# TO REVERSE THE GRADIENT OR NOT: AN EMPIRICAL COMPARISON OF ADVERSARIAL AND MULTI-TASK LEARNING IN SPEECH RECOGNITION

\*Yossi Adi<sup>1,2</sup>, \*Neil Zeghidour<sup>2,3</sup>, Ronan Collobert<sup>2</sup>, Nicolas Usunier<sup>2</sup>, Vitaliy Liptchinsky<sup>2</sup>, Gabriel Synnaeve<sup>2</sup>

<sup>1</sup>Bar-Ilan University. <sup>2</sup>Facebook AI Research. <sup>3</sup>CoML, ENS/INRIA.

## ABSTRACT

Transcribed datasets typically contain speaker identity for each instance in the data. We investigate two ways to incorporate this information during training: Multi-Task Learning and Adversarial Learning. In multi-task learning, the goal is speaker prediction; we expect a performance improvement with this joint training if the two tasks of speech recognition and speaker recognition share a common set of underlying features. In contrast, adversarial learning is a means to learn representations invariant to the speaker. We then expect better performance if this learnt invariance helps generalizing to new speakers. While the two approaches seem natural in the context of speech recognition, they are incompatible because they correspond to opposite gradients back-propagated to the model. In order to better understand the effect of these approaches in terms of error rates, we compare both strategies in controlled settings. Moreover, we explore the use of additional un-transcribed data in a semi-supervised, adversarial learning manner to improve error rates. Our results show that deep models trained on big datasets already develop invariant representations to speakers without any auxiliary loss. When considering adversarial learning and multi-task learning, the impact on the acoustic model seems minor. However, models trained in a semi-supervised manner can improve error-rates. Index Terms: automatic speech recognition, adversarial learning, multi-task learning, neural networks

## 1. INTRODUCTION

The minimal components of a speech recognition dataset are audio recordings and their corresponding transcription. They also typically contain an additional information, which is the anonymous identifier of the speaker corresponding to each utterance. As speaker variations are one of the most challenging aspects of speech processing, this information can be leveraged to improve the performance and robustness of speech recognition systems. Traditionally, the identity of the speakers has been used to extract speaker representations and use it as an additional input to the acoustic model [1, 2]. This approach requires an additional feature extraction process of the input signal. Recently, two approaches have been proposed to leverage speaker information as another supervision to the acoustic model; Multi-Task Learning (MT) [3, 4] and Adversarial Learning (AL) [5]. In the context of speech recognition, a way of performing MT or AL is to add a speaker classification branch in parallel of the main branch trained for transcription. All layers below the fork of the two branches receive gradients from both transcription loss and speaker loss.

In MT our goal is to jointly transcribe the speech signal together with classifying the speaker identity. Previous work has explored using auxiliary tasks such as gender or context [6], as well as speaker classification [7, 8]. For example, the authors in [7] propose to classify the speaker identity in addition to estimating the phoneme-state posterior probabilities used for ASR, they presented a non-negligible improvement in terms of *Phoneme Error Rate* following the above approach.

An opposite approach is to assume that good acoustic representations should be invariant to any task that is not the speech recognition task, in particular the speaker characteristics. A method to learn such invariances is AL: the branch of the speaker classification task is trained to reduce its classification error, however the main branch is trained to maximize the loss of this speaker classifier. By doing so, it learns invariance to speaker characteristics. This approach has been previously used for speech recognition to learn invariance to noise conditions [9], speaker identity [10, 11, 12] and accent [13]. The authors in [11], proposed to minimize the speaker classification loss, and simultaneously maximize the speaker classification loss. This approach achieves a significant improvement in terms of *Word Error Rate*.

Our study arises from two observations. First, both the MT and AL approaches described above have been used with the same purpose despite being fundamentally opposed, and this raises the question of which one is the most appropriate choice to improve the performance of speech recognition

Work conducted while Yossi Adi was an Intern at Facebook AI Research. \*Equal contribution

systems. Secondly, they only differ by a simple computational step which consists in reversing the gradient at the fork between the speaker branch and the main branch. This allows for controlled experiments where both approaches can be compared with equivalent architectures and number of parameters.

In this work we focus on letter based acoustic models. In this setting, we are given audio recordings and their transcripts, and train a neural network to output letters with a sequence-based criterion. During training, we also have access to the identity of the speaker of each utterance. We perform a systematic comparison between MT and AL for the task of large vocabulary speech recognition on the Wall Street Journal dataset (WSJ) [14]. Knowing whether the speaker information should be used to impact low level or high level representations is also a valuable information, so for each approach we experiment with three levels to fork the speaker branch to: lower layer, middle and upper layer.

**Our Contribution:** (i) We observed that deep models trained on big datasets, already develop invariant representations to speakers with neither AL nor MT; (ii) Both AL and MT do not have a clear impact on the acoustic model with respect to error rates; and (iii) Using additional, un-transcribed speaker labeled data, (i.e. the speaker is known, but without the transcription), seems promising. It improves Letter-Error Rates (LER), even though these improvements did not transfer into Word-Error Rates (WER) improvements.

## 2. ADVERSARIAL VS. MULTITASK

Given a training set of n transcribed acoustic utterances coupled with their speaker identity, our goal is to utilize these speaker labels to improve transcription loss.

Formally, let  $S = (X_i, Y_i, s_i)_{i=0}^n$  be the training set, in which each example is composed of three elements: a sequence of m acoustic features  $X_i = \{x_1, \dots, x_m\}$  where each  $x_j$  is a d-dimensional vector, a sequence of k characters  $Y_i = \{y_1, \dots, y_k\}$  which are not aligned with  $X_i$ , and a speaker label  $s_i$ . In the following subsections we present two common approaches to add the speaker identity as another supervision to the network.

### 2.1. Adversarial Learning

As speech recognition tasks should not depend on a specific set of speakers, in AL, we would like to learn an acoustic representation which is *speaker-invariant* and at the same time *characters-discriminative*. For that purpose, we consider the first k layers of the network as an encoder, denoted by  $E_r$  with parameters  $\theta_r$ , which maps the acoustic features  $X_i$  to  $R_i =$  $\{r_1, \dots, r_l\}$  where l can be smaller than m. Then, the representations  $R_i$  are fed into a decoder network  $D_y$  with parameters  $\theta_y$ , to output the characters posteriors  $p(Y_i|X_i; \theta_r, \theta_y)$ . Furthermore, we introduce another decoder network to classify the speaker, denote  $D_s$ . The above decoder, maps  $R_i$  to the speaker posteriors  $p(s_i|X_i; \theta_r, \theta_y)$ .

In order to make the representation  $\mathbf{R}$  speaker-invariant, we train  $E_r$  and  $D_s$  jointly using adversarial loss, where we optimize  $\theta_r$  to maximize speaker classification loss, and at the same time optimizing  $\theta_s$  to minimize the speaker classification loss. Recall, we would like to make the representation  $\mathbf{R}$ characters-discriminative, hence, we further optimize  $\theta_y$  and  $\theta_r$  to minimize the transcription loss. Therefore, the total loss is constructed as follows,

$$\mathcal{L}(\theta_r, \theta_y, \theta_s) = \mathcal{L}_{acoustic}(\theta_r, \theta_y) - \lambda \mathcal{L}_{spk}(\theta_r, \theta_s) \qquad (1)$$

where  $\mathcal{L}_{acoustic}(\theta_r, \theta_y)$  is the transcription loss,  $\mathcal{L}_{spk}(\theta_r, \theta_s)$  is the speaker loss, and  $\lambda$  is a trade-off parameter which controls the balance between the two. This minimax game will enhance the discriminative properties of  $D_s$  and  $D_y$  while at the same time push  $E_r$  towards generating speaker-invariant representation.

As suggested in [5, 11], we optimize the parameters using Stochastic Gradient Descent where we reverse the gradients of  $\partial \mathcal{L}_{spk}/\partial \theta_r$  [5].

## 2.2. Multi-Task Learning

Similarly to AL, in MT we consider an encoder  $E_r$  with parameters  $\theta_r$  and two decoders  $D_y$  and  $D_s$  with parameters  $\theta_y$  and  $\theta_s$  respectively. However, in MT, our goal is to minimize both  $\mathcal{L}_{acoustic}(\theta_r, \theta_y)$  and  $\mathcal{L}_{spk}(\theta_r, \theta_s)$ . Therefore, the objective in MT case can be expressed as follows,

$$\mathcal{L}(\theta_r, \theta_y, \theta_s) = \mathcal{L}_{acoustic}(\theta_r, \theta_y) + \lambda \cdot \mathcal{L}_{spk}(\theta_r, \theta_s) \quad (2)$$

Practically, on forward propagation, both AL and MT act the same way. Yet, on back-propagation they act exactly the opposite. In AL, we reverse the gradients from  $D_s$  before back-propagating into  $E_r$ . In contrast, in MT we sum together the gradients from  $D_s$  and  $D_y$ . Figure 1 depicts an illustration of the described model.

#### 3. MODELS

Our acoustic models are based on Gated Convolutional Neural Networks (Gated ConvNets) [15], fed with 40 log-mel filterbank energies extracted every 10 ms with a 25 ms sliding window. Gated ConvNets stack 1D convolutions with Gated Linear Units (GLUs). 1D ConvNets were introduced early in the speech community, and are also referred as Time-Delay Neural Networks [16].

We set  $\mathcal{L}_{acoustic}(\theta_r, \theta_y)$  to be the AutoSeg Criterion (ASG) criterion [17, 18], which is similar to the Connectionist Temporal Classification (CTC) criterion [19].

Although speech utterances contain time dependent transcriptions, the speaker labels should be applied to the whole



**Fig. 1**. An illustration of our architecture. We first feed the network a sequence of acoustic features, then we take the output of some intermediate representation an use it also to classify the speaker.

sequence. Hence, when optimizing for multi-task or adversarial learning we add a speaker branch which aggregates over all the time frames to classify the speaker. This pooling operation aggregates the time-dependent representations into a single sequence-level representation  $s = g(\cdot)$ . Two straightforward aggregation methods are to take the sum over all frames:  $g(\cdot) = \sum_t r_t$  where  $r_t$  is the representation at time frame t, or the max  $g(\cdot) = \max_t(r_t)$ , where the max is taken over each dimension. We propose to use a trade-off solution between these two methods, which is the LogSumExp [20],  $s = \frac{1}{\tau} \log \left(\frac{1}{T} \sum_t e^{\tau \cdot r_t}\right)$ , where  $\tau$  is a hyper-parameter controls the trade-off between the sum and the max; with high values of  $\tau$  the LogSumExp is similar to the max, while it tends to a sum as  $\tau \to 0$ . In practice, we used  $\tau = 1$ .

We set  $L_{speaker}(\theta_r, \theta_s)$  to be the Negative Log Likelihood (NLL) loss function. The overall network, together with  $D_y$ , is trained by back-propagation. At inference time,  $D_s$ , the speaker classification branch is discarded.

# 4. EXPERIMENTAL RESULTS

In the following section we describe our experimental results. First, we provide all the technical details and setups in Subsection 4.1. Then, we describe our baselines and analyze their ability to encode information about the speaker in Subsection 4.2. Lastly, on Subsection 4.3 we present our results on the WSJ dataset [14].

## 4.1. Setups

All models were composed of 17 gated convolutional layers followed by a weight normalization [21]. We used a dropout layer after every GLU layer with dropout rate of 0.25. We optimized the models using SGD with two different learning rate values, we used learning rate of 1.4 for  $E_r$  and  $D_y$  and learning rate of 0.1 for  $D_s$ . We did not use momentum or weight decay. All the models were trained using WSJ dataset which contains 283 speakers with 81.5 hours of speaking time.

We computed WER using our own one-pass decoder, which performs a simple beam-search with beam thresholding, histogram pruning and language model smearing [22]. We did not implement any sort of model adaptation before decoding, nor any word graph rescoring. Our decoder relies on KenLM [23] for the language modeling part, using a 4gram LM trained on the standard data of WSJ [14]. For LER, we used Viterbi decoding as in [17].

The speaker branch,  $D_s$ , was composed of one gated convolutional layer of width 5 and 200 feature maps with weight normalization, followed by a linear layer. For MT models we used a static  $\lambda$  value of 0.5. For AL models, we slowly increased the  $\lambda$  from 0 to 0.2 using the following update scheme,  $\lambda_i = 2/(1 + e^{-p_i}) - 1$ , where  $p_i$  is a scaled version of the epoch number. The above technique was successfully explored in previous studies [24, 5]. In practice, we observed that reversing the gradients from the beginning of training can cause the network to diverge. To avoid that, we first trained the models without back-propagating the gradients from  $D_s$ to  $E_r$ , (i.e.  $\lambda = 0$ ). Then, we trained only  $D_s$ , and finally, we trained both of them jointly. We found the above procedure crucial for AL, however in MT the performance is equivalent to standard training. Thus, for comparison, we followed the above approach in both settings.

## 4.2. Representation Analysis

We would like to assess to what extent the network encodes the speaker identity? The motivation for answering this question is two-fold (i) to analyze the networks' behavior with respect to the speaker information; and (ii) having a measurement of how much speaker information is encoded in the representation; this can provide an intuitive insight into how much error rates can be improved using speaker-adversarial learning.

To tackle this question, we follow a similar approach to the one proposed in [25, 26]. We first trained a baseline network without the speaker classification branch, then we used it to extract representations of the speech signal from different layers in the network. Finally, we use these representations only, to train a model to classify the speaker identity. After training, we measure the model's accuracy to analyze the presence of speaker identity in the representations.

The basic premise of this approach is that if we cannot optimize a classifier to predict speaker identity based on representation from the model, then this property is not encoded in the representation, or rather, not encoded in an effective way, considered the way it will be used. Notice that we are not interested in improving speaker classification but rather to have a comparison between the models based on the clas-



**Fig. 2**. Speaker classification accuracy using the representation from different layers of the network. Results are reported for Baseline, MT and AL models using three different settings, after the second, eighth, and fifteenth layers.

sifier's classification accuracy.

For that purpose, we defined three settings, where we add the speaker branch after: (i) two convolutional layers, denote as IN, (ii) eight convolutional layers, denote as MID, and (iii) 15 convolutional layers, denote as OUT. For each setting we trained a classification model for 10 epochs, using the same architecture as the speaker classification branch.

Figure 2 presents the speaker classification accuracy for the baseline together with MT and AL using representations from IN , MID and OUT . Notice, we observed a similar behavior in AL, MT, and the baseline, where the classification gets harder when extracting representations from deeper layers in the network. This implies that the network already develops speaker invariant representations during training. Even though, adding additional speaker loss can push it further and improve speaker invariance.

# 4.3. WSJ Experiments

Following our experiments in subsection 4.2, we trained MT and AL in three versions, IN , MID and OUT . Table 1 summarizes the results of our models, the baseline, and state-of-the-art systems.

Using both AL and MT seems to have a minor effect on transcription modeling when training on WSJ dataset. Both approaches achieve equal or somewhat worse LER than the baseline on the development set, and improve over the test set, with AL performs slightly better than MT. When considering WER, both AL and MT improve results over the baseline on the test set, however none of them improve over the development set. These results are somewhat counter-intuitive since both approaches attain similar results, however can be considered as opposite paradigms. One explanation is that both methods act as another regularization term. In other words, although the two approaches push the gradients towards opposite directions, when treated carefully, they both impose an-

 Table 1. Letter Error Rates and Word Error Rates on the WSJ dataset.

Model		Nov93 dev		Nov92 eval	
Туре	Layer	LER	WER	LER	WER
Human [27]		-	8.1	-	5.0
BLSTMs [28]		-	6.6	-	3.5
convnet BLSTM w/ addi-		-	4.4	-	3.6
tional data [27]					
Our Baseline		7.2	9.7	4.9	5.9
Mult.	In	7.2	9.9	4.8	5.8
	Mid	7.3	9.8	4.9	5.9
	Out	7.3	10.1	4.8	5.7
Adv.	In	7.2	9.9	4.7	5.7
	Mid	7.3	10.0	4.8	5.8
	Out	7.2	9.8	4.7	5.6
Semi.	In	6.8	9.8	4.6	6.2
	Mid	6.8	9.7	4.7	6.0
	Out	6.9	10.0	4.6	6.3

other constraint on the models' parameters [29, 30, 31].

## 4.4. Adversarial Semi-Supervised Learning

Lastly, we wanted to investigate whether the observed improvements were limited by the relatively small amount of speakers (283) in the train set of WSJ. For that purpose, we ran a semi-supervised experiment were we control the amount of transcribed data by using only speaker labeled examples. We add  $\approx 35$  hours of speech using another 132 speakers, from the Librispeech corpus [32], and train the network in an adversarial fashion. For the new training examples we do not use the transcriptions, but only the speaker identity, hence optimizing the speaker loss only. Results are summarized on Table 1. Using additional speaker labeled data, seems beneficial in terms of LER: it improves over the baseline and previous AL results in all settings.

On the other hand, when considering WER, none of the settings improve over the baseline. This misalignment between the LER and WER results can be either because we constrained the model's outputs by LM at decoding time, hence WER results are greatly affected by the type of LM used, or because the LER improvements are on meaningless parts in the sequence. The analysis of this observations as well as jointly optimizing acoustic model and LM are left for future research.

## 5. CONCLUSION AND FUTURE WORK

In this paper, we studied AL and MT in the context of speech recognition. We showed that deep models already learn speaker-invariant representations, but can still benefit from semi-supervised speaker adversarial learning.

For future work, we would like to examine the effect of semi-supervised learning for low-resource languages.

#### 6. REFERENCES

- A. Senior and I. Lopez-Moreno, "Improving dnn speaker independence with i-vector inputs," in Acoustics, Speech and Signal Processing (ICASSP), 2014 IEEE International Conference on. IEEE, 2014, pp. 225–229.
- [2] V. Peddinti, G. Chen, D. Povey, and S. Khudanpur, "Reverberation robust acoustic modeling using i-vectors with time delay neural networks," in *INTERSPEECH*, 2015.
- [3] R. Caruana, "Multitask learning," in *Learning to learn*, pp. 95–133. Springer, 1998.
- [4] R. Collobert, J. Weston, L. Bottou, M. Karlen, K. Kavukcuoglu, and P. Kuksa, "Natural language processing (almost) from scratch," *Journal of Machine Learning Research*, vol. 12, pp. 2493–2537, 2011.
- [5] Y. Ganin et al., "Domain-adversarial training of neural networks," *The Journal of Machine Learning Research*, vol. 17, no. 1, pp. 2096–2030, 2016.
- [6] G. Pironkov, S. Dupont, and T. Dutoit, "Multi-task learning for speech recognition: an overview," in *Proc. of ESANN*, 2016.
- [7] G. Pironkov, S. Dupont, and T. Dutoit, "Speaker-aware long short-term memory multi-task learning for speech recognition," in *Proc. of EUSIPCO*, 2016.
- [8] Z. Tang, L. Li, and D. Wang, "Multi-task recurrent model for speech and speaker recognition," in *Proc. of APSIPA*, 2016.
- [9] Dmitriy S. et al., "Invariant representations for noisy speech recognition," *CoRR*, vol. abs/1612.01928, 2016.
- [10] Taira Tsuchiya, Naohiro Tawara, Testuji Ogawa, and Tetsunori Kobayashi, "Speaker invariant feature extraction for zeroresource languages with adversarial learning," in *Proc. of ICASSP*, 2018.
- [11] Z. Meng, J. Li, Z. Chen, Y. Zhao, V. Mazalov, Y. Gong, et al., "Speaker-invariant training via adversarial learning," *arXiv*:1804.00732, 2018.
- [12] G. Saon et al., "English conversational telephone speech recognition by humans and machines," arXiv:1703.02136, 2017.
- [13] Sining Sun, Ching-Feng Yeh, Mei-Yuh Hwang, Mari Ostendorf, and Lei Xie, "Domain adversarial training for accented speech recognition," *Proc. of ICASSP*, 2018.
- [14] D. B Paul and J. M Baker, "The design for the wall street journal-based csr corpus," in *Proc. of the workshop on Speech* and Natural Language, 1992.
- [15] Y. N Dauphin, Aa Fan, M. Auli, and D. Grangier, "Language modeling with gated convolutional networks," arXiv:1612.08083, 2016.
- [16] V. Peddinti, D. Povey, and S. Khudanpur, "A time delay neural network architecture for efficient modeling of long temporal contexts," in *INTERSPEECH*, 2015.
- [17] R. Collobert, C. Puhrsch, and G. Synnaeve, "Wav2letter: an end-to-end convnet-based speech recognition system," *arXiv*:1609.03193, 2016.
- [18] V. Liptchinsky, G. Synnaeve, and R. Collobert, "Letter-based speech recognition with gated convnets," *arXiv*:1712.09444, 2017.

- [19] A. Graves, S. Fernández, F. Gomez, and J. Schmidhuber, "Connectionist temporal classification: labelling unsegmented sequence data with recurrent neural networks," in *Proc. of ICML*, 2006.
- [20] D. Palaz, G. Synnaeve, and R. Collobert, "Jointly learning to locate and classify words using convolutional networks.," in *Proc. of INTERSPEECH*, 2016.
- [21] T. Salimans and D. Kingma, "Weight normalization: A simple reparameterization to accelerate training of deep neural networks," in *Proc. of NIPS*, 2016.
- [22] V. Steinbiss, B. Tran, and H. Ney, "Improvements in beam search," in *Proc. of ICSLP*, 1994.
- [23] K. Heafield, I. Pouzyrevsky, J. H Clark, and P. Koehn, "Scalable modified kneser-ney language model estimation," in *Proc.* of ACL, 2013.
- [24] G. Lample, N. Zeghidour, N. Usunier, A. Bordes, L. Denoyer, et al., "Fader networks: Manipulating images by sliding attributes," in *Proc. of NIPS*, 2017.
- [25] Y. Adi, E. Kermany, Y. Belinkov, O. Lavi, and Y. Goldberg, "Fine-grained analysis of sentence embeddings using auxiliary prediction tasks," *Proc. of ICLR*, 2016.
- [26] Y Adi, E Kermany, Y Belinkov, O Lavi, and Y Goldberg, "Analysis of sentence embedding models using prediction tasks in natural language processing," *IBM Journal of Research* and Development, vol. 61, no. 4, pp. 3–1, 2017.
- [27] Dario Amodei et al., "Deep speech 2 : End-to-end speech recognition in english and mandarin," in *Proc. of ICML*, 2016.
- [28] W. Chan and I. Lane, "Deep recurrent neural networks for acoustic modelling," arXiv:1504.01482, 2015.
- [29] S. Ruder, "An overview of multi-task learning in deep neural networks," arXiv:1706.05098, 2017.
- [30] M. Nski, L. Simon, and F. Jurie, "An adversarial regularisation for semi-supervised training of structured output neural networks," in *Proc. of NIPS*, 2017.
- [31] D. Cohen, B. Mitra, K. Hofmann, and W B. Croft, "Cross domain regularization for neural ranking models using adversarial learning," *arXiv preprint arXiv:1805.03403*, 2018.
- [32] V. Panayotov, G. Chen, D. Povey, and S. Khudanpur, "Librispeech: an asr corpus based on public domain audio books," in *Proc. of ICASSP*, 2015.