REINFORCEMENT LEARNING OF SPEECH RECOGNITION SYSTEM BASED ON POLICY GRADIENT AND HYPOTHESIS SELECTION

Taku Kato, Takahiro Shinozaki

Tokyo Institute of Technology, School of Engineering, Kanagawa, Japan

ABSTRACT

Automatic speech recognition (ASR) systems have achieved high recognition performance for several tasks. However, the performance of such systems is dependent on the tremendously costly development work of preparing vast amounts of task-matched transcribed speech data for supervised training. The key problem here is the cost of transcribing speech data. The cost is repeatedly required to support new languages and new tasks. Assuming broad network services for transcribing speech data for many users, a system would become more self-sufficient and more useful if it possessed the ability to learn from very light feedback from the users without annoying them. In this paper, we propose a general reinforcement learning framework for ASR systems based on the policy gradient method. As a particular instance of the framework, we also propose a hypothesis selection-based reinforcement learning method. The proposed framework provides a new view for several existing training and adaptation methods. The experimental results show that the proposed method improves the recognition performance compared to unsupervised adaptation.

Index Terms— reinforcement learning, policy gradient method, hypothesis selection, deep neural network, speech recognition

1. INTRODUCTION

Today's ASR systems heavily rely on supervised training using large amounts of task-matched training data to achieve high recognition performance. To prepare labeled speech data, a large transcription cost is required. This is particularly a problem for resource-limited languages. However, even for resource-rich languages, a significant factor that limits the application area of ASR is the additional transcription cost required to support new tasks that are different from the initial training condition.

When considering network applications of ASR with many users, one strategy to improve the system performance without incurring development cost is to utilize feedback from the users while providing recognition results to them. Ogata et al. developed a web service called Podcastle that uses a speech recognizer to automatically transcribe speech contents in podcasts such that the users can read and search them [1, 2]. The system includes a user interface that allows the users to correct the recognition errors word by word. By gathering the corrected transcriptions, the ASR system can be re-estimated and improved by using any supervised model training or adaptation methods [3, 4, 5, 6]. For this system, the motivation for the users to fix the errors in the automatic transcriptions is to contribute to sharing the contents that they like. However, a considerable amount of effort is required to produce a correct transcription, and the user contribution would be limited to those contents that have enthusiastic listeners. If users are only asked about the recognition quality rather than the corrections of the errors in the transcriptions, and the system could utilize the scalar feedback to update the model by reinforcement learning, it would greatly reduce the effort required by the users. By reducing the effort of users, larger applications would become possible.

Reinforcement learning is based on the common sense idea that if an action is followed by an improvement in the state of affairs, then the tendency to produce that action is strengthened [7]. The two major formalizations of reinforcement learning are value-based methods including Q-learning approaches [8, 9, 10], and policy-based methods including policy gradient methods [11, 12]. In this paper, we first formulate a very general reinforcement learning framework for ASR systems based on the policy gradient method. Then, we propose a reinforcement learning method following the framework, where the feedback is based on hypothesis selection by the users.

The remainder of this paper is organized as follows. We first briefly review the application of reinforcement learning in speech information processing in Section 2, and the policy gradient method in Section 3. We then explain our proposed method in Section 4 and our implementation for experiments in Section 5. The experimental setup is described in Section 6, and the results are shown in Section 7. Finally, the conclusions are presented in Section 8.

2. RELATED WORK

There have been many studies that apply reinforcement learning to speech dialogue systems to improve dialogue control [13, 14, 15]. For source enhancement, Koizumi et al. have proposed a Q-learning-based method for a DNN-based system [16]. In their method, the speech enhancement performance was improved based on feedback from human evaluators about the perceptual quality of the enhanced speech. However, studies that apply reinforcement learning to ASR systems are limited, as noted in [17].

In studies on ASR, Nisida et al. have proposed a method that tunes an update coefficient τ of the MAP adaptation for GMM-HMM [18]. Their method used a confidence measure (CM) obtained from the result of Viterbi decoding of an utterance as the reward. Therefore, there was no human interaction. A small τ was used for speech segments with high confidence, and a large τ was used for segments with low confidence. Molina et al. have proposed a twopass decoding method that was also based on CM [17]. The idea was to reinforce the phone models in the second pass if they had a high confidence value, whereas they were weakened if they had low confidence. In the algorithm, the choice of the phone models in the decoding process is regarded as an action of reinforcement learning in a broad sense. The CM was estimated in the first pass, and it was used in the second pass by adding the value to the acoustic likelihood. The algorithm was for a decoding process, and the acoustic model was not updated. These methods were based on intuitive ideas

This work was supported by JSPS KAKENHI Grant Number 26280055 and 17K20001.

to modify the model update or decoding process based on CM. However, their connections to the major formalizations of reinforcement learning methods were not explained. In the same sense, the twopass unsupervised adaptation algorithms that reject low confidence hypotheses (e.g. [19]) may also be seen as a type of reinforcement learning.

3. POLICY GRADIENT METHOD

As the general setup for the policy gradient method-based reinforcement learning, a system has a set of actions and a policy function fthat takes a state s and returns a probability distribution $P_f(a|s)$ of an action a to take. The policy function is parameterized by a set of parameters θ . From $P_f(a|s)$, an action is sampled and executed. According to the action, the system gets a scalar reward $r_s(a)$.

The goal of the learning is to maximize the expected reward $\mathbb{E}[r_s(a)] = \sum_a P_f(a|s) r_s(a)$ with respect to θ . The maximization can be performed by applying the gradient ascent method. However, the key points here are that, while the reward $r_s(a)$ can be evaluated given the choice of the action, there may not exist an analytical functional form of the reward, and enumerating all possible actions may not be tractable. Therefore, we need a scheme to evaluate the gradient as follows, which is parallel to the derivation process of the natural evolution strategy using the *log*-trick [20, 21].

$$\nabla_{\boldsymbol{\theta}} \mathbb{E} \left[r_{\boldsymbol{s}} \left(\boldsymbol{a} \right) \middle| \boldsymbol{\theta} \right] = \nabla_{\boldsymbol{\theta}} \sum_{a} P_{f} \left(\boldsymbol{a} \middle| \boldsymbol{s} \right) r_{\boldsymbol{s}} \left(\boldsymbol{a} \right)$$
$$= \sum_{a} P_{f} \left(\boldsymbol{a} \middle| \boldsymbol{s} \right) \left(\frac{\nabla_{\boldsymbol{\theta}} P_{f} \left(\boldsymbol{a} \middle| \boldsymbol{s} \right)}{P_{f} \left(\boldsymbol{a} \middle| \boldsymbol{s} \right)} \right) r_{\boldsymbol{s}} \left(\boldsymbol{a} \right)$$
$$= \mathbb{E} \left[r_{\boldsymbol{s}} \left(\boldsymbol{a} \right) \nabla_{\boldsymbol{\theta}} \log P_{f} \left(\boldsymbol{a} \middle| \boldsymbol{s} \right) \right]. \tag{1}$$

Equation (1) means that $r_s(a) \nabla_{\theta} \log P_f(a|s)$ is an unbiased estimator of the gradient $\nabla_{\theta} \mathbb{E}[r_s(a)|\theta]$. Given the estimate of the gradient, the parameter update formula is obtained as follows.

$$\hat{\boldsymbol{\theta}} = \boldsymbol{\theta} + \epsilon r_{\boldsymbol{s}} \left(a \right) \nabla_{\boldsymbol{\theta}} \log P_f \left(a | \boldsymbol{s} \right), \tag{2}$$

where ϵ (> 0) is the learning rate. The same formulation holds when the reward is a conditional probability of r given a.

4. PROPOSED METHOD

We assume a situation where an ASR system is used to serve a vast number of general users over the Internet. The users input speech data that they want to transcribe. Such data would include recordings of school lectures, invited talks, presentations, and meetings. More interactive applications, such as voice input for email, can also be the target. The users want a reasonably good transcript quickly and easily, and they do not have time to correct all the recognition errors word by word. The user interface is equipped with a mechanism that allows the users to provide a scalar evaluation score for the recognition result as user feedback. There are several design choices about what types of scores we expect the users to provide intentionally or unintentionally, but we assume that it is given in an utterance basis.

To formulate a reinforcement learning framework for statistical ASR systems, we regard the whole system as a policy function that takes a feature sequence of an utterance as the input s and returns a probability distribution of a word sequence l of recognition hypothesis as an action. In particular, when the recognition system is based on an acoustic model $P_{AM}(s|l)$ and a language model $P_{LM}(l)$, the (unnormalized) probability distribution is given by Equation (3).

$$P(\boldsymbol{l}|\boldsymbol{s}) = \frac{P_{AM}(\boldsymbol{s}|\boldsymbol{l}) P_{LM}(\boldsymbol{l})}{P(\boldsymbol{s})} \propto P_{AM}(\boldsymbol{s}|\boldsymbol{l}) P_{LM}(\boldsymbol{l}). \quad (3)$$

If we further assume that we only want to update the acoustic model and it is a DNN-HMM, and we only want to update the DNN parameters θ to better predict the posterior probability of HMM states, then the gradient in Equation (2) becomes independent of the language model. Moreover, it is further decomposed to each time frame, and becomes:

$$r_{s}(\boldsymbol{l}) \frac{\partial \log P_{AM}(\boldsymbol{l}_{t}|\boldsymbol{s}_{t})}{\partial \boldsymbol{\theta}},$$
 (4)

where s_t is an acoustic feature vector at time frame t, and l_t is the HMM state aligned to that frame. Equation (4) indicates that the update formula for the reinforcement learning of DNN-HMM using the policy gradient method is simply a reward weighted version of normal cross-entropy based back-propagation. The update formula satisfies the criterion of the REINFORCE algorithm having the form shown in Equation (5) [22],

$$(r-b)\frac{\partial \log g(i)}{\partial \theta},$$
 (5)

where r is the reward, b is the reinforcement baseline, g is the probability function over the item i, and θ is the parameters of the neural network.

If we use CM as reward and round it to a binary value of 1.0 and 0.0, we can now clearly state that the conventional unsupervised adaptation with the hypothesis rejection mechanism mentioned in Section 2 is an example of the policy gradient-based reinforcement learning if the hypothesis is obtained by the sampling.

To utilize human feedback, the most direct measure of the recognition performance is the word accuracy. However, asking general users to evaluate word accuracy would not be realistic. Even for users with a technical background in ASR, it is time consuming to calculate. To avoid this problem, we propose a hypothesis selection-based reinforcement learning method in which we prepare two recognition systems. One system is the subject for the reinforcement learning, and the other is used as a rival. For each input utterance, a recognition hypothesis is sampled from each of the systems, and both of them are presented to the user. Then, the user selects the better hypothesis among them. In this case, the selection itself is the feedback to the system: 1 is the feedback when the hypothesis of the first system is selected, and 0 is the feedback otherwise. Based on the binary reward r, we update the DNN using the weighted gradient defined in Equation (6).

$$(1+\alpha)\left(r-\frac{\alpha}{1+\alpha}\right)\frac{\partial\log P_{AM}\left(\boldsymbol{l}_{t}|\boldsymbol{s}_{t}\right)}{\partial\boldsymbol{\theta}},$$
 (6)

where α ($0 \le \alpha \le 1$) is a scalar constant. The coefficient ($1 + \alpha$) is constant and can be seen as a part of the learning rate. Choosing $\alpha = 0$ corresponds to updating the parameters only when the hypothesis is selected.

5. IMPLEMENTATION WITH APPROXIMATIONS

To implement the proposed hypothesis selection-based reinforcement learning, we made some approximations in our experiments. First, we used a Viterbi decoding as in normal ASR systems to find the best hypothesis rather than sampling a hypothesis from the posterior distribution. Second, instead of preparing a separate rival system, we used the *n*-th best hypothesis (1 < n) of the same system as the rival hypothesis, where *n* is a constant. We refer to the best hypothesis as the Candidate 1 hypothesis $(l^{(1)})$ and the rival hypothesis as the Candidate 2 hypothesis $(l^{(2)})$. Since both of the hypotheses come from the same model, we used both of them in a symmetric

Training set	labeled	10 hours	
	unlabeled	50 + 50 + 50 + 50 hours	
Evaluation set		2 hours	
Vocabulary size		72k words	
Hypothes	is 1 Hypothesis	1 Hypothesis 1	Hypothesis 1
Hypothesis 2 Hypothesis		2 Hypothesis 2	Hypothesis 2
->: Decoding			
: Iraining			
Unlabeled		Unlabeled	Unlabeled
batch 1 batch 2		batch 3	batch 4
Labeled		Ť	Ť
data			
L Initial L RL1 RL2 RL3			
Supervised model			\checkmark
learning Stage 0	Stage 1	Stage 2	Stage 3

 Table 1. CSJ data used for the experiments.

Fig. 1. Reinforcement learning process. RLk indicates a model made by applying reinforcement learning k times. RLk is made at stage k and used to decode large batch #k+1

manner in the gradient update as shown in Equation (7).

$$(1+\alpha)\left(r-\frac{\alpha}{1+\alpha}\right)\frac{\partial\log P_{AM}\left(\boldsymbol{l}_{t}^{(1)}|\boldsymbol{s}_{t}\right)}{\partial\boldsymbol{\theta}} + (1+\alpha)\left((-r)-\frac{-1}{1+\alpha}\right)\frac{\partial\log P_{AM}\left(\boldsymbol{l}_{t}^{(2)}|\boldsymbol{s}_{t}\right)}{\partial\boldsymbol{\theta}}.$$
 (7)

This corresponds to collecting two feedbacks for two actions at the same time. For example, assuming $\alpha = 0$, we compute the gradient using the Candidate 1 hypothesis when it is selected (i.e. $\alpha = 1$) with weight 1, and we use the Candidate 2 hypothesis with weight 1 otherwise. Third, the parameter update by the reinforcement learning was performed based on large batches rather than an utterance by utterance update. This is mainly for the purpose of quick implementation.

For a more rigorous implementation of the sampling from the unnormalized posterior, beam sampling could be used [23]. Another strategy of preparing a rival system would be to use the same system from a randomly selected previous stage of update, as in AlphaGo [24]. By rewriting, Equation (7) becomes Equation (8). In this form, it can be seen that the hypothesis selection method is similar to discriminative training [25] in that it tries to increase the difference of the likelihood of the selected hypothesis (corresponding to correct the hypothesis) and the other hypothesis (the denominator lattice). However, the selected hypothesis is not a reference and usually contains errors, and it is within the formulation of the expected reward.

$$\begin{cases}
\frac{\partial \log P_{AM}(l_t^{(1)}|s_t)}{\partial \theta} - \alpha \frac{\partial \log P_{AM}(l_t^{(2)}|s_t)}{\partial \theta} & (r=1) \\
\frac{\partial \log P_{AM}(l_t^{(2)}|s_t)}{\partial \theta} - \alpha \frac{\partial \log P_{AM}(l_t^{(1)}|s_t)}{\partial \theta} & (r=0).
\end{cases}$$
(8)

6. EXPERIMENTAL SETUP

4

We performed the experiments using data from the Corpus of Spontaneous Japanese (CSJ) [26], and based on the CSJ recipe¹ in the Kaldi speech recognition toolkit [27]. In our experiments, we made two subsets from the original CSJ training data. The first subset contained 10 hours of data, and it was used as a labeled training set



Fig. 2. Number of stages and WERs of the large batch data. At stage k, the RLk model is used to decode large batch #k+1. When supervised training was performed, the WER at stage 3 was 19.3%.



Fig. 3. Number of stages and WERs of the evaluation set.

to train an initial baseline system. The other subset had 200 hours in total, and it was further divided into four subsets, each of which contained 50 hours of data. These four subsets were used as the unlabeled large batches for the reinforcement learning assuming the corresponding transcripts were not given to the system. Additionally, the standard evaluation set of CSJ, including two hours of speech data was used to evaluate the updated models using the same data set. Table 1 summarizes these data sets. The feedback from the users was simulated by evaluating the word error rates (WER) of the hypotheses from the system using the reference labels, and then performing the hypothesis selection based on the true WER. To simulate selection errors caused by users, experiments that introduced random swapping of the selected and unselected hypotheses were performed.

The input acoustic features for the DNN were 40 dimensional fMLLR features. The fMLLR features were computed using lattices, where the lattices were made by forced aligning the true labels for the training set, and by decoding the speech data for the large batches and for the evaluation set. The size of the input layer of the DNN was 1400 (spliced by +/- 17 frames). The DNNs had 6 hidden layers with a sigmoid activation function. They had 1905 units per hidden layer and 812 units for the output softmax layer.

The DNN-HMM of the baseline system was trained by pretraining and fine-tuning using the 10-hour labeled training data. For the large batch based reinforcement learning, the initial learning rates for the batches were set to 0.004, 0.002, 0.001 and 0.0005 for stages 1, 2, 3, and 4, respectively. The 10-hour labeled training data was always used by mixing it with the unlabeled large batches. The learning rate controls for the training data set and for the large batches were based on cross-validation using 10% of the labeled trained data as the held out set. The learning rate was halved when the improvement in cross-entropy on a cross-validation set fell below 1% in an epoch. The upper limit of the number of iterations in each epoch was set to 7.

Figure 1 shows the outline of the reinforcement learning process. The unlabeled large batch #1 was decoded using the initial baseline

¹https://github.com/kaldi-asr/kaldi/tree/master/egs/csj



Fig. 4. Relationship between hypotheses selection error rate and WER of the selected hypotheses.



Fig. 5. Number of stages and WERs of the large batches when there is 15% hypotheses selection error.

DNN-HMM model. The Candidate 1 hypotheses were the best results in the N-best list, and the Candidate 2 hypotheses were either 10th or fifth results in the list. The N-best list was created from a decoded lattice. After the first updated model (RL1) was made using large batch #1, it was used to recognize large batch #2. Based on the recognition results, the model was updated, making the next model (RL2). This process was repeated for all the large batches. For comparison purposes, unsupervised adaptation was performed, where the model was updated using the Candidate 1 hypothesis with and without the CM based hypothesis rejection. We used the averaged word posterior as CM obtained from a confusion network [28], and set the rejection threshold to keep 75% of the hypotheses based on a preliminary experiment.

7. RESULTS

Figure 2 shows the WER of the successively updated models based on unsupervised adaptation and reinforcement learning using the large batches sequentially. At stage 0, the initial baseline model is used to decode large batch #1. The hypothesis selection is performed for the model update, and the WERs are evaluated based on all the 1-best result. Therefore, differences of the WERs arise from Stage 1. In the figure, "initial model" indicates the WERs of the large batches using the baseline initial model. The unsupervised adaptations gave better results than the non-updated initial model. A slight improvement was obtained by using the confidence based hypothesis rejection compared to using all the hypotheses. As the lower-bound of the WER based on supervised adaptation using manual transcripts, 19.3% was obtained at state 3. For the reinforcement learning, 10th-best results were used as the Candidate 2 hypotheses. By using reinforcement learning, a larger improvement than the unsupervised adaptations was obtained when the coefficient α was chosen from 0.0 to 0.5. Choosing α greater than 0 means both of the hypotheses were used. The lowest WER was obtained when α was 0.5. When α was larger than 0.5, the second hypothesis affected the gradient too much and WER greatly increased. At stage 3, WER



Fig. 6. Number of stages and WERs of the large batches when the 5th and 10th best hypotheses were used as the Candidate 2 results. 15% selection error rate is simulated.

slightly increased except when $\alpha = 0.5$, including the unsupervised adaptation. This was partly because our learning rate reducing strategy was not optimal, and partly because the fourth batch simply contained relatively difficult utterances to recognize, as it is seen that the WER using the initial model was also higher compared to the other large batches.

To evaluate the updated models using the same data set, Figure 3 shows WERs of the common evaluation set. The WER by the initial baseline model was 26.4%, and the unsupervised adaptation gave 0.4% absolute improvement at the 4th stage when no CM was used and 0.6% when CM was used. Consistent improvement was observed by the reinforcement learning with $\alpha = 0.5$, and it gave the lowest WER of 25.5% at the 4th stage.

Figure 4 shows the simulated results of the relation between the selection error rate by the users and the WERs of the selected hypotheses. When the selection error rate is equal to or lower than 20%, we can expect lower WER in the selected hypotheses than the Candidate 1 hypotheses. Based on this analysis, we next investigated the performance of the reinforcement learning when there were 15% errors in the hypotheses selection. Figure 5 shows the WERs. The WER of stage 0 is the same as that of the Figure 2. It is confirmed that the reinforcement learning with $\alpha = 0.5$ still outperformed the unsupervised adaptation. At the 3rd stage, a slight increase in WER was observed both for the unsupervised adaptation and the reinforcement learning due to the same reason as before.

Finally, we have evaluated the performance of the reinforcement learning when the 5th-best results were used as the Candidate 2 hypotheses instead of the 10th-best results. Figure 6 shows the WERs with 15% selection errors. For reinforcement learning, α was set to 0.5. While the improvement became small, reinforcement learning still gave better results than the unsupervised adaptation.

8. CONCLUSION

In this paper, we have proposed a policy gradient-based reinforcement learning framework for ASR systems, and also have proposed a hypothesis selecting-based reinforcement learning method as a particular instance of the framework. In the experiments, we have shown that the proposed method reduces WER compared to the unsupervised adaptations. The tendencies were the same when 15% of simulated noise in the hypothesis selection was introduced, while the improvement became slightly smaller. When the number of stages was increased, there was a tendency for the WER to increase in both the unsupervised adaptation and the reinforcement learning in several cases. Future work includes addressing the problem of overtraining by adjusting the strategy for the learning rate and the number of iterations in each stage, and improving the performance by investigating more effective ways to update the model.

9. REFERENCES

- J. Ogata, M. Goto, and K. Eto, "Automatic transcription for a web 2.0 service to search podcasts," in *Proc. Interspeech*, 2007, pp. 2617–2620.
- [2] M. Goto, J. Ogata, and K. Eto, "Podcastle: A web 2.0 approach to speech recognition research," in *Proc. Interspeech*, 2007, pp. 2397–2400.
- [3] C. J. Leggetter and P. C. Woodland, "Maximum likelihood linear regression for speaker adaptation of continuous density hidden Markov models," *Computer Speech and Language*, vol. 9, pp. 171–185, 1995.
- [4] J.-L. Gauvain and C.-H. Lee, "Maximum a posteriori estimation for multivariate Gaussian mixture observations of Markov chains," *IEEE Transactions on Speech and Audio Processing*, vol. 2, pp. 291–298, 1994.
- [5] S. Mirsamadi and J. Hansen, "A study on deep neural network acoustic model adaptation for robust far-field speech recognition," in *Proc. Interspeech*, 2015, pp. 2430–2434.
- [6] K. Yao, D. Yu, F. Seide, H. Su, L. Deng, and Y. Gong, "Adaptation of context-dependent deep neural networks for automatic speech recognition," in *Proc. IEEE Workshop on Spoken Language Technology (SLT)*, 2012.
- [7] R. S. Sutton, A. G. Barto, and R. J. Williams, "Reinforcement learning is direct adaptive optimal control," *IEEE Control Systems*, vol. 12, no. 2, pp. 19–22, April 1992.
- [8] Christopher J.C.H. Watkins and Peter Dayan, "Technical note: Q-learning," *Machine Learning*, vol. 8, no. 3, pp. 279–292, May 1992.
- [9] G. A. Rummery and M. Niranjan, "On-line q-learning using connectionist systems," Tech. Rep., 1994.
- [10] Volodymyr Mnih, Koray Kavukcuoglu, David Silver, Alex Graves, Ioannis Antonoglou, Daan Wierstra, and Martin Riedmiller, "Playing atari with deep reinforcement learning," in *NIPS Deep Learning Workshop*. 2013.
- [11] Richard S. Sutton, David McAllester, Satinder Singh, and Yishay Mansour, "Policy gradient methods for reinforcement learning with function approximation," in *Proceedings of the* 12th International Conference on Neural Information Processing Systems, 1999, NIPS'99, pp. 1057–1063.
- [12] Volodymyr Mnih, Adria Puigdomenech Badia, Mehdi Mirza, Alex Graves, Timothy Lillicrap, Tim Harley, David Silver, and Koray Kavukcuoglu, "Asynchronous methods for deep reinforcement learning," in *Proceedings of The 33rd International Conference on Machine Learning*, Maria Florina Balcan and Kilian Q. Weinberger, Eds. 2016, vol. 48 of *Proceedings of Machine Learning Research*, pp. 1928–1937, PMLR.
- [13] D. Lu, T. Nishimoto, and N. Minematsu, "Decision of response timing for incremental speech recognition with reinforcement learning," in 2011 IEEE Workshop on Automatic Speech Recognition Understanding, Dec 2011, pp. 467–472.
- [14] P. H. Su, D. Vandyke, M. Gasic, D. Kim, N. Mrksic, T. H. Wen, and S. J. Young, "Learning from real users: rating dialogue success with neural networks for reinforcement learning in spoken dialogue systems.," in *Proc. INTERSPEECH*, 2015, pp. 2007–2011.

- [15] F. Wang, "A multi-agent reinforcement learning algorithm for disambiguation in a spoken dialogue system," in *Proceedings* of the 2010 International Conference on Technologies and Applications of Artificial Intelligence. 2010, TAAI '10, pp. 116– 123, IEEE Computer Society.
- [16] Y. Koizumi, K. Niwa, Y. Hioka, K. Kobayashi, and Y. Haneda, "DNN-based source enhancement self-optimized by reinforcement learning using sound quality measurements," in *Proc. ICASSP*, March 2017, pp. 81–85.
- [17] C. Molina, N. B. Yoma, F. Huenupan, C. Garreton, and J. Wuth, "Maximum entropy-based reinforcement learning using a confidence measure in speech recognition for telephone speech," *IEEE Transactions on Audio, Speech, and Language Processing*, vol. 18, no. 5, pp. 1041–1052, July 2010.
- [18] M. Nishida, Y. Mamiya, Y. Horiuchi, and A. Ichikawa, "Online incremental adaptation based on reinforcement learning for robust speech recognition," in *Proc. Interspeech*, 2004, pp. 1985–1988.
- [19] D. Charlet, "Confidence-measure-driven unsupervised incremental adaptation for HMM-based speech recognition," in *Proc. ICASSP*, 2001, vol. 1, pp. 357–360 vol.1.
- [20] D. Wierstra, T. Schaul, T. Glasmachers, Y. Sun, J. Peters, and J. Schmidhuber, "Natural evolution strategies," *J. Mach. Learn. Res.*, vol. 15, no. 1, pp. 949–980, 2014.
- [21] N. Hansen, S. D. Müller, and P. Koumoutsakos, "Reducing the time complexity of the derandomized evolution strategy with covariance matrix adaptation (CMA-ES)," *Evolutionary Computation*, vol. 11, no. 1, pp. 1–18, 2003.
- [22] R. J. Williams, "Simple statistical gradient-following algorithms for connectionist reinforcement learning," *Machine Learning*, vol. 8, no. 3, pp. 229–256, May 1992.
- [23] Jurgen Van Gael, Yunus Saatci, Yee Whye Teh, and Zoubin Ghahramani, "Beam sampling for the infinite hidden Markov model," in *Proceedings of the 25th International Conference* on Machine Learning. 2008, ICML '08, pp. 1088–1095, ACM.
- [24] D. Silver, A. Huang, C. J. Maddison, A. Guez, L. Sifre, G. van den Driessche, J. Schrittwieser, I. Antonoglou, V. Panneershelvam, M. Lanctot, S. Dieleman, D. Grewe, J. Nham, N. Kalchbrenner, I. Sutskever, T. Lillicrap, M. Leach, K. Kavukcuoglu, T. Graepel, and D. Hassabis, "Mastering the game of Go with deep neural networks and tree search," *Nature*, vol. 529, no. 7587, pp. 484–489, Jan. 2016.
- [25] K. Vesely, A. Ghoshal, L. Burget, and D. Povey, "Sequencediscriminative training of deep neural networks," in *Proc. Interspeech*, 2013, pp. 2345–2349.
- [26] S. Furui, K. Maekawa, and H. Isahara, "A Japanese national project on spontaneous speech corpus and processing technology," in *Proc. ASR'00*, 2000, pp. 244–248.
- [27] D. Povey, A. Ghoshal, G. Boulianne, L. Burget, O. Glembek, N. Goel, M. Hannemann, P. Motlicek, Y. Qian, P. Schwarz, J. Silovski, G. Stemmer, and K. Veseli, "The Kaldi speech recognition toolkit," in *Proc. ASRU*, 2011.
- [28] Karel Vesely, Mirko Hannemann, and Lukas Burget, "Semisupervised training of deep neural networks," in Automatic Speech Recognition and Understanding (ASRU), 2013 IEEE Workshop on. IEEE, 2013, pp. 267–272.