp. 8091–8095.
- [93] S. Goldwater, T. L. Griffiths, and M. Johnson, "A bayesian framework for word segmentation: Exploring the effects of context," *Cognition*, vol. 112, no. 1, pp. 21–54, 2009.
- [94] R. Algayres, M. S. Zaiem, B. Sagot, and E. Dupoux, "Evaluating the reliability of acoustic speech embeddings," in *INTERPEECH*, 2020.
- [95] M. Johnson, T. Griffiths, and S. Goldwater, "Adaptor grammars: A framework for specifying compositional nonparametric bayesian models," in *NIPS*, 2007, vol. 19, pp. 641–648.
- [96] S. Goldwater, T. Griffiths, and M. Johnson, "A bayesian framework for word segmentation: Exploring the effects of context," *Cognition*, vol. 112, p. 21–54, 2009.
- [97] S. Bhati, J. Villalba, P. Żelasko, L. Moro-Velazquez, and N. Dehak, "Segmental contrastive predictive coding for unsupervised word segmentation," *arXiv preprint arXiv:2106.02170*, 2021.
- [98] O. Scharenborg, L. Besacier, A. W. Black, M. Hasegawa-Johnson, F. Metzke, G. Neubig, S. Stüker, P. Godard, M. Müller, L. Ondel, S. Palaskar, P. Arthur, F. Ciannella, M. Du, E. Larsen, D. Merx, R. Riad, L. Wang, and E. Dupoux, "Linguistic unit discovery from multi-modal inputs in unwritten languages: Summary of the "speaking rosetta" JSALT 2017 workshop," in *ICASSP*, 2018, pp. 4979–4983.
- [99] Z. Wu, O. Watts, and S. King, "Merlin: An open source neural network speech synthesis system," in *Speech Synthesis Workshop*. ISCA, 2016, pp. 202–207.
- [100] K. Lakhotia, E. Kharitonov, W.-N. Hsu, Y. Adi, A. Polyak, B. Bolte, T.-A. Nguyen, J. Copet, A. Baevski, A. Mohamed, and E. Dupoux, "On generative spoken language modeling from raw audio," *TACL*, vol. 9, pp. 1336–1354, 2021.
- [101] J. Kahn, M. Riviere, W. Zheng, E. Kharitonov, Q. Xu, P. Mazare, J. Karadayi, V. Liptchinsky, R. Collobert, C. Fuegen, T. Likhomanenko,

- G. Synnaeve, A. Mohamed, and E. Dupoux, "Libri-light: A benchmark for asr with limited or no supervision," *ICASSP*, pp. 7669–7673, 2020.
- [102] M. Rivière, A. Joulin, P.-E. Mazaré, and E. Dupoux, "Unsupervised pretraining transfers well across languages," in *ICASSP*, 2020, pp. 7414–7418.
- [103] E. Keuleers and M. Brysbaert, "Wuggy: A multilingual pseudoword generator," *Behavior research methods*, vol. 42, no. 3, pp. 627–633, 2010.
- [104] A. Warstadt, A. Parrish, H. Liu, A. Mohanany, W. Peng, S.-F. Wang, and S. R. Bowman, "Blimp: A benchmark of linguistic minimal pairs for english," *TACL*, pp. 377–392, 2019.
- [105] Y.-A. Chung and J. Glass, "Speech2vec: A sequence-to-sequence framework for learning word embeddings from speech," in *INTER-SPEECH*, 2018, pp. 811–815.
- [106] D. Yang and D. M. Powers, *Verb similarity on the taxonomy of WordNet*. Masaryk University, 2006.
- [107] E. Agirre, E. Alfonseca, K. Hall, J. Kravalova, M. Pasca, and A. Soroa, "A study on similarity and relatedness using distributional and wordnet-based approaches," in *Proc. of NAACL*, 2009, pp. 19–27.
- [108] G. A. Miller and W. G. Charles, "Contextual correlates of semantic similarity," *Language and cognitive processes*, vol. 6, no. 1, pp. 1–28, 1991.
- [109] H. Rubenstein and J. B. Goodenough, "Contextual correlates of synonymy," *Comm. of the ACM*, vol. 8, no. 10, pp. 627–633, 1965.
- [110] M.-T. Luong, R. Socher, and C. D. Manning, "Better word representations with recursive neural networks for morphology," in *Proceedings of CoNLL*, 2013, pp. 104–113.
- [111] F. Hill, R. Reichart, and A. Korhonen, "Simlex-999: Evaluating semantic models with (genuine) similarity estimation," *Computational Linguistics*, vol. 41, no. 4, pp. 665–695, 2015.
- [112] D. Gerz, I. Vulić, F. Hill, R. Reichart, and A. Korhonen, "Simverb-3500: A large-scale evaluation set of verb similarity," *arXiv preprint arXiv:1608.00869*, 2016.
- [113] S. Baker, R. Reichart, and A. Korhonen, "An unsupervised model for instance level subcategorization acquisition," in *Proceedings of EMNLP*, 2014, pp. 278–289.
- [114] E. Bruni, G. Boleda, M. Baroni, and N.-K. Tran, "Distributional semantics in technicolor," in *Proceedings of ACL*, 2012, pp. 136–145.
- [115] K. Radinsky, E. Agichtein, E. Gabrilovich, and S. Markovitch, "A word at a time: computing word relatedness using temporal semantic analysis," in *Proceedings of ACM WWW*, 2011, pp. 337–346.
- [116] G. Halawi, G. Dror, E. Gabrilovich, and Y. Koren, "Large-scale learning of word relatedness with constraints," in *Proceedings of ACM SIGKDD*, 2012, pp. 1406–1414.
- [117] V. Panayotov, G. Chen, D. Povey, and S. Khudanpur, "LibriSpeech: an ASR corpus based on public domain audio books," in *ICASSP*, 2015, pp. 5206–5210.
- [118] D. F. Harwath, A. Torralba, and J. R. Glass, "Unsupervised learning of spoken language with visual context," in *NIPS*, 2016, pp. 1858–1866.
- [119] R. Arandjelovic and A. Zisserman, "Look, listen and learn," in *Proceedings of IEEE ICCV*, 2017, pp. 609–617.
- [120] G. Chrupała, L. Gelderloos, and A. Alishahi, "Representations of language in a model of visually grounded speech signal," in *Proceedings of ACL*, 2017, pp. 613–622.
- [121] L. Gwilliams, T. Linzen, D. Poeppel, and A. Marantz, "In spoken word recognition, the future predicts the past," *Journal of Neuroscience*, vol. 38, no. 35, pp. 7585–7599, 2018.
- [122] B. Beekhuizen, B. C. Armstrong, and S. Stevenson, "Probing lexical ambiguity: Word vectors encode number and relatedness of senses," *Cognitive Science*, vol. 45, no. 5, p. e12943, 2021.
- [123] F. De Saussure, *Course in General Linguistics*. McGraw-Hill Book Company, New York-Toronto-London, 1916.
- [124] R. Van Gijn and F. Zúñiga, "Word and the americanist perspective," *Morphology*, vol. 24, no. 3, pp. 135–160, 2014.
- [125] S. Chen, C. Wang, Z. Chen, Y. Wu, S. Liu, Z. Chen, J. Li, N. Kanda, T. Yoshioka, X. Xiao *et al.*, "Wavlm: Large-scale self-supervised pre-training for full stack speech processing," *IEEE Journal of Selected Topics in Signal Processing*, 2022.
- [126] E. Kharitonov, A. Lee, A. Polyak, Y. Adi, J. Copet, K. Lakhotia, T.-A. Nguyen, M. Rivière, A. Mohamed, E. Dupoux, and W. Hsu, "Text-free prosody-aware generative spoken language modeling," in *ACL*, 2021, pp. 8666–8681.
- [127] T. A. Nguyen, E. Kharitonov, J. Copet, Y. Adi, W.-N. Hsu, A. Elkhakhy, P. Tomasello, R. Algayres, B. Sagot, A. Mohamed, and E. Dupoux, "Generative spoken dialogue language modeling," 2022. [Online]. Available: <https://arxiv.org/abs/2203.16502>
- [128] M. de Seyssel, M. Lavechin, Y. Adi, E. Dupoux, and G. Wisniewski, "Probing phoneme, language and speaker information in unsupervised speech representations," in *INTER-SPEECH*, 2022, pp. arXiv–2203.
- [129] J. Millet and E. Dunbar, "Perceptimatic: A human speech perception benchmark for unsupervised subword modelling," *arXiv preprint arXiv:2010.05961*, 2020.
- [130] M. Rivière and E. Dupoux, "Towards unsupervised learning of speech features in the wild," in *IEEE SLT Workshop*, 2021, pp. 156–163.
- [131] A. Tjandra, S. Sakti, and S. Nakamura, "Speech-to-speech translation between untranscribed unknown languages," in *IEEE ASRU Workshop*, 2019, pp. 593–600.
- [132] Y. Jia, R. J. Weiss, F. Biadys, W. Macherey, M. Johnson, Z. Chen, and Y. Wu, "Direct speech-to-speech translation with a sequence-to-sequence model," *arXiv preprint arXiv:1904.06037*, 2019.
- [133] A. Lee, P.-J. Chen, C. Wang, J. Gu, X. Ma, A. Polyak, Y. Adi, Q. He, Y. Tang, J. Pino *et al.*, "Direct speech-to-speech translation with discrete units," *arXiv preprint arXiv:2107.05604*, 2021.
- [134] Y. Jia, M. T. Ramanovich, T. Remez, and R. Pomerantz, "Translation 2: Robust direct speech-to-speech translation," *arXiv preprint arXiv:2107.08661*, 2021.
- [135] A. Lee, H. Gong, P.-A. Duquenne, H. Schwenk, P.-J. Chen, C. Wang, S. Popuri, J. Gu, and W.-N. Hsu, "Textless speech-to-speech translation on real data," *arXiv preprint arXiv:2112.08352*, 2021.



Ewan Dunbar is an Assistant Professor in the French Linguistics program at the University of Toronto and an Affiliated Scientist with the Cognitive Machine Learning (CoML) team at the Ecole Normale Supérieure (ENS) in Paris and INRIA. He holds a PhD in Linguistics from the University of Maryland, College Park, an MA in Linguistics from the University of Toronto, and a BSc in Linguistics and Computer Science, also from the University of Toronto. His research focuses on speech perception, psycholinguistics, and speech technology. He is co-organizer of the Zero Resource Speech Challenge (2017–2021).



Nicolas Hamilakis is a Research Engineer for the Cognitive Machine Learning (CoML) team at the Ecole Normale Supérieure (ENS) in Paris and INRIA. He obtained an undergraduate degree in mathematics and computer science at the Pierre & Marie Curie University and a Masters in the engineering of programs and algorithms at Université de Paris. His work is a mix of data management and software engineering of experimental research tools. His interests are in software engineering and in how to provide better, more elegant solutions and reproducible results. He worked on the organizational committee of the 2021 edition of the challenge including the construction of a pipeline for evaluating submissions.



Emmanuel Dupoux is a Professor at the École des Hautes Études en Sciences Sociales (EHESS) and Research Scientist at Meta AI Labs. He directs the Cognitive Machine Learning (CoML) team at the Ecole Normale Supérieure (ENS) in Paris and INRIA. His education includes a doctorate in Cognitive Science (EHESS), a Masters in Computer Science (Orsay University) and an undergraduate degree in Applied Mathematics (Pierre & Marie Curie University). He is co-organizer of the Zero Resource Speech Challenge series (2015–2021). He is a CIFAR LMB and a ELLIS Fellow. He has authored 150 articles in peer reviewed outlets in cognitive science and language technology.