



# Bayesian Subspace Hidden Markov Model for Acoustic Unit Discovery

Lucas Ondel, Hari Krishna Vydana, Lukáš Burget, Jan Černocký

Brno University of Technology

{iondel, vydana, burget, cernocky}@fit.vutbr.cz

## Abstract

This work tackles the problem of learning a set of language specific acoustic units from unlabeled speech recordings given a set of labeled recordings from other languages. Our approach may be described by the following two steps procedure: first the model learns the notion of acoustic units from the labelled data and then the model uses its knowledge to find new acoustic units on the target language. We implement this process with the Bayesian Subspace Hidden Markov Model (SHMM), a model akin to the Subspace Gaussian Mixture Model (SGMM) where each low dimensional embedding represents an acoustic unit rather than just a HMM's state. The subspace is trained on 3 languages from the GlobalPhone corpus (German, Polish and Spanish) and the AUs are discovered on the TIMIT corpus. Results, measured in equivalent Phone Error Rate, show that this approach significantly outperforms previous HMM based acoustic units discovery systems and compares favorably with the Variational Auto Encoder-HMM.

**Index Terms:** Bayesian Inference, Hidden Markov Model, Subspace Model, Variational Bayes, Low-resource languages, Acoustic Unit Discovery

## 1. Introduction

State-of-the-art Automatic Speech Recognition (ASR) systems rely upon very large amount of speech recordings paired with textual transcriptions. While this approach has proven to be very successful, it is however limited to the very few languages having enough resources to train an ASR system. Due to the cost of data collection and transcription, broadening the range of speech technologies to any language remains an unreachable objective. Parallel to the mainstream ASR, there has been a growing interest in the paradigm of unsupervised learning of speech [1]. Unsupervised speech learning attempts to use machine learning techniques to extract various information (phonetic content, speaker identity, ...) from unlabeled recordings. While this is considerably harder than standard ASR, solving this problem would have a considerable impact on the field by reducing the amount of human labour necessary to build a full fledged ASR pipeline. It is also important to emphasize that the linguistic diversity is diminishing worldwide. Many languages are now considered endangered and risk to disappear in a near future. Affordable speech technologies could be a precious tool to help linguists and communities to document and preserve these languages.

This work focuses on the specific task of acoustic unit discovery (AUD). Given a collection of unlabeled recordings in a specific language, the task is to learn a set of basic speech units (also called pseudo-phones) to describe the language. AUD algorithms have to solve three problems: to segment the speech, to cluster the segments into units and to infer how many units are necessary to describe the language. Several approaches have been proposed relying upon Bayesian non-parametric version of

the Hidden Markov Model (HMM) [2, 3, 4]. An important recent extension of this model is the VAE-HMM [5, 6, 7] which combines the traditional HMM with Variational Auto Encoder [8]. However, most of the AUD algorithms are prone to model speaker/channel or any non-phonetic variability. To address this issue, we propose the Bayesian Subspace HMM (SHMM). The SHMM is an HMM based AUD model in which the parameters of each unit is constrained to be in the phonetic subspace of the total parameter space. This restriction forces the AUD model to focus on the phonetic content of the speech signal and to ignore irrelevant information.

## 2. Model

### 2.1. Standard Acoustic Unit Discovery

Let  $\mathbf{X} = (\mathbf{x}_1, \dots, \mathbf{x}_N)$  be the sequence of  $N$  observed speech frames and  $U = \{u_1, \dots, u_P\}$  be the set of  $P$  acoustic units.  $\mathbf{v} = (v_1, \dots, v_N)$ ,  $v_i \in U$  is a sequence of variables indicating to which unit each speech frame is associated, and  $\mathbf{Z} = (\mathbf{z}_1, \dots, \mathbf{z}_N)$  are model-dependent latent variables. We consider generative models for which the complete likelihood of the data factorizes as:

$$p(\mathbf{X}, \mathbf{Z} | \mathbf{v}) = \prod_{n=1}^N p(\mathbf{x}_n, \mathbf{z}_n | v_n) \quad (1)$$

and the likelihood of a speech frame for a given unit is member of the exponential family of distribution:

$$p(\mathbf{x}_n, \mathbf{z}_n | v_n = u) = \exp \{ \boldsymbol{\eta}_u^T T(\mathbf{x}_n, \mathbf{z}_n) - A(\boldsymbol{\eta}_u) \} \quad (2)$$

where  $\boldsymbol{\eta}_u \in \mathcal{H}$  is the  $D$ -dimensional vector of natural parameters corresponding to one acoustic unit,  $T(\mathbf{x}_n, \mathbf{z}_n)$  are the sufficient statistics and  $A(\boldsymbol{\eta}_u)$  is the (log-)normalization constant of the density. Note that the nature of the model for the units (HMM, GMM, Linear Dynamical Model, ...) will depend on the value of  $\mathbf{z}_n$  and the sufficient statistics  $T$ . In this work we consider that each unit is modeled by an HMM with a GMM for each state's emission but it can be replaced by any model satisfying Eq. 1 and Eq. 2. Previous works [2, 3, 6] use special cases of this model to perform the AUD. More precisely, one can understand AUD as finding a set of vectors  $\boldsymbol{\eta}_{u_1}, \dots, \boldsymbol{\eta}_{u_p}$  such that the likelihood of the observation is maximized<sup>1</sup>. This search is difficult because speech recordings encode many factors other than the phonetic information (speaker identity, emotions, environment, ...) and the AUD algorithm may maximize the likelihood while modeling non-phonetic information.

### 2.2. Subspace HMM

To avoid the AUD model to capture non-phonetic information, we proposed the Subspace HMM (SHMM) which constrains

<sup>1</sup>These algorithm also learn the number of acoustic units  $P$  needed to fit the data




Figure 1: (a) Directed Acyclic Graph of a Generalized Subspace Model. Dashed lines represent deterministic relationship between variables. SHMM, JFA, SGMM are special cases of this model. In this work each embedding  $\mathbf{h}_u$  encodes the parameters of one HMM corresponding to an acoustic unit. (b) Illustration of the subspace model for acoustic units. Each point of the plane corresponds to the parameters of an acoustic unit model and the blue line represents the subspace defined by  $f(\mathbf{W}^T \mathbf{h} + \mathbf{b})$ . Given an acoustic unit model corresponding to the sound aa, moving its parameters along the subspace will change the model to represent another unit/phone (ow, z in this example). Conversely, moving the parameters away from the phonetic subspace will push the model to capture non-phonetic information (for instance speaker gender).

the parameters of the acoustic units to live in the phonetic space. This model extends the unsupervised HMM by assuming that the phonetic information of a language is contained in a subspace of the total parameters space. Formally, it is defined as:

$$\boldsymbol{\eta}_u = f(\mathbf{W}^T \mathbf{h}_u + \mathbf{b}) \quad (3)$$

where  $f : \mathbb{R}^D \mapsto \mathcal{H}$  is a differentiable function. We further refine this subspace model by introducing a prior over the subspace's parameters:

$$W_{r,c} \sim \mathcal{N}(0, \sigma_{W_{r,c}}^2) \quad (4)$$

$$\mathbf{b} \sim \mathcal{N}(\mathbf{0}, \mathbf{I}) \quad (5)$$

$$\mathbf{h}_u \sim \mathcal{N}(\mathbf{0}, \mathbf{I}) \quad (6)$$

As depicted in Fig. 1b, the bases of  $\mathbf{W}$  span the subspace containing the phonetic variability. Since the parameters of the acoustic units are constrained to live in a low dimensional subspace, the AUD algorithm can be seen as finding the set of embeddings  $\mathbf{h}_{u_1}, \dots, \mathbf{h}_{u_P}$  which maximizes the likelihood of the observations. By constraining the search in the phonetic subspace, we therefore force the algorithm to ignore non-phonetic source of variability.

Note that Subspace Gaussian Mixture Model [9], Joint Factor Analysis [10], Subspace Multinomial Model [11], etc. are special cases of Eq. 3. In fact, Eq. 3 is the general form of any subspace model for which the complete likelihood is a member of the exponential family of distributions. We denote Eq. 3 as the *Generalized Subspace Model* (GSM) of which the Subspace HMM, like other aforementioned models, is just a special instance. The graphical representation of the (GSM) is depicted in Fig. 1a. To complete our definition of the SHMM, we need to specify the mapping  $f$  from  $\mathbb{R}^D$  to the natural parameters space  $\mathcal{H}$ . In our setting, each unit is modeled by a HMM with a 3 states left-to-right topology and each state has a GMM emissions with  $K$  Gaussian components with diagonal covariance matrix. For convenience, we introduce the vector  $\boldsymbol{\psi} = \mathbf{W}^T \mathbf{h} + \mathbf{b}$  which can be decomposed into three parts  $\boldsymbol{\psi} = (\boldsymbol{\psi}_1, \boldsymbol{\psi}_2, \boldsymbol{\psi}_3)^T$ .  $\boldsymbol{\psi}_i$  is the vector of parameters (before the

mapping  $f$ ) associated with the  $i$ th HMM state.  $\boldsymbol{\psi}_i$  further decomposes into  $\boldsymbol{\psi}_i = (\boldsymbol{\psi}_i^\pi, \boldsymbol{\psi}_{i,1}^\mu, \dots, \boldsymbol{\psi}_{i,K}^\mu, \boldsymbol{\psi}_{i,1}^\Sigma, \dots, \boldsymbol{\psi}_{i,K}^\Sigma)^T$  where  $\boldsymbol{\psi}_i^\pi$  is the vector encoding the parameters of the mixture's weights and  $\boldsymbol{\psi}_{i,j}^\mu$  and  $\boldsymbol{\psi}_{i,j}^\Sigma$  are the vectors encoding the parameters of the mean and covariance matrix of the  $j$ th Gaussian component, respectively. We set  $f$  such that:

$$\pi_{i,j} = \frac{\exp\{\boldsymbol{\psi}_{i,j}^\pi\}}{1 + \sum_{k=1}^{K-1} \exp\{\boldsymbol{\psi}_{i,k}^\pi\}} \quad (7)$$

$$\boldsymbol{\mu}_{i,j} = \boldsymbol{\psi}_{i,j}^\mu \quad (8)$$

$$\boldsymbol{\Sigma}_{i,j} = \text{diag}(\exp\{\boldsymbol{\psi}_{i,j}^\Sigma\}) \quad (9)$$

where  $\exp$  is the elementwise exponential function. One could also include the transition probabilities of the HMM but we kept them as fixed parameters in this work.

### 2.3. Estimating the phonetic subspace

Unlike previous AUD algorithms, our model requires to specify the phonetic subspace (parameterized by  $\mathbf{W}$  and  $\mathbf{b}$ ) before searching the acoustic units. This is a "chicken or egg" problem since we need the phonetic subspace to find the pseudo-phones of the language and we need to know the phones of a language to estimate the subspace. However, this problem can be alleviated by observing that many languages in the world have common phones. It is reasonable to believe that the phonetic subspace of language is well approximated by a phonetic subspace estimated from one or several other languages for which we have labeled data. Interestingly, this rationale naturally fits the Bayesian approach of the problem of AUD. Given unlabeled set of observation  $\mathbf{X}^{(t)}$  in a target language  $t$ , previous Bayesian AUD algorithms try to estimate the inventory of (pseudo-)phones  $U^{(t)}$  of the target language by estimating:

$$p(U^{(t)} | \mathbf{X}^{(t)}) = \frac{p(\mathbf{X}^{(t)} | U^{(t)}) p(U^{(t)})}{p(\mathbf{X}^{(t)})} \quad (10)$$

If we now assume the phonetic subspace to be estimated from the observations  $\mathbf{X}^{(p)}$  of another language  $p$  with known inven-

tory of phones  $U^{(p)}$ , the problem can be reformulated as:

$$p(U^{(t)}|\mathbf{X}^{(t)}, L^{(p)}, S) = \frac{p(\mathbf{X}^{(t)}|U^{(t)}, L^{(p)}, S)p(U^{(t)}|L^{(p)}, S)}{p(\mathbf{X}^{(t)}|L^{(p)}, S)} \quad (11)$$

$$L^{(p)} = \{\mathbf{X}^{(p)}, U^{(p)}\} \quad (12)$$

$$S = \{\mathbf{W}, \mathbf{b}\} \quad (13)$$

The term  $p(U^{(t)}|L^{(p)}, S)$  may be seen as some educated/informative prior which embeds the notion of phone into the AUD algorithm. This educated prior needs to be estimated as well which leads to a two steps procedure for the SHMM AUD algorithm. First, given the labeled data of one or several languages, the prior over the acoustic units is estimated. Informally speaking, we force the model to learn "what is a phone". Second the unlabeled data of the target language is clustered into pseudo-phones given the phonetic knowledge acquired by the model during the first step.

#### 2.4. Training

The two steps of the training (learning the prior and clustering the units) are carried out by optimizing the same objective function except that when estimating the prior, the acoustic unit transcription of each utterance is known. The presence or absence of the transcription will be reflected in  $p(\mathbf{v})$ . When there is no transcription,  $p(\mathbf{v})$  can be understood as a "pseudo-phone" loop (see [3] for details) and when the transcription is known then  $p(\mathbf{v})$  is just the inference graph used for forced alignment in a traditional HMM based ASR system.

Since the estimation of the exact posterior of the model's parameters is intractable, we use the Variational Bayes (VB) objective function to find an approximate posterior:

$$\mathcal{L}[q] = \langle \ln p(\mathbf{X}|\Xi, \Theta) \rangle_q - \text{D}_{\text{KL}}(q(\Xi, \Theta)||p(\Xi, \Theta)) \quad (14)$$

$$\Xi = \{\mathbf{Z}, \mathbf{v}\} \quad (15)$$

$$\Theta = \{\mathbf{W}, \mathbf{b}, \mathbf{h}_{u_1}, \dots, \mathbf{h}_{u_p}\} \quad (16)$$

where  $\langle \dots \rangle_q$  denote the expectation w.r.t. the distribution  $q$  and  $\text{D}_{\text{KL}}$  is the Kullback-Leibler divergence. Eq. 14 is not tractable for arbitrary distribution  $q$  we therefore consider the restricted set of distributions with the following *mean-field* factorization and the given parameterization:

$$q(\Xi, \Theta) = q(\Xi; \phi)q(\Theta; \zeta) \quad (17)$$

$$\zeta = \{\mathbf{m}, \lambda\} \quad (18)$$

$$q(\Theta; \zeta) = \mathcal{N}(\mathbf{m}, \text{diag}(\exp\{\lambda\})) \quad (19)$$

The parameters  $\phi$  of the variational posterior over  $\Xi$  will depend on the type of the model of the acoustic unit. For the case of an HMM, this is the probability to be in a particular state given the sequence of observations. Under these restrictions the optimization reduces to:

$$\phi^*, \zeta^* = \arg \max_{\phi, \zeta} \mathcal{L}(\phi, \zeta) \quad (20)$$

Since we assume each unit to be modeled by an HMM,  $\phi^*$  has an analytical solution which can be efficiently calculated using the *forward-backward* algorithm [12].  $\zeta^*$  has no analytical solution but can be found through a stochastic gradient ascent scheme. Noting that  $\nabla_{\phi} \mathcal{L}_{\zeta}(\phi^*) = \mathbf{0}$  we have:

$$\nabla_{\zeta} \mathcal{L}(\phi^*, \zeta) = \nabla_{\zeta} \mathcal{L}_{\phi^*}(\zeta) + \nabla_{\zeta} \phi^* \nabla_{\phi} \mathcal{L}_{\zeta}(\phi^*) \quad (21)$$

$$= \nabla_{\zeta} \mathcal{L}_{\phi^*}(\zeta) \quad (22)$$

Finally, we approximate  $\nabla_{\zeta} \mathcal{L}(\phi^*, \zeta) \approx \nabla_{\zeta} \mathcal{L}'(\phi^*, \zeta)$  by using the so called *re-parameterization* trick introduced in [8]:

$$\epsilon_l \sim \mathcal{N}(\mathbf{0}, \mathbf{I}) \quad (23)$$

$$\Theta_l = \mathbf{m} + \text{diag}(\exp\{\frac{\lambda}{2}\})\epsilon_l \quad (24)$$

$$\mathcal{L}(\phi, \zeta) \approx \frac{1}{L} \sum_{l=1}^L \ln p(\mathbf{X}|\Xi, \Theta_l) \quad (25)$$

$$- \text{D}_{\text{KL}}(q(\Xi, \Theta)||p(\Xi, \Theta)) = \mathcal{L}'(\phi, \zeta) \quad (26)$$

In practice we use the ADAM optimizer [13] to update  $\zeta$  and we use  $L = 10$  samples to compute the empirical expectation. The parameters  $\phi$  are re-estimated every 1000 updates of  $\zeta$ .

## 3. Experiments

### 3.1. Data, Features and Metrics

We conducted our experiments with the TIMIT [14] database and 3 languages from the GlobalPhone corpus [15]: German (GE), Polish (PO) and Spanish (SP). For each of the three GlobalPhone languages, we kept only 3000 randomly selected utterances. We used two sets of features: (i) the MFCC features concatenated with their first and second derivatives (ii) the Multi-Lingual bottleNeck (MBN) features trained on 17 Babel's languages [16]. The set of languages used to train the MBN features does not include English, German, Polish or Spanish. Both set of features were extracted at a rate of 100 Hz. For the case the MBN features, the audio signal was down-sampled to 8kHz.

We evaluated the different AUD algorithms in terms of phonetic segmentation and equivalent Phone Error Rate (eq. PER) ([17, 5]). For the phonetic segmentation we used the standard Recall, Precision and F-score measured against the timing provided in the TIMIT database with the 61 original phones. We tolerated boundary shifted by +- 2 frames (20 milliseconds). To compute the eq. PER, we mapped each acoustic unit to one of the 61 phones it overlaps the most with. Then, we reduced the reference transcription and proposed transcription to the 39 phone set [18] and computed the PER.

### 3.2. Estimating the phonetic subspace

First, we ran a controlled experiment to assess whether the SHMM is able to properly learn the phonetic subspace of a language. In this experiment, we used the MBN features and each HMM state had 8 Gaussian components. First, we trained a Bayesian HMM phone recognizer on the 48 phone set with a flat phonotactic language model on the traditional TIMIT training set (no SA\* utterances) and decoded on the test set mapping the phones to the 39 phone set. This phone recognizer achieved 36.4 % Phone Error Rate (PER). This number is very high since we have removed crucial elements of the traditional ASR pipeline (language model, context-dependent phones, ...) in order to evaluate the quality the acoustic model. For comparison, we trained a monophone system with a flat phonotactic language model using the Kaldi toolkit [9] which yielded 37.3 % PER. We then trained an SHMM based phone recognizer with varying subspace dimension using the same training and testing setup as the baseline HMM. We used the baseline model to provide the first estimate of  $\phi$  which we modified so that all the Gaussian components within a state have equal responsibility. We pre-trained the subspace for 15000 updates before updating  $\phi$  then we re-estimated  $\phi$  after every 1000 updates of  $\zeta$  for 30

Table 1: Comparison of the SHMM against other AUD models in terms of phonetic segmentation (Recall, Precision, F-score) and equivalent Phone Error Rate (%).

Model	Features	Prior Language	Recall	Precision	F-score	eq. PER
HMM [5]	MFCC + $\Delta$ + $\Delta\Delta$	None	-	-	-	65.4
VAE-HMM [5]	MFCC + $\Delta$ + $\Delta\Delta$	None	-	-	-	58.9
VAE-BHMM [6]	log-mel FBANK + $\Delta$ + $\Delta\Delta$	None	-	-	-	56.57
HMM	MFCC + $\Delta$ + $\Delta\Delta$	None	66.47	57.81	61.84	64.92
HMM	MBN	None	63.98	54.21	58.69	68.25
SHMM	MFCC + $\Delta$ + $\Delta\Delta$	GE	<b>75.74</b>	<b>78.98</b>	<b>77.32</b>	58.89
SHMM	MFCC + $\Delta$ + $\Delta\Delta$	GE+PO	73.94	74.47	74.20	58.23
SHMM	MFCC + $\Delta$ + $\Delta\Delta$	GE+PO+SP	75.03	74.00	74.51	56.91
SHMM	MBN	GE	56.57	69.34	62.31	55.14
SHMM	MBN	GE+PO	59.18	69.12	63.76	54.1
SHMM	MBN	GE+PO+SP	60.89	68.41	64.43	<b>49.2</b>

iterations. Results, shown in Fig. 2, indicate that the SHMM




Figure 2: PER of the SHMM for varying subspace dimension.

is perfectly able to learn the phonetic subspace of a language by compressing the 3861-dimensional<sup>2</sup> parameter space to a subspace as small as 30 dimensions and yet achieving the same PER as the HMM baseline.

### 3.3. Acoustic Unit Discovery

We now consider the case of unsupervised learning of speech where English is assumed to be a low-resourced language. In this setup, we use the complete TIMIT set (training, development and testing set including the SA\* utterances) as the corpus from which to extract acoustic units. In this experiment, all the HMM/SHMM have 4 Gaussian components per state. Our baseline is the HMM based AUD system described in [3] and the VAE-(B)HMM based AUD system proposed in [5, 6]. We compare these baselines with 3 SHMM based AUD models for which the posterior of the phonetic subspace  $q(\mathbf{W}, \mathbf{b})$  was estimated using: (i) German language (ii) German and Polish languages (iii) German, Polish and Spanish languages. For each case the phonetic subspace had 35, 70 and 100 dimensions respectively. Note that the choice of the languages and the order of combination was arbitrary and it is likely that choosing languages closely related to the target language would be beneficial. We considered all the phones of all the languages to be unique and didn't merge them while estimating the subspace. The posteriors of the embeddings  $q(\mathbf{h}_u)$  corresponding to the German, Polish and Spanish phones were discarded before the AUD clustering.

The results are presented in Table 1 and differ significantly depending on the input features. The SHMM always benefits from learning the phonetic subspace in terms of eq.

<sup>2</sup>3 states  $\times$  (8 Gaussian  $\times$  2  $\times$  80 + 7). 80 is the features dimension, 2 accounts for the mean and the diagonal of the covariance matrix and 7 is the dimension of the per-state mixture weights.

PER. Interestingly, the baseline HMM fails to benefit from the MBN features as it underperforms compared to the HMM trained on MFCC features. The SHMM, thanks to its subspace, learns from other languages to fully exploit the discriminatively trained features. Regarding the segmentation evaluation, the SHMM better segments the speech compared to the simple HMM. However, we observe that using more than one language does not necessarily improve the segmentation. Also, contrary to the eq. PER, the MBN features do not seem to be ideal to get accurate segmentation.

Finally, we tried to label the TIMIT corpus with a HMM phone-recognizer (MBN features) trained on German, German and Polish and German, Polish and Spanish and we interpreted the output phones as acoustic units. For these 3 models the eq. PER was 61.22% (GE), 66.47% (GE+PO) and 71.96% (GE+PO+SP). Contrary to the SHMM, this naive approach does not benefit from having more languages.

## 4. Conclusions

We proposed a new model for AUD: the Subspace HMM. Unlike other AUD models the SHMM is trained in a supervised fashion on one or several languages to learn the notion of "phone". This phonetic knowledge is encoded into a non-linear subspace of the total parameter space. Then, the SHMM searches a set of acoustic units in this subspace which maximizes the likelihood of the observations of the target language. The SHMM outperforms the HMM based AUD and is competitive with the VAE-HMM. When using discriminatively trained features, the SHMM achieves 49.2% equivalent PER on TIMIT without any supervision in the target language.

## 5. Acknowledgements

The work was supported by Czech National Science Foundation (GACR) project "NEUREM3" No. 19-26934X, Czech Ministry of Interior project No. VI20152020025 "DRAPAK", and Czech Ministry of Education, Youth and Sports from the National Programme of Sustainability (NPU II) project "IT4Innovations excellence in science - LQ1602". This work was also supported by the Office of the Director of National Intelligence (ODNI), Intelligence Advanced Research Projects Activity (IARPA) MATERIAL program, via Air Force Research Laboratory (AFRL) contract # FA8650-17-C-9118. The views and conclusions contained herein are those of the authors and should not be interpreted as necessarily representing the official policies, either expressed or implied, of ODNI, IARPA, AFRL or the U.S. Government.

The authors would like to thank Kateřina Žmolíková for her corrections and insightful comments.

## 6. References

- [1] J. R. Glass, "Towards unsupervised speech processing," in *ISSPA*. IEEE, 2012, pp. 1–4.
- [2] C. Lee and J. R. Glass, "A nonparametric bayesian approach to acoustic model discovery," in *ACL (1)*. The Association for Computer Linguistics, 2012, pp. 40–49.
- [3] L. Ondel, L. Burget, and J. Cernocky, "Variational inference for acoustic unit discovery," in *Procedia Computer Science*, vol. 2016, no. 81. Elsevier Science, 2016, pp. 80–86.
- [4] L. Ondel, L. Burget, J. Černocký, and S. Kesiraju, "Bayesian phonotactic language model for acoustic unit discovery," in *Proceedings of ICASSP 2017*. IEEE Signal Processing Society, 2017, pp. 5750–5754.
- [5] J. Ebberts, J. Heymann, L. Drude, T. Glarner, R. Haeb-Umbach, and B. Raj, "Hidden markov model variational autoencoder for acoustic unit discovery," in *Interspeech 2017, 18th Annual Conference of the International Speech Communication Association, Stockholm, Sweden, August 20-24, 2017*, 2017, pp. 488–492. [Online]. Available: [http://www.isca-speech.org/archive/Interspeech\\_2017/abstracts/1160.html](http://www.isca-speech.org/archive/Interspeech_2017/abstracts/1160.html)
- [6] T. Glarner, P. Hanebrink, J. Ebberts, and R. Haeb-Umbach, "Full bayesian hidden markov model variational autoencoder for acoustic unit discovery," in *Interspeech*. ISCA, 2018, pp. 2688–2692.
- [7] L. Ondel, P. Godard, L. Besacier, E. Larsen, M. Hasegawa-Johnson, O. Scharenborg, E. Dupoux, L. Burget, F. Yvon, and S. Khudanpur, "Bayesian Models for Unit Discovery on a very Low Resource Language," in *IEEE International Conference on Acoustics, Speech and Signal Processing*, ser. ICASSP, Calgary, Canada, 2018. [Online]. Available: <sources/Ondel18bayesian.pdf>
- [8] D. P. Kingma and M. Welling, "Auto-encoding variational bayes," in *ICLR*, 2014.
- [9] D. Povey, A. Ghoshal, G. Boulianne, N. Goel, M. Hannemann, Y. Qian, P. Schwarz, and G. Stemmer, "The kaldi speech recognition toolkit," in *In IEEE 2011 workshop*, 2011.
- [10] P. Kenny, "Joint factor analysis of speaker and session variability: Theory and algorithms," Tech. Rep., 2005.
- [11] S. Kesiraju, L. Burget, I. Szóke, and J. Černocký, "Learning document representations using subspace multinomial model," in *Proceedings of Interspeech 2016*. International Speech Communication Association, 2016, pp. 700–704.
- [12] L. R. Rabiner, "A tutorial on hidden markov models and selected applications in speech recognition," in *PROCEEDINGS OF THE IEEE*, 1989, pp. 257–286.
- [13] D. P. Kingma and J. Ba, "Adam: A method for stochastic optimization," in *ICLR*, 2015.
- [14] J. S. Garofolo, L. F. Lamel, W. M. Fisher, J. G. Fiscus, D. S. Pallett, and N. L. Dahlgren, "Darpa timit acoustic phonetic continuous speech corpus cdrom," 1993.
- [15] T. Schultz, N. T. Vu, and T. Schlippe, "Globalphone: A multilingual text & speech database in 20 languages." in *ICASSP*. IEEE, 2013, pp. 8126–8130.
- [16] R. Fer, P. Matejka, F. Grezl, O. Plchot, K. Vesely, and J. H. Cernocky, "Multilingually trained bottleneck features in spoken language recognition," *Computer Speech & Language*, vol. 46, no. Supplement C, pp. 252 – 267, 2017.
- [17] 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.
- [18] H.-W. H. Kai-Fu Lee, "Speaker-independent phone recognition using hidden markov models," vol. IEEE Transactions on acoustics, speech and signal processing, no. 37, pp. 1641–1648, 1989.