

EXTRACTING DOMAIN INVARIANT FEATURES BY UNSUPERVISED LEARNING FOR ROBUST AUTOMATIC SPEECH RECOGNITION

Wei-Ning Hsu, James Glass

MIT Computer Science and Artificial Intelligence Laboratory
Cambridge, MA 02139, USA
{wnhsu, glass}@mit.edu

ABSTRACT

The performance of automatic speech recognition (ASR) systems can be significantly compromised by previously unseen conditions, which is typically due to a mismatch between training and testing distributions. In this paper, we address robustness by studying domain invariant features, such that domain information becomes transparent to ASR systems, resolving the mismatch problem. Specifically, we investigate a recent model, called the Factorized Hierarchical Variational Autoencoder (FHVAE). FHVAEs learn to factorize sequence-level and segment-level attributes into different latent variables without supervision. We argue that the set of latent variables that contain segment-level information is our desired domain invariant feature for ASR. Experiments are conducted on Aurora-4 and CHiME-4, which demonstrate 41% and 27% absolute word error rate reductions respectively on mismatched domains.

Index Terms— robust speech recognition, factorized hierarchical variational autoencoder, domain invariant representations

1. INTRODUCTION

Recently, neural network-based acoustic models [1, 2, 3] have greatly improved the performance of automatic speech recognition (ASR) systems. Unfortunately, it is well known (e.g., [4]) that ASR performance can degrade significantly when testing in a domain that is mismatched from training. A major reason is that speech data have complex distributions and contain information about not only linguistic content, but also speaker identity, background noise, room characteristics, etc. Among these sources of variability, only a subset are relevant to ASR, while the rest can be considered as a nuisance and therefore hurt the performance if the distributions of these attributes are mismatched between training and testing.

To alleviate this issue, some robust ASR research focuses on mapping the out-of-domain data to in-domain data using enhancement-based methods [5, 6, 7], which generally requires parallel data from both domains. Another popular strategy is to train an ASR system with as large, and as diverse a dataset as possible [8, 9]; however, this strategy is not feasible when the labeled data are not available for all domains. Alternatively, robustness can also be achieved by training using features that are domain invariant [10, 11, 12, 13, 14]. In this case, we would not have domain mismatch issues, because domain information is now transparent to the ASR system.

In this paper, we consider the same highly adverse scenario as in [4], where both clean and noisy speech are available, but the transcripts are only available for clean speech. We study the use of a recently proposed model, called Factorized Hierarchical Variational Autoencoder (FHVAE) [15], for learning domain invariant

ASR features without supervision. FHVAE models learn to factorize sequence-level attributes and segment-level attributes into different latent variables. By training an ASR system on the latent variables that encode segment-level attributes, and testing the ASR in mismatched domains, we demonstrate that these latent variables contain linguistic information and are more domain invariant. Comprehensive experiments study the effect of different FHVAE architectures, training strategies, and the use of derived domain features on the robustness of ASR systems. Our proposed method is evaluated on Aurora-4 [16] and CHiME-4 [17] datasets, which contain artificially corrupted noisy speech and real noisy speech respectively. The proposed FHVAE-based feature reduces the absolute word error rate (WER) by 27% to 41% compared to filter bank features, and by 14% to 16% compared to variational autoencoder-based features. We have released the code of FHVAEs described in the paper.¹

The rest of the paper is organized as follows. In Section 2, we introduce the FHVAE model and a method to extract domain invariant features. Section 3 describes the experimental setup, while Section 4 presents results and discussion. We conclude our work in Section 5.

2. LEARNING DOMAIN INVARIANT FEATURES

2.1. Modeling a Generative Process of Speech Segments

As mentioned above, generation of speech data often involves many independent factors, which are however unseen in the unsupervised setting. It is therefore natural to describe such a generative process using a latent variable model, where a latent variable \mathbf{z} is first sampled from a prior distribution, and a speech segment \mathbf{x} is then sampled from a distribution conditioned on \mathbf{z} . In [18], a convolutional variational autoencoder (VAE) is proposed to model such process; by assuming the prior to be a diagonal Gaussian, it is shown that the VAE automatically learns to model independent attributes regarding generation, such as the speaker identity and the linguistic content, using orthogonal latent subspaces. This result provided a mechanism of potentially learning domain invariant features for ASR by discovering latent variables that do not contain domain information.

2.2. Extracting Domain Invariant Features from FHVAEs

The generation of sequential data often involves multiple independent factors operating at different scales. For instance, the speaker identity affects the fundamental frequency (F0) at the utterance level, while the phonetic content affects spectral characteristics at the segment level. As a result, sequence-level attributes, such as F0 and

¹<https://github.com/wnhsu/FactorizedHierarchicalVAE>

volume, tends to have a smaller amount of variation within an utterance, compared to between utterances, while the other attributes, such as spectral contours, tend to have similar amounts of variation within and between utterances.

Based on this observation, FHVAEs [15] formulate the generative process of sequential data with a factorized hierarchical graphical model that imposes sequence-dependent priors and sequence-independent priors to different sets of latent variables. Specifically, given a dataset $\mathcal{D} = \{\mathbf{X}^{(i)}\}_{i=1}^M$ consisting of M i.i.d. sequences, where $\mathbf{X}^{(i)} = \{\mathbf{x}^{(i,n)}\}_{n=1}^{N^{(i)}}$ is a sequence of $N^{(i)}$ segments (sub-sequence), a sequence \mathbf{X} of N segments is assumed to be generated from a random process that involves latent variables $\mathbf{Z}_1 = \{\mathbf{z}_1^{(n)}\}_{n=1}^N$, $\mathbf{Z}_2 = \{\mathbf{z}_2^{(n)}\}_{n=1}^N$, and $\boldsymbol{\mu}_2$ as follows: (1) an s -vector $\boldsymbol{\mu}_2$ is drawn from a prior distribution $p(\boldsymbol{\mu}_2) = \mathcal{N}(\boldsymbol{\mu}_2 | \mathbf{0}, \frac{2}{\mu_2} \mathbf{I})$; (2) N i.i.d. *latent segment variables* $\{\mathbf{z}_1^{(n)}\}_{n=1}^N$ and *latent sequence variables* $\{\mathbf{z}_2^{(n)}\}_{n=1}^N$ are drawn from a sequence-independent prior $p(\mathbf{z}_1) = \mathcal{N}(\mathbf{z}_1 | \mathbf{0}, \frac{2}{z_1} \mathbf{I})$ and a sequence-dependent prior $p(\mathbf{z}_2 | \boldsymbol{\mu}_2) = \mathcal{N}(\mathbf{z}_2 | \boldsymbol{\mu}_2, \frac{2}{z_2} \mathbf{I})$ respectively; (3) N i.i.d. speech segments $\{\mathbf{x}^{(n)}\}_{n=1}^N$ are drawn from a condition distribution $p(\mathbf{x} | \mathbf{z}_1, \mathbf{z}_2) = \mathcal{N}(\mathbf{x} | f_{\mu_x}(\mathbf{z}_1, \mathbf{z}_2), \text{diag}(f_{\frac{2}{x}}(\mathbf{z}_1, \mathbf{z}_2)))$, whose mean and diagonal variance are parameterized by neural networks. The joint probability for a sequence is formulated in Eq. 1:

$$p(\boldsymbol{\mu}_2) \prod_{n=1}^N p(\mathbf{x}^{(n)} | \mathbf{z}_1^{(n)}, \mathbf{z}_2^{(n)}) p(\mathbf{z}_1^{(n)}) p(\mathbf{z}_2^{(n)} | \boldsymbol{\mu}_2). \quad (1)$$

Based on this formulation, $\boldsymbol{\mu}_2$ can be regarded as a summarization of sequence-level attributes for a sequence, and \mathbf{z}_2 is encouraged to encode sequence-level attributes for a segment that are similar within an utterance. Consequently, \mathbf{z}_1 encodes the residual segment-level attributes for a segment, such that \mathbf{z}_1 and \mathbf{z}_2 together provide sufficient information for generating a segment.

Since exact posterior inference is intractable, FHVAEs introduce an inference model $q(\mathbf{Z}_1^{(i)}, \mathbf{Z}_2^{(i)}, \boldsymbol{\mu}_2^{(i)} | \mathbf{X}^{(i)})$ as formulated in Eq. 2 that approximates the true posterior $p(\mathbf{Z}_1^{(i)}, \mathbf{Z}_2^{(i)}, \boldsymbol{\mu}_2^{(i)} | \mathbf{X}^{(i)})$:

$$q(\boldsymbol{\mu}_2^{(i)}) \prod_{n=1}^{N^{(i)}} q(\mathbf{z}_1^{(i,n)} | \mathbf{x}^{(i,n)}, \mathbf{z}_2^{(i,n)}) q(\mathbf{z}_2^{(i,n)} | \mathbf{x}^{(i,n)}), \quad (2)$$

from which we observe that inference of $\mathbf{z}_1^{(i,n)}$ and $\mathbf{z}_2^{(i,n)}$ only depends on the corresponding segment $\mathbf{x}^{(i,n)}$; in particular, the posteriors, $q(\mathbf{z}_1 | \mathbf{x}, \mathbf{z}_2) = \mathcal{N}(\mathbf{z}_1 | g_{\mu_{z_1}}(\mathbf{x}, \mathbf{z}_2), \text{diag}(g_{\frac{2}{z_1}}(\mathbf{x}, \mathbf{z}_2)))$ and $q(\mathbf{z}_2 | \mathbf{x}) = \mathcal{N}(\mathbf{z}_2 | g_{\mu_{z_2}}(\mathbf{x}), \text{diag}(g_{\frac{2}{z_2}}(\mathbf{x})))$, are approximated with diagonal Gaussian distributions whose mean and diagonal variance are also parameterized by neural networks. On the other hand, $q(\boldsymbol{\mu}_2^{(i)})$ is modeled as an isotropic Gaussian, $\mathcal{N}(\boldsymbol{\mu}_2^{(i)} | g_{\mu_{\boldsymbol{\mu}_2}}(i), \frac{2}{\mu_2} \mathbf{I})$, where $g_{\mu_{\boldsymbol{\mu}_2}}(i)$ is a trainable lookup table of the posterior mean of $\boldsymbol{\mu}_2$ for each training sequence. Estimation of $\boldsymbol{\mu}_2$ for testing sequences can be found in [15].

As pointed out in [4], nuisance attributes regarding ASR, such as speaker identity, room geometry, and background noise, are generally consistent within an utterance. If we treat each utterance as a sequence, these attributes then become sequence-level attributes, which would be encoded by \mathbf{z}_2 and $\boldsymbol{\mu}_2$. As a result, \mathbf{z}_1 encodes the residual linguistic information and is invariant to these nuisance attributes, which is our desired domain invariant ASR feature.

2.3. Training FHVAE and Preventing S-Vector Collapsing

As in other generative models, FHVAEs aim to maximize the marginal likelihood of the observed dataset; due to the intractability of the exact posterior, FHVAEs optimize the *segment variational lower bound*, $\mathcal{L}(\cdot, \cdot; \mathbf{x}^{(i,n)})$, which is formulated as follows:

$$\begin{aligned} & \mathbb{E}_q(\mathbf{z}_1^{(i,n)}, \mathbf{z}_2^{(i,n)} | \mathbf{x}^{(i,n)}) \log p(\mathbf{x}^{(i,n)} | \mathbf{z}_1^{(i,n)}, \mathbf{z}_2^{(i,n)}) \\ & - \mathbb{E}_q(\mathbf{z}_2^{(i,n)} | \mathbf{x}^{(i,n)}) D_{KL}(q(\mathbf{z}_1^{(i,n)} | \mathbf{x}^{(i,n)}, \mathbf{z}_2^{(i,n)}) || p(\mathbf{z}_1^{(i,n)})) \\ & - D_{KL}(q(\mathbf{z}_2^{(i,n)} | \mathbf{x}^{(i,n)}) || p(\mathbf{z}_2^{(n)} | g_{\mu_{\boldsymbol{\mu}_2}}(i))) + \frac{1}{N} \log p(g_{\mu_{\boldsymbol{\mu}_2}}(i)). \end{aligned}$$

Notice that if the $\boldsymbol{\mu}_2$ are the same for all utterances, an FHVAE would then degenerate to a vanilla VAE. To prevent $\boldsymbol{\mu}_2$ from collapsing, we can add an additional discriminative objective, $\log p(i | \mathbf{z}_2^{(i,n)})$, that encourages the discriminability of \mathbf{z}_2 regarding which utterance the segment is drawn from. Specifically, we define it as $\log p(\mathbf{z}_2^{(i,n)} | g_{\mu_{\boldsymbol{\mu}_2}}(i)) - \log \prod_{j=1}^M p(\mathbf{z}_2^{(i,n)} | g_{\mu_{\boldsymbol{\mu}_2}}(j))$. By combining the two objectives with a weighting parameter λ , we obtain the *discriminative segment variational lower bound*:

$$\mathcal{L}^{dis}(\cdot, \cdot; \mathbf{x}^{(i,n)}) = \mathcal{L}(\cdot, \cdot; \mathbf{x}^{(i,n)}) + \lambda \log p(i | \mathbf{z}_2^{(i,n)}). \quad (3)$$

3. EXPERIMENT SETUP

To evaluate the effectiveness of the proposed method on extracting domain invariant features, we consider domain mismatched ASR scenarios. Specifically, we train an ASR system using a clean set, and test the system on both a clean and noisy set. The idea is that one would observe a smaller performance discrepancy between different domains if the feature representation is more domain invariant. We next introduce the datasets, as well as the model architectures and training configurations for the experiments.

3.1. Dataset

We use Aurora-4 [16] as the primary dataset for our experiments. Aurora-4 is a broadband corpus designed for noisy speech recognition tasks based on the Wall Street Journal (WSJ0) corpus [19]. Two microphone types, clean/channel are included, and six noise types are artificially added to both microphone types, which results in four conditions: clean(A), channel(B), noisy(C), and channel+noisy(D). We use the multi-condition development set for training the VAE and FHVAE models, because the development set contains both noise labels and speaker labels for each utterance, which are used in *Exp. Index 5*, while the training set only contains speaker labels. The ASR system is trained on the clean `train_si84_clean` set and evaluated on the multi-condition `test_eval192` set.

To verify our proposed method on a non-artificial dataset, we repeat our experiments on the CHiME-4 [17] dataset, which contains real distant-talking recordings in noisy environments. We use the original 7,138 clean utterances and the 1,600 single channel real noisy utterances in the training partition to train the VAE and FHVAE models. The ASR system is trained on the original clean training set and evaluated on the CHiME-4 development set.

3.2. VAE/FHVAE Setup and Training

The VAE is trained with stochastic gradient descent using a mini-batch size of 128 without clipping to minimize the negative variational lower bound plus an L_2 -regularization with weight 10^{-4} .

Exp. Index	Feature	Setting		Seq. Label	WER (%) Avg.	WER (%) by Condition			
		#Layers	#Units			A	B	C	D
1	FBank	-	-	-	65.64	3.21	61.61	51.78	82.39
	\mathbf{z}	1/1	256/256	-	44.79	4.22	38.16	36.11	59.63
	\mathbf{z}	1/1	512/256	-	40.31	4.35	33.83	34.43	53.77
	\mathbf{z}_1	1/1	256/256	10	26.58	4.54	19.28	20.85	38.50
2	\mathbf{z}_1	1/1	256/256	10	26.58	4.54	19.28	20.85	38.50
	\mathbf{z}_1	2/2	256/256	10	25.54	4.11	16.90	20.62	38.58
	\mathbf{z}_1	3/3	256/256	10	24.30	4.91	15.44	22.83	36.63
3	\mathbf{z}_1	1/1	128/128	10	34.66	5.06	26.70	25.39	49.09
	\mathbf{z}_1	1/1	256/256	10	26.58	4.54	19.28	20.85	38.50
	\mathbf{z}_1	1/1	512/512	10	26.97	5.32	18.18	23.13	40.01
4	\mathbf{z}_1	1/1	256/256	0	33.30	4.86	25.67	25.46	46.97
	\mathbf{z}_1	1/1	256/256	5	30.55	4.63	22.66	23.33	43.96
	\mathbf{z}_1	1/1	256/256	10	26.58	4.54	19.28	20.85	38.50
	\mathbf{z}_1	1/1	256/256	15	29.92	5.01	20.82	24.79	44.03
	\mathbf{z}_1	1/1	256/256	20	32.64	5.57	25.48	24.53	45.66
5	\mathbf{z}_1	1/1	256/256	10	26.58	4.54	19.28	20.85	38.50
	\mathbf{z}_1	1/1	256/256	10	32.27	4.33	23.89	28.96	45.86
	\mathbf{z}_1	1/1	256/256	10	34.95	4.39	27.27	32.22	48.20
6	\mathbf{z}_1	1/1	256/256	10	26.58	4.54	19.28	20.85	38.50
	$\mathbf{z}_1 - \mu_2$	1/1	256/256	10	43.61	5.08	42.47	27.55	53.85

Table 1. Aurora-4 test_eval92 set word error rate of acoustic models trained on different features.

The Adam [20] optimizer is used with $\beta_1 = 0.95$, $\beta_2 = 0.999$, $\epsilon = 10^{-8}$, and initial learning rate of 10^{-3} . Training is terminated if the lower bound on the development set does not improve for 50 epochs. The FHVAE is trained with the same configuration and optimization method, except that the loss function is replaced with the negative discriminative segment variational lower bound.

Seq2Seq-VAE [4] and Seq2Seq-FHVAE [15] architectures with LSTM units are used for all experiments. We let the latent space of the VAEs contain 64 dimensions. Since the FHVAE models have two latent spaces, we let each of them be 32 dimensional. Other hyperparameters are explored in our experiments. Inputs to VAE/FHVAE, \mathbf{x} , are chunks of 20 consecutive speech frames randomly drawn from utterances, where each frame is represented as 80 dimensional filter bank (FBank) energies. To extract features from the VAE and FHVAE for ASR training, for each utterance, we compute and concatenate the posterior mean and variance of chunks shifted by one frame, which generates a sequence of new features that are 19 frames shorter than the original sequence. We pad the first frame and the last frame at each end to match the original length.

3.3. ASR Setup and Training

Kaldi [21] is used for feature extraction, decoding, forced alignment, and training of an initial HMM-GMM model on the original clean utterances. The recipe provided by the CHiME-4 challenge (`run_gmm.sh`) and the Kaldi Aurora-4 recipe are adapted by only changing the training data being used. The Computational Network Toolkit (CNTK) [22] is used for neural network-based acoustic model training. For all experiments, the same LSTM acoustic model [23] with the architecture proposed in [24] is applied, which has 1,024 memory cells and a 512-node projection layer for each LSTM layer, and 3 LSTM layers in total.

Following the training setup in [25], LSTM acoustic models are trained with a cross-entropy criterion, using truncated backpropagation-through-time (BPTT) [26] to optimize. Each

BPTT segment contains 20 frames, and each mini-batch contains 80 utterances, since we find empirically that 80 utterances has similar performance to 40 utterances. A momentum of 0.9 is used starting from the second epoch [3]. Ten percent of the training data is held out as a validation set to control the learning rate. The learning rate is halved when no gain is observed after an epoch. The same language model is used for decoding for all experiments.

4. EXPERIMENTAL RESULTS AND DISCUSSION

In this section, we report the experimental results on both datasets, and provide insights on the outcome. Table 1 and 2 summarize the results on Aurora-4 and CHiME-4 respectively. For both tables, different experiments are separated by double horizontal lines and indexed by the *Exp. Index* on the first column. The second column, *Feature*, refers to the frame representations used for training ASR models. The third to the sixth column explains the model configuration and the discriminative training weight for VAE or FHVAE models. We separate the encoder and decoder parameters by “/” in the third and the fourth column. Averaged and by-condition word error rate (WER) are shown in the rest of the columns.

4.1. Baseline

We start with establishing Aurora-4 baseline results trained on different types of feature representations, including (1) FBank, (2) latent variable, \mathbf{z} , extracted from the VAE, and (3) latent segment variable, \mathbf{z}_1 , extracted from the FHVAE. Because each FHVAE model has two encoders, to have a fair comparison between VAE and FHVAE models, we also consider a VAE model with 512 hidden units at each encoder layer. The results are shown in Table 1 *Exp. Index 1*. As mentioned, condition A is the matched domain, while conditions B, C, and D are all mismatched domains.

FBank degrades significantly in the mismatched conditions, producing between 49% to 79% absolute WER increase. On the other

Exp. Index	ASR Feature	Setting			WER (%)		WER (%) by Noise Type			
		#Layers	#Units	Seq. Label	Clean	Noisy	BUS	CAF	PED	STR
1	FBank	-	-	-	19.37	87.69	95.56	92.05	78.77	84.37
	\mathbf{z}	1/1	512/256	-	19.47	73.95	70.10	91.45	64.26	69.99
	\mathbf{z}_1	1/1	256/256	10 uttid	19.57	67.94	71.96	79.37	59.32	61.11
2	\mathbf{z}_1	1/1	256/256	10 uttid	19.57	67.94	71.96	79.37	59.32	61.11
	\mathbf{z}_1	2/2	256/256	10 uttid	19.73	62.44	71.28	71.86	52.46	54.18
	\mathbf{z}_1	3/3	256/256	10 uttid	19.52	60.39	69.13	66.24	51.22	54.96

Table 2. CHiME-4 development set word error rate of acoustic models trained on different features.

hand, both VAE and FHVAE models improve the performance in the mismatched domains by a large margin, with only a slight degradation in the matched domain. In particular, the features learned by the FHVAE consistently outperform the VAE features in all mismatched conditions by 14% absolute WER reduction.

We believe that this experiment verifies that FHVAEs can successfully retain domain invariant linguistic features in \mathbf{z}_1 , while encode domain related information into \mathbf{z}_2 . In contrast, as the results suggests, VAEs encode all the information into a single set of latent variables, \mathbf{z} , which still contain domain related information that can hurt ASR performance on the mismatched domains.

4.2. Comparing Model Architectures

We next explore the optimal FHVAE architectures for extracting domain invariant features. In particular, we study the effect of the number of hidden units at each layer and the number of layers. Results of each variant are listed in Table 1 *Exp. Index 2* and *Exp. Index 3* respectively. Regarding the averaged WER, the model with 256 hidden units at each layer and in total three layers achieves the lowest WER (24.30%). Interestingly, if we break down the WER by condition, it can be observed that increasing the FHVAE model capacity (i.e. increasing number of layers or hidden units) helps reducing the WER in the noisy condition (B), but deteriorates channel-mismatching condition (C) above 256 hidden units and 2 layers.

4.3. Effect of FHVAE Discriminative Training

Speaker verification experiments in [15] suggest that discriminative training facilitates factorizing segment-level attributes and sequence-level attributes into two sets of latent variables. Here we study the effect of discriminative training on learning robust ASR features, and show the results in Table 1 *Exp. Index 4*. When $\lambda = 0$, the model is not trained with the discriminative object. While increasing the discriminative weight from 0 to 10, we observe consistent improvement in all 4 conditions due to better factorization of segment and sequence information; however, when further increasing the weight to 20, the performance starts to degrade. This is because the discriminative object can inversely affects the modeling capacity by constraining the expressibility of the latent sequence variables.

4.4. Choice of Sequence Label

A core idea of FHVAE is to learn sequence-specific priors to model the generation of sequence-level attributes, which have a smaller amount of variation within a sequence. Suppose we treat each utterance as one sequence, then both speaker and noise information belongs to sequence-level attributes, because they are consistent within an utterance. Alternatively, we consider two FHVAE models that learn speaker-specific priors and noise-specific priors respectively.

This can be easily achieved by concatenating sequences of the same speaker label or noise label, and treating it as one sequence used for FHVAE training. We report the results in Table 1 *Exp. Index 5*.

It may at first seem surprising that utilizing supervised information in this fashion does not improve performance. We believe that concatenating utterances actually discards some useful information with respect to learning domain invariant features. FHVAEs use latent segment variables to encode attributes that are not consistent within a sequence. By concatenating speaker utterances, noise information is no longer consistent within sequences, and would thus be encoded into latent segment variables; similarly, latent segment variables would not be speaker invariant in the other case.

4.5. Use of S-Vector

Lastly, we study the use of s-vectors, μ_2 , derived from the FHVAE model, which can be seen as a summarization of sequence-level attributes of an utterance. We apply the same procedure as i-vector based speaker adaptation [27]: For each utterance, we first estimate its s-vector, and then concatenate s-vectors with the feature representation of each frame to generate the new feature sequence.

Results are shown in Table 1 *Exp. Index 6*, from which we observe a significant degradation of WER that is similar to those of the VAE models. This is reasonable because \mathbf{z}_1 and μ_2 in combination actually contains similar information as the latent variable \mathbf{z} in VAE models, and the degradation is due to the mismatch between the distributions of μ_2 in the training and testing sets.

4.6. Verifying Results on CHiME4

In this section, we repeat the baseline and the layer experiments on the CHiME-4 dataset, in order to verify the effectiveness of the FHVAE and the optimality of the FHVAE architecture on a non-artificial dataset. The results are shown in Table 2. From *Exp. Index 1*, we see that the same trend applies to the CHiME-4 dataset, where the latent segment variables from the FHVAE outperform those from the VAE, and both latent variable representations outperform FBank features. For the FHVAE architectures, a 7% absolute WER decrease is achieved by increasing the number of encoder/decoder layers from 1 to 3, which is also consistent with the trends we saw on Aurora-4.

5. CONCLUSION AND FUTURE WORK

In this paper, we conduct comprehensive experiments on studying the use of FHVAE models domain invariant ASR features extractors. Our feature demonstrates superior robustness in mismatched domains compared to FBank and VAE-based features by achieving 41% and 27% absolute WER reduction on Aurora-4 and CHiME-4 respectively. In the future, we plan to study FHVAE-based augmentation methods similar to [4].

6. REFERENCES

- [1] Tara N Sainath, Brian Kingsbury, George Saon, Hagen Soltau, Abdel-rahman Mohamed, George Dahl, and Bhuvana Ramabhadran, "Deep convolutional neural networks for large-scale speech tasks," *Neural Networks*, vol. 64, pp. 39–48, 2015.
- [2] Haşim Sak, Félix de Chaumont Quitry, Tara Sainath, Kanishka Rao, et al., "Acoustic modelling with cd-ctc-smbr lstm rnns," in *Automatic Speech Recognition and Understanding (ASRU), 2015 IEEE Workshop on*. IEEE, 2015, pp. 604–609.
- [3] Wei-Ning Hsu, Yu Zhang, and James Glass, "A prioritized grid long short-term memory rnn for speech recognition," in *Spoken Language Technology Workshop (SLT), 2016 IEEE*. IEEE, 2016, pp. 467–473.
- [4] Wei-Ning Hsu, Yu Zhang, and James Glass, "Unsupervised domain adaptation for robust speech recognition via variational autoencoder-based data augmentation," in *Automatic Speech Recognition and Understanding (ASRU), 2017 IEEE Workshop on*. IEEE, 2017.
- [5] Arun Narayanan and DeLiang Wang, "Ideal ratio mask estimation using deep neural networks for robust speech recognition," in *Acoustics, Speech and Signal Processing (ICASSP), 2013 IEEE International Conference on*. IEEE, 2013, pp. 7092–7096.
- [6] Yusuf Isik, Jonathan Le Roux, Zhuo Chen, Shinji Watanabe, and John R Hershey, "Single-channel multi-speaker separation using deep clustering," *arXiv preprint arXiv:1607.02173*, 2016.
- [7] Xue Feng, Yaodong Zhang, and James Glass, "Speech feature denoising and dereverberation via deep autoencoders for noisy reverberant speech recognition," in *Acoustics, Speech and Signal Processing (ICASSP), 2014 IEEE International Conference on*. IEEE, 2014, pp. 1759–1763.
- [8] Jinyu Li, Dong Yu, Jui-Ting Huang, and Yifan Gong, "Improving wideband speech recognition using mixed-bandwidth training data in cd-dnn-hmm," in *Spoken Language Technology Workshop (SLT), 2012 IEEE*. IEEE, 2012, pp. 131–136.
- [9] Michael L Seltzer, Dong Yu, and Yongqiang Wang, "An investigation of deep neural networks for noise robust speech recognition," in *Acoustics, Speech and Signal Processing (ICASSP), 2013 IEEE International Conference on*. IEEE, 2013, pp. 7398–7402.
- [10] Brian ED Kingsbury, Nelson Morgan, and Steven Greenberg, "Robust speech recognition using the modulation spectrogram," *Speech communication*, vol. 25, no. 1, pp. 117–132, 1998.
- [11] Richard M Stern and Nelson Morgan, "Features based on auditory physiology and perception," *Techniques for Noise Robustness in Automatic Speech Recognition*, p. 193227, 2012.
- [12] Oriol Vinyals and Suman V Ravuri, "Comparing multilayer perceptron to deep belief network tandem features for robust asr," in *Acoustics, Speech and Signal Processing (ICASSP), 2011 IEEE International Conference on*. IEEE, 2011, pp. 4596–4599.
- [13] Tara N Sainath, Brian Kingsbury, and Bhuvana Ramabhadran, "Auto-encoder bottleneck features using deep belief networks," in *Acoustics, Speech and Signal Processing (ICASSP), 2012 IEEE International Conference on*. IEEE, 2012, pp. 4153–4156.
- [14] Sining Sun, Binbin Zhang, Lei Xie, and Yanning Zhang, "An unsupervised deep domain adaptation approach for robust speech recognition," *Neurocomputing*, 2017.
- [15] Wei-Ning Hsu, Yu Zhang, and James Glass, "Unsupervised learning of disentangled and interpretable representations from sequential data," in *Advances in Neural Information Processing Systems*, 2017.
- [16] David Pearce, *Aurora working group: DSR front end LVCSR evaluation AU/384/02*, Ph.D. thesis, Mississippi State University, 2002.
- [17] Emmanuel Vincent, Shinji Watanabe, Aditya Arie Nugraha, Jon Barker, and Ricard Marxer, "An analysis of environment, microphone and data simulation mismatches in robust speech recognition," *Computer Speech & Language*, 2016.
- [18] Wei-Ning Hsu, Yu Zhang, and James Glass, "Learning latent representations for speech generation and transformation," in *Interspeech*, 2017, pp. 1273–1277.
- [19] John Garofalo, David Graff, Doug Paul, and David Pallett, "Csr-i (wsj0) complete," *Linguistic Data Consortium, Philadelphia*, 2007.
- [20] Diederik Kingma and Jimmy Ba, "Adam: A method for stochastic optimization," *arXiv preprint arXiv:1412.6980*, 2014.
- [21] Daniel Povey, Arnab Ghoshal, Gilles Boulianne, Lukas Burget, Ondrej Glembek, Nagendra Goel, Mirko Hannemann, Petr Motlicek, Yanmin Qian, Petr Schwarz, et al., "The kaldı speech recognition toolkit," in *IEEE 2011 workshop on automatic speech recognition and understanding*. IEEE Signal Processing Society, 2011, number EPFL-CONF-192584.
- [22] Dong Yu, Adam Eversole, Mike Seltzer, Kaisheng Yao, Zhiheng Huang, Brian Guenter, Oleksii Kuchaiev, Yu Zhang, Frank Seide, Huaming Wang, et al., "An introduction to computational networks and the computational network toolkit," Tech. Rep., Tech. Rep. MSR, Microsoft Research, 2014, <http://codebox/cntk>, 2014.
- [23] Hasim Sak, Andrew W Senior, and Françoise Beaufays, "Long short-term memory recurrent neural network architectures for large scale acoustic modeling.," in *Interspeech*, 2014, pp. 338–342.
- [24] Yu Zhang, Guoguo Chen, Dong Yu, Kaisheng Yao, Sanjeev Khudanpur, and James Glass, "Highway long short-term memory RNNs for distant speech recognition," in *2016 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2016, pp. 5755–5759.
- [25] Wei-Ning Hsu, Yu Zhang, Ann Lee, and James R Glass, "Exploiting depth and highway connections in convolutional recurrent deep neural networks for speech recognition.," in *INTER-SPEECH*, 2016, pp. 395–399.
- [26] Ronald J Williams and Jing Peng, "An efficient gradient-based algorithm for on-line training of recurrent network trajectories," *Neural computation*, vol. 2, no. 4, pp. 490–501, 1990.
- [27] George Saon, Hagen Soltau, David Nahamoo, and Michael Picheny, "Speaker adaptation of neural network acoustic models using i-vectors.," in *ASRU*, 2013, pp. 55–59.