



A Temporal Extension of Latent Dirichlet Allocation for Unsupervised Acoustic Unit Discovery

Werner van der Merwe, Herman Kamper, Johan du Preez

MediaLab, E&E Engineering, Stellenbosch University

20695470@sun.ac.za, kamperh@sun.ac.za, dupreez@sun.ac.za

Abstract

Latent Dirichlet allocation (LDA) is widely used for unsupervised topic modelling on sets of documents. No temporal information is used in the model. However, there is often a relationship between the corresponding topics of consecutive tokens. In this paper, we present an extension to LDA that uses a Markov chain to model temporal information. We use this new model for acoustic unit discovery from speech. As input tokens, the model takes a discretised encoding of speech from a vector quantised (VQ) neural network with 512 codes. The goal is then to map these 512 VQ codes to 50 phone-like units (topics) in order to more closely resemble true phones. In contrast to the base LDA, which only considers how VQ codes co-occur within utterances (documents), the Markov chain LDA additionally captures how consecutive codes follow one another. This extension leads to an increase in cluster quality and phone segmentation results compared to the base LDA. Compared to a recent vector quantised neural network approach that also learns 50 units, the extended LDA model performs better in phone segmentation but worse in mutual information.

Index Terms: unsupervised learning, acoustic unit discovery, Bayesian inference, latent Dirichlet allocation.

1. Introduction

The goal in unsupervised acoustic unit discovery (AUD) is to find a finite set of phone-like units from unlabelled speech that resembles the phonetic inventory of a language [1,2]. Annotating speech is expensive and often not possible for many low-resource languages. Unsupervised AUD could be a way to get around this bottleneck.

Recently vector-quantised (VQ) neural networks have performed well in unsupervised AUD [3–5]. Approaches include the VQ variational autoencoder (VAE) [6] and contrastive predictive coding (CPC) [7] models. The output of these models is a discrete encoding of speech at a fixed rate. Despite good performance, the number of VQ codes in these models (512 for the VQ model in this paper) is often far more than the number of true phone units used in a language (in the order of 50). To more closely match the phonetic inventory of a language, the larger set of VQ codes therefore need to be mapped to a smaller set [8]. Our goal in this paper is to develop a new model for this task; the model will specifically operate on encodings with 512 VQ codes and attempt to map these down to a fixed number of units (50 in this paper).

Our new approach is an extension of latent Dirichlet allocation (LDA). LDA is a topic model that has been used across different domains [9]. In natural language processing, an LDA is commonly used to find latent topics in sets of documents. The LDA model gains information from two main assumptions. Firstly, each document is comprised of a subset of topics. Secondly, a topic has a subset of words associated with that topic.

We can extend the document assumptions of an LDA model towards AUD of speech. Firstly, an utterance of speech contains a subset of phones. Secondly, there exists a subset of VQ codes that corresponds to a phone. In the LDA terminology, an utterance would therefore correspond to a document, a VQ code to a word token, and a phone unit to a topic. In this way we could apply LDA to map the larger set of VQ codes down to a smaller set of phone units, thereby performing AUD. However, the LDA assumptions alone might not be enough: LDA considers each utterance (document) as a bag-of-VQ-codes (bag-of-words) and disregards any order information.

The contribution of this paper is a new model that extends LDA with temporal information. Concretely, our extended LDA model includes a Markov chain between adjacent phone units to model transitions. To perform AUD, a two-step approach is therefore followed. Firstly, a VQ-VAE model [10] is used to extract utterances with 512 VQ codes. Secondly, the Markov chain LDA model uses these VQ codes to discover 50 latent phone-like units.

The AUD approach we introduce differs from other AUD implementations that also use VQ codes as input but do not reduce the number of possible phone-like units [11, 12]. We use Bayesian inference for our model to capture uncertainty and incorporate useful assumptions [13]. Bayesian inference has also been used in several other leading AUD approaches [14–19]. Many of these models attempt to learn a hidden Markov model for each acoustic unit [14, 16, 18, 19]. However, none of these operate on state-of-the-art VQ codes as we do here but rather take conventional acoustic representations as input.

We compare the base LDA and expanded Markov chain LDA to an established CPC-Big + K -means model [3] that clusters CPC representations directly into 50 units. We compare the different approaches both in speaker-dependent and speaker-independent settings. The evaluation considers phone segmentation and clustering quality. Segmentation considers correct boundary placement while clustering considers whether the units are grouped according to the true phone categories. The phone segmentation scores for the Markov chain LDA is better than that of the base LDA and CPC-Big + K -means model. Mutual information (measuring cluster quality) with the Markov chain LDA is better than the base LDA but worse than that of the CPC-Big + K -means model. The extended Markov chain LDA therefore provides a clear benefit in segmentation for AUD.

2. Latent Dirichlet allocation

LDA is a generative probabilistic model used for topic discovery on collections of discrete data, such as the words in sets of documents [9]. The model assumes that each document, based on the observed words in the document has underlying topics as latent variables. Documents that contain similar observed words will have overlapping topics. The LDA model is a bag-of-words

model where the order of words in a document is disregarded. LDA assumes the number of topics and words are known.

2.1. Model description

A cluster graph is a representation of a probabilistic graphical model (PGM) that is closely related to factor graphs [20–22]. We start by formulating LDA within this generalised representation, since this enables us to easily extend it with a Markov chain (Section 3). In a cluster graph, two connected clusters, Φ_a and Φ_b , have a sepset, $S_{a,b}$, of shared random variables between them [23, 24]. The cluster graph of the general LDA model can be seen in Figure 1. M , N , K and V are respectively the number of documents, words per document, latent topics and vocabulary size.

The model is defined as follows:

- The M number of Φ_θ clusters are the per-document Dirichlet prior distributions, $p(\theta_m)$, of topics with parameter vector α .
- The K number of Φ_φ clusters are the per-topic Dirichlet prior distributions, $p(\varphi_k)$, of words with parameter vector β .
- The $M \cdot N$ number of Φ_Z clusters are categorical distributions, $p(Z|\theta)$ with K categories.
- The cluster Φ_W is a conditional categorical distribution, $p(W|Z, \theta)$, with V categories.

2.2. Inference

The LDA model we describe is Bayesian in origin and the resulting distributions are intractable. The LDA proposed by [9] performed inference using a convexity-based variational approach. We use the Lauritzen-Spiegelhalter or loopy belief update (LBU) algorithm [25] to perform inference. LBU sends messages between clusters in the graph by updating the cluster, Ψ , and sepset, ψ , beliefs. Due to the loopy nature of the graph, belief propagation will converge towards an approximation [26].

The **initialisation** of the model is performed as follows:

- Model the internal cluster factors, Φ_a, Φ_b, \dots
- Construct the sepsets, $S_{a,b}$.
- Construct the messages, $\mu_{a,b}$.
- Calculate cluster beliefs, Ψ_a , from all incoming messages and internal cluster factors:

$$\Psi_a = \Phi_a \prod_i \mu_{i,a} \quad (1)$$

- Calculate sepset beliefs, $\psi_{a,b}$, from the message pairs between clusters:

$$\psi_{a,b} = \mu_{a,b} \mu_{b,a} \quad (2)$$

The **message passing** between clusters using LBU is performed as follows:

1. Update the sepset belief by marginalising out the cluster belief:

$$\psi'_{a,b} = \sum_{\setminus S_{a,b}} \Psi_a \quad (3)$$

2. Update the cluster belief by replacement of the old sepset belief with the new one:

$$\Psi'_b = \Psi_b \frac{\psi'_{a,b}}{\psi_{a,b}} \quad (4)$$

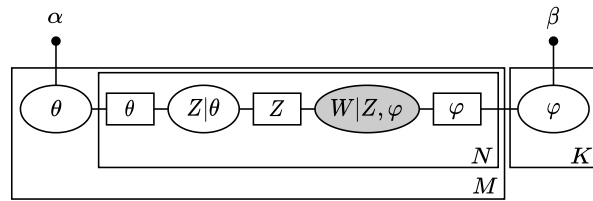


Figure 1: *The cluster graph of latent Dirichlet allocation model.*

3. Measure convergence using the Kullback-Leibler (KL) divergence [27] between the new and old sepset belief:

$$\text{KL}(\psi'_{ab} || \psi_{ab}) \quad (5)$$

4. Replace the old sepset belief with the new sepset belief:

$$\psi_{ab} = \psi'_{ab} \quad (6)$$

The cluster graph forms parallel paths for each document. This allows the parallelisation of message passing. A forward and backward pass is done through the model for each iteration of message passing until convergence.

2.3. Application towards acoustic unit discovery

Bayesian models have widely been used as an approach towards unsupervised AUD [14, 18, 19, 28, 29]. Bayesian models often do not require large amounts of data and perform well in low resource or unsupervised settings [13, 20]. Strides have been made by using neural networks to find phone-like units in speech [3, 7, 10–12, 30]. VQ-VAE and VQ-CPC models have been some of the predominant methods used [6, 7]. These output VQ codes have 512 possible values, far exceeding the expected number of acoustic units in a language. Further, the fixed-length codes are often oversegmented and do not capture the appropriate phone boundaries.

These discrete VQ codes can be used as input for an LDA model. The assumptions that are used by an LDA model is that the written words in a document contribute to a mixture of latent topics. An utterance is comprised of a subset of acoustic units. An analogous assumption would be that the codes in an utterance contribute to a mixture of latent acoustic units or phones. From this point on we will use the AUD terminology of utterances, phones and VQ codes instead of the LDA terminology of documents, topics and words.

In this paper, LDA was used to reduce fixed-length VQ-VAE [10] codes with 512 codebook entries towards 50 discrete codes. The choice of 50 discrete codes allows for a fair comparison to the benchmark CPC-Big + K -means model [3]. The phonetic vocabulary size is set to a fixed integer K , and the model does not infer this value as done in other studies [28, 29]. The output, Z_n , is a categorical distribution of possible latent phones. We use the category with the highest probability to output discrete acoustic units. In some practical applications, one might prefer keeping the soft speech units as done in [31].

While this assumption should contribute information towards achieving AUD, it alone is not sufficient. The LDA model is a bag-of-VQ-codes model and considers no temporal information. There is no modelling of phone length or VQ code transition probabilities. Further loss of information occurs by only using the discrete VQ codes disregarding the information present in the vector codebook values. Due to the unsupervised nature of the model, uninformative priors were used.

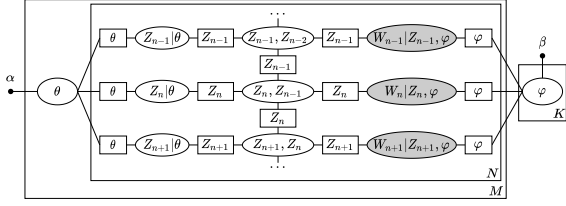


Figure 2: The cluster graph of the Markov chain LDA model.

3. Markov Chain LDA

We extended the base LDA model to include temporal information about the VQ codes. Due to the modular nature of PGMs, the inclusion of new information into a model is often straightforward [21]. The cluster graph representation we use provides flexibility to expand the model. The temporal dependency can be included as a new cluster and inserted into the cluster graph structure while retaining the original clusters.

Oversegmentation of phones is one of the hurdles in AUD. Since the codes used have a frame length of 20 ms, shorter than most speech units, we expect a single phone to mostly be a sequence of the same codes. If the system oversegments an input utterance many of these longer sequences are broken and too many phone boundaries are inserted. The base LDA also suffers from not knowing that two consecutive phone-like units are likely to belong to the same phone. To address this, we model the transition probability between consecutive phone-like units.

3.1. First-order Markov chain model

To include temporal information into the base LDA model, we have included a first-order discrete-time Markov chain into the LDA model. The new cluster, $\Phi_{Z_n|Z_{n-1}}$ is a two-dimensional discrete joint distribution with K^2 probabilities. The cluster models the transition probability of the latent phone codes to have a higher probability of repeating. The distribution used has values:

$$\Phi_{Z_n, Z_{n-1}} = \begin{cases} \frac{a}{K^2 + K(a-1)}, & Z_n = Z_{n-1} \\ \frac{1}{K^2 + K(a-1)}, & Z_n \neq Z_{n-1} \end{cases} \quad (7)$$

Where a is a non-negative scalar. Values of $a > 1$ decrease the probability of a code differing from the previous code. This reduces the segmentation of phones and encourages longer sequences of the same codes. The conditional probability of $p(Z_n|Z_{n-1})$ is a times as likely to transition to the same phone than a different phone.

$$p(Z_n|Z_n = Z_{n-1}) = a \cdot p(Z_n|Z_n \neq Z_{n-1}) \quad (8)$$

The expanded LDA model can be seen in Figure 2. The new cluster was included between the two categorical clusters $\Phi_{W_n|Z_n, \varphi_k}$ and $\Phi_{Z_n|\theta_m}$. The system is still parallelised over each document and can be executed asynchronously. Each of the new clusters is connected with a Markov chain to the contiguous cluster.

3.2. Message passing

The message passing order now becomes more complex due to the insertion of the Markov chain. The message passing per iteration is described in Algorithm 1. The paths along each observed code are no longer parallel and the order of message passing will influence the inference in the graph. The per utterance loop can still be parallelised.

Algorithm 1 The Markov chain LDA message schedule.

for each utterance m in $\{1, 2 \dots M\}$ **do**

Forward Pass

for each code n in $\{1, 2 \dots N\}$ **do**

$\Phi_{\theta_m} \rightarrow \Phi_{Z_n|\theta_m}$

$\Phi_{Z_n|\theta_m} \rightarrow \Phi_{Z_n|Z_{n-1}}$

$\Phi_{Z_n|Z_{n-1}} \rightarrow \Phi_{Z_{n+1}|Z_n}$

$\Phi_{Z_n|Z_{n-1}} \rightarrow \Phi_{W_n|Z_n, \varphi_k}$

for each phone k in $\{1, 2 \dots K\}$ **do**

$\Phi_{W_n|Z_n, \varphi_k} \rightarrow \Phi_{\varphi_k}$

end for

end for

Backward Pass

for each code n in $\{N \dots 2, 1\}$ **do**

for each phone k in $\{1, 2 \dots K\}$ **do**

$\Phi_{\varphi_k} \rightarrow \Phi_{W_n|Z_n, \varphi_k}$

end for

$\Phi_{W_n|Z_n, \varphi_k} \rightarrow \Phi_{Z_n|Z_{n-1}}$

$\Phi_{Z_n|Z_{n-1}} \rightarrow Z_{n-1}|Z_{n-2}$

$\Phi_{Z_n|Z_{n-1}} \rightarrow Z_n|\theta_m$

$\Phi_{Z_n|\theta_m} \rightarrow \Phi_{\theta_m}$

end for

end for

4. Experimental setup

We compare the base LDA and Markov chain LDA against more established methods of doing AUD. A fair comparison would be a system that also uses discrete codes as input and outputs 50 discrete phone-like units. The CPC-Big + K -means 50 model [3] provides a good benchmark to measure against. This model uses K -means clustering of the CPC representation vectors to find 50 clusters. The frame size is 10 ms compared to the 20 ms used by the systems we describe. The K -means method has access to the information of Euclidean distances between vector codebook values which the LDA systems do not have access to.

Dataset. The Buckeye speech corpus [32] of American English with 12 speakers is used for testing purposes. The dataset is conversational speech with male and female speakers. Speaker dependant results for each of the 12 speakers were calculated and we report the mean scores. Speaker independent results used the combined data set of 12 speakers. For each input utterance, continuous features within the VQ-VAE is discretised to the closest codebook entry, producing an encoding with 512 codes. The two LDA systems use these varying length utterances of VQ-VAE codes as input and produces an encoding of the input speech into 50 phone-like units.

Metrics. The two important aspects of AUD we report is phone segmentation and clustering quality. We compare the segmented boundaries of phones to the annotated ground truth labels. Phone segmentation measures the accuracy of the boundaries. We describe four metrics to measure the phone segmentation. A perfectly segmented result would score 100% on all four of these metrics. A tolerance of 20 ms from the label boundary was allowed to determine if a result boundary was correct. The precision, P , is the proportion of correct result boundaries out of the all predicted result boundaries. The recall, R , is the proportion of correct result boundaries out of all the label boundaries. The F-score, F_1 uses the harmonic mean of the precision and recall to give a score for segmentation. The F-score often gives high scores for oversegmented results. The R-value [33] gives a score for segmentation that penalises oversegmentation. The R-value

Table 1: Phone segmentation results (%) on Buckeye speech test set for speaker dependent and speaker independent results.

Model	P	R	F_1	$R\text{-val.}$
<i>Speaker dependent</i>				
CPC-Big + km-50 [3]	35.3	94.0	51.3	-44.4
VQ-VAE [10]	32.1	97.6	48.1	-80.7
LDA	36.0	94.9	52.00	-45.0
Markov chain LDA	55.1	72.4	62.2	55.6
<i>Speaker independent</i>				
CPC-Big + km-50 [3]	35.5	93.9	51.6	-42.5
VQ-VAE [10]	32.0	97.7	48.2	-76.2
LDA	36.8	93.3	52.8	-33.4
Markov chain LDA	55.4	66.4	60.4	61.6

can be negative if the results are severely oversegmented.

For cluster quality, we report on the cluster purity, singleton percentage and normalised mutual information (NMI). The cluster purity [34] gives a measure of a segment of codes overlapping with a single phone. Oversegmentation is not penalised and would increase the cluster purity. The singleton percentage is the number of singleton segments, a segment of one code, divided by the total number of codes. We want to reduce the number of singletons since a phone is rarely 20 ms long. The NMI measures the shared information of the segmented codes and annotated labels as a percentage of the information contained in the annotated labels.

Parameters. Due to the priors being probability distributions we need to include parameters for these distributions. As with most Bayesian systems, these parameters should allow the LDA models to model uncertainty. Based on validation set results the following parameters are used for testing purposes:

- $\alpha = \{1, 1 \dots 1\}$, flat Dirichlet parameters are used for the document Dirichlets.
- $\beta = \{10^{-4}, 10^{-4} \dots 10^{-4}\}$, sparse Dirichlet parameters are used for the topic Dirichlets.
- $a = 10$, the scalar in Equation 7 is set to prefer longer sequences of phones.

5. Results

The results are compared on metrics that measure the phone segmentation and cluster quality. The input, VQ-VAE [10] codes, of the LDA and Markov chain LDA is included. Special attention should be given to the performance of the Markov chain LDA compared to the base LDA since this would reveal whether the temporal expansion has contributed towards improved AUD.

5.1. Phone segmentation

As seen in Table 1, the CPC-Big + K -means, VQ-VAE input and the LDA model have poor segmentation results with high recall but negative $R\text{-val.}$ scores. The Markov chain LDA yields the highest F_1 and $R\text{-val.}$ scores for both the speaker dependant and speaker independent tests. In this case the $R\text{-val.}$ is positive because of the lower oversegmentation.

5.2. Cluster quality

In terms of cluster quality, the Markov chain LDA system performed better on all metrics than the base LDA as seen in Table 2.

Table 2: Clustering quality results (%) on Buckeye speech test set for speaker dependent and speaker independent results.

Model	Purity	Singletons	NMI
<i>Speaker dependent</i>			
CPC-Big + km-50 [3]	33.1	9.3	39.6
VQ-VAE [10]	34.0	26.7	38.4
LDA	24.1	20.2	26.6
Markov chain LDA	29.1	4.5	30.6
<i>Speaker independent</i>			
CPC-Big + km-50 [3]	32.3	9.3	37.8
VQ-VAE [10]	31.8	27.0	34.3
LDA	21.5	18.9	22.0
Markov chain LDA	27.3	4.0	24.7

It also has a lower singleton percentage than both the CPC-Big + K -means and LDA models. This is noteworthy because singletons often reflect incorrect boundary placement. There is a loss in NMI between the VQ-VAE and the Markov chain LDA of 7.8% and 9.7%. A loss in NMI is expected when reducing 512 possible input codes towards 50 possible output codes. The CPC-Big + K -means model performed the best in terms of NMI.

5.3. Discussion

The comparison of the Markov chain LDA to the base LDA model shows that the inclusion of temporal information improves AUD performance: it leads to an increase in phone segmentation quality in terms of F_1 and $R\text{-val.}$ scores that outcompete the CPC-Big + K -means model. The temporal information leads to an increase in NMI and a decrease in singletons compared to the base LDA. Despite good segmentation results from the Markov chain LDA, its NMI is still low in comparison to the CPC-Big + K -means model.

6. Conclusions and Future Work

We have shown that the inclusion of temporal information in an LDA model leads to an increase in mutual information. Concretely, the mutual information from our Markov chain LDA model was lower than that of the CPC-Big + K -means method we compared against, but gave better phone segmentation results. The Markov chain LDA achieved better phone segmentation results than the base LDA and the CPC-Big + K -means method. The extension of LDA has been shown to perform better in AUD than the base LDA model.

We speculate that the use of uninformative priors contributes towards the low mutual information. The information contained in the continuous output of the VQ codes, such as the vector space distance of a code to a cluster centre, could be used to give the system a more informative prior. Further extensions of the model could include forcing singleton cancellation and doing phone length modelling. Another extension would be a non-parametric model that infers the phonetic vocabulary size K .

7. Acknowledgements

This work is supported in part by the National Research Foundation of South Africa (grant no. 120409). Experiments were performed on the NICIS CHPC, South Africa. We thank Christiaan Jacobs and Matthew Baas for helpful suggestions.

8. References

- [1] L. Besacier, E. Barnard, A. Karpov, and T. Schultz, "Automatic speech recognition for under-resourced languages: A survey," *Speech communication*, vol. 56, pp. 85–100, 2014.
- [2] G. Adda, S. Stüker, M. Adda-Decker, O. Ambouroue, L. Besacier, D. Blachon, H. Bonneau-Maynard, P. Godard, F. Hamlaoui, D. Idiatov, G.-N. Kouarata, L. Lamel, E.-M. Makasso, A. Rialland, M. Van de Velde, F. Yvon, and S. Zerbian, "Breaking the unwritten language barrier: The BULB project," in *Proc. SLTU*, 2016.
- [3] 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," in *NeurIPS SAS Workshop*, 2020.
- [4] E. Dunbar, R. Algayres, J. Karadayi, M. Bernard, J. Benjumea, X.-N. Cao, L. Miskic, C. Dugrain, L. Ondel, A. Black, L. Besacier, S. Sakti, and E. Dupoux, "The Zero Resource Speech Challenge 2019: TTS without T," in *Proc. Interspeech*, 2019.
- [5] E. Dunbar, J. Karadayi, M. Bernard, X.-N. Cao, R. Algayres, L. Ondel, L. Besacier, S. Sakti, and E. Dupoux, "The Zero Resource Speech Challenge 2020: Discovering discrete subword and word units," in *Proc. Interspeech*, 2020.
- [6] A. van den Oord, O. Vinyals, and K. Kavukcuoglu, "Neural discrete representation learning," in *Proc. NeurIPS*, 2017.
- [7] A. Van den Oord, Y. Li, and O. Vinyals, "Representation learning with contrastive predictive coding," *arXiv preprint arXiv:1807.03748*, 2018.
- [8] S. Takahashi, S. Sakti, and S. Nakamura, "Unsupervised neural-based graph clustering for variable-length speech representation discovery of zero-resource languages," in *Proc. Interspeech*, 2021.
- [9] D. M. Blei, A. Y. Ng, and M. I. Jordan, "Latent Dirichlet allocation," *Journal of Machine Learning Research*, vol. 3, pp. 993–1022, 2003.
- [10] B. van Niekerk, L. Nortje, and H. Kamper, "Vector-quantized neural networks for acoustic unit discovery in the ZeroSpeech 2020 challenge," in *Proc. Interspeech*, 2020.
- [11] H. Kamper and B. van Niekerk, "Towards unsupervised phone and word segmentation using self-supervised vector-quantized neural networks," in *Proc. Interspeech*, 2021.
- [12] H. Kamper, "Word segmentation on discovered phone units with dynamic programming and self-supervised scoring," *arXiv preprint arXiv:2202.11929*, 2022.
- [13] C. M. Bishop, *Pattern Recognition and Machine Learning (Information Science and Statistics)*. Springer-Verlag, 2006.
- [14] C.-y. Lee and J. Glass, "A nonparametric Bayesian approach to acoustic model discovery," in *Proc. ACL*, 2012.
- [15] M. Heck, S. Sakti, and S. Nakamura, "Feature optimized DPGMM clustering for unsupervised subword modeling: A contribution to ZeroSpeech 2017," in *Proc. ASRU*, 2017.
- [16] J. Ebbens, J. Heymann, L. Drude, T. Glarner, R. Haeb-Umbach, and B. Raj, "Hidden Markov model variational autoencoder for acoustic unit discovery," in *Proc. Interspeech*, 2017.
- [17] M. Heck, S. Sakti, and S. Nakamura, "Learning supervised feature transformations on zero resources for improved acoustic unit discovery," *IEICE Transactions on Information and Systems*, vol. E101.D, pp. 205–214, 2018.
- [18] 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 *Proc. ICASSP*, 2018.
- [19] L. Ondel, H. K. Vydana, L. Burget, and J. Černocký, "Bayesian subspace hidden Markov model for acoustic unit discovery," in *Proc. Interspeech*, 2019.
- [20] D. Barber, *Bayesian Reasoning and Machine Learning*. Cambridge University Press, 2012.
- [21] D. Koller and N. Friedman, *Probabilistic graphical models: principles and techniques*. MIT press, 2009.
- [22] S. Streicher and J. du Preez, "Graph coloring: Comparing cluster graphs to factor graphs," in *Proc. SAWACMMM*, 2017.
- [23] S. Streicher, W. Brink, and J. du Preez, "A probabilistic graphical model approach to the structure-and-motion problem," in *Proc. PRASA-RobMech*, 2016.
- [24] S. Streicher, "Incremental inference on higher-order probabilistic graphical models applied to constraint satisfaction problems," *arXiv preprint arXiv:2202.12916*, 2022.
- [25] S. L. Lauritzen and D. J. Spiegelhalter, "Local computations with probabilities on graphical structures and their application to expert systems," *Journal of the Royal Statistical Society*, vol. 50, no. 2, pp. 157–224, 1988.
- [26] G. R. Shafer and P. P. Shenoy, "Probability propagation," *Annals of mathematics and Artificial Intelligence*, vol. 2, no. 1, pp. 327–351, 1990.
- [27] S. Kullback and R. A. Leibler, "On information and sufficiency," *Annals of Mathematics and Artificial Intelligence*, vol. 22, no. 1, pp. 79–86, 1951.
- [28] L. Ondel, L. Burget, and J. Černocký, "Variational inference for acoustic unit discovery," in *Proc. SLTU*, 2016.
- [29] L. Ondel, L. Burget, J. Černocký, and S. Kesiraju, "Bayesian phonotactic language model for acoustic unit discovery," in *Proc. ICASSP*, 2017.
- [30] 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, "Unsupervised acoustic unit discovery for speech synthesis using discrete latent-variable neural networks," in *Proc. Interspeech*, 2019.
- [31] B. van Niekerk, M.-A. Carbonneau, J. Zaidi, M. Baas, H. Seute, and H. Kamper, "A comparison of discrete and soft speech units for improved voice conversion," *arXiv preprint arXiv:2111.02392*, 2021.
- [32] M. A. Pitt, K. Johnson, E. Hume, S. Kiesling, and W. Raymond, "The buckeye corpus of conversational speech: Labeling conventions and a test of transcriber reliability," *Speech Communication*, vol. 45, no. 1, pp. 89–95, 2005.
- [33] O. J. Räsänen, U. K. Laine, and T. Altsaar, "An improved speech segmentation quality measure: the R-value," in *Proc. Interspeech*, 2009.
- [34] C. Manning, P. Raghavan, and H. Schütze, "Introduction to information retrieval," *Natural Language Engineering*, vol. 16, no. 1, pp. 100–103, 2010.