



O-1: Self-training with Oracle and 1-best Hypothesis

Murali Karthick Baskar, Andrew Rosenberg, Bhuvana Ramabhadran, Kartik Audhkhasi

Google Inc.,

{mkbaskar, rosenberg, bhuv}@google.com

Abstract

We introduce O-1, a new self-training objective to reduce training bias and unify training and evaluation metrics for speech recognition. O-1 is a faster variant of Expected Minimum Bayes Risk (EMBR), that boosts the oracle hypothesis and can accommodate both supervised and unsupervised data. We demonstrate the effectiveness of our approach in terms of recognition on publicly available SpeechStew datasets and a large-scale, in-house data set. On Speechstew, the O-1 objective closes the gap between the actual and oracle performance by 80% relative compared to EMBR which bridges the gap by 43% relative. O-1 achieves 13% to 25% relative improvement over EMBR on the various datasets that SpeechStew comprises of, and a 12% relative gap reduction with respect to the oracle WER over EMBR training on the in-house dataset. Overall, O-1 results in a 9% relative improvement in WER over EMBR, thereby speaking to the scalability of the proposed objective for large-scale datasets. **Index Terms:** Self-training, EMBR, O-1, ASR, speech recognition, discriminative training.

1. Introduction

The use of labeled and unlabeled data in training has led to a suite of co-training and self-training [1] methods in speech and language processing. The underlying principle in these algorithms is an iterative learning paradigm where a model learns from a ground-truth distribution as well as from its own predictions at the previous iteration with appropriate regularization [2]. These algorithms serve to mitigate three key problems. First, *exposure bias* in autoregressive models [3], where the model is conditioned only from the supervised training data (teacher forcing) [4] and thus never learns from its own predictions. This leads to early convergence but is mismatched with the inference step, where no ground-truth is available. Scheduled sampling [5] mitigates this issue by exposing the model to its past/own predictions at random instances. Second, *mismatches between training objectives and evaluation metrics* such as, maximum likelihood estimation (MLE) vs Word Error Rate (WER) [6, 3] which cause performance degradation of the model. While the former results in error accumulation during inference, the latter ignores the lattice of predictions which can be utilized to learn from unsupervised data. Thirdly, *label bias* in e2e systems occurs when the model cannot relearn from new labels and remains biased to past labels. This label bias is mostly subtle and is studied and addressed in [7] with the use of global normalization.

Prior work on Automated Speech Recognition (ASR) has shown that, Minimum Bayes risk (MBR) approach is capable of handling the exposure bias [8, 9], training-inference criterion mismatch [10, 11] and the label bias problem [7, 12]. MBR is a self-training paradigm (commonly referred to as discriminative

training) [13] and is either termed as expected MBR (EMBR) [14, 15, 16] when used to reduce expected word errors or state MBR (sMBR) when minimizing the state level expected errors [17, 18, 19, 20]. Despite the benefits of EMBR, it suffers from two main drawbacks:

- EMBR formulation does not directly impact the oracle hypothesis (best WER), but instead focusses on expected error reduction by boosting multiple hypotheses scores.
- The training duration is relatively high compared to MLE training as batch size gets scaled by the beam size.

In this work, we propose a novel self-training objective, referred to as O-1. It is a variant of EMBR formulated to discriminatively boost only the Oracle hypothesis from the 1-best hypothesis. The primary contributions of this work are:

- We introduce, the O-1 objective which allows the efficient use of both supervised data and unsupervised data. O-1 is integrated with hard distillation [21] for unsupervised training. To the best of our knowledge, O-1 is the first variant of EMBR which can perform unsupervised learning.
- Optimized training time (and computational costs) compared to EMBR, thereby paving the way to train for longer duration and more data.
- Effective self-training paradigm that uses WER to select the best sequence predicted by the model as the target label sequence.
- The proposed objective results in a significant reduction of the gap between 1-best WER and the oracle WER. Experiments conducted on the set of public corpora in Speechstew show consistent reduction in WER, by 13-25 % relative.

2. Related Work

Exposure bias [5, 22, 23] is well-known as a central problem for auto-regressive models such as language models. Scheduled sampling [5] is a form of self-training proposed to alleviate this training-inference bias [24] caused by the maximum-likelihood (MLE) based training objective. Beam search optimization (MLE) based training objective. Beam search optimization was proposed to reduce exposure bias in [25]. It collects and improves the scores of each ground truth token and the previous token leading to incorrect prediction in each beam search step. Recent work on reducing exposure bias for ASR using recurrent neural network transducers (RNN-T) claims that the direct application of scheduled sampling to RNN-T architectures is challenging as it is not a generative architecture [26], unlike self-attention based transformer like encoder-decoder architectures [27]. To mitigate this bias, the authors sample the target label from the top-k candidates from a language model instead of the model itself. In [28], the EMBR computation is optimized for RNN-T architectures by using alignment scores for target

hypotheses. The bias introduced during training-inference with LM rescoring is addressed in [16] using the EMBR objective. Prefix level EMBR is proposed in [29, 30, 15] to impose correct prefixes in the beam by boosting prefixes with less errors. Scheduled sampling uses a single hypothesis and EMBR uses multiple hypotheses in their objectives, but neither of the above approaches select the oracle hypothesis and boosts its score. The proposed O-1 mitigates this constraint and has shown consistent improvements in recognition performance.




Figure 1: Visualization of O-1 training defined in Equation (4) that boosts the score of the oracle hypothesis to the top of the beam.

3. O-1: Proposed Approach

In this work, we use Recurrent Neural Network Transducers [31] are used as the framework for EMBR and the proposed O-1 objectives. Figure 1 visualizes the training procedure corresponding to Equation (4).

3.1. EMBR

Given an input speech sequence \mathbf{X} , the encoder network, prediction and joint network in an RNN-T model help to compute the alignment probability $p(\mathbf{a} | \mathbf{X})$ with the help of forward-backward algorithm. The RNN-T training objective \mathcal{L}_{rnt} is defined as:

$$\mathcal{L}_{rnt} = -\log p(\mathbf{Y} | \mathbf{X}) = -\log \sum_{\mathbf{a} \in \mathcal{B}^{-1}(\mathbf{Y})} p(\mathbf{a} | \mathbf{X}) \quad (1)$$

Here, \mathcal{B} operation removes all blank labels in a given sequence and $p(\mathbf{Y} | \mathbf{X})$ is the probability of the predicted hypothesis \mathbf{Y} . Based on the objective in Equation (1), the EMBR loss [28] for RNNT is computed as:

$$\mathcal{L}_{embr} = \sum_{\mathbf{Y}_i \in \mathcal{Y}} \hat{p}(\mathbf{Y}_i | \mathbf{X}) W(\mathbf{Y}_i, \mathbf{Y}^*) \quad (2)$$

where $\hat{p}(\mathbf{Y}_i | \mathbf{X}) = \text{softmax}(p(\mathbf{Y}_i | \mathbf{X}))$ is the probability of i^{th} hypothesis in the beam obtained by applying a softmax function on $\log p$ of the RNN-T model to normalize over the n-best hypotheses. Here, \mathcal{Y} denotes the n-best list and $W(\mathbf{Y}_i, \mathbf{Y}^*)$ is the number of erroneous words computed using an edit distance between each hypothesis in the beam \mathbf{Y}_i and the ground-truth \mathbf{Y}^* . In the partial derivative of the EMBR objective below,

$$\frac{\partial \mathcal{L}_{embr}}{\partial \log p(\mathbf{Y}_i | \mathbf{X})} = \hat{p}(\mathbf{Y}_i | \mathbf{X}) (W(\mathbf{Y}_i, \mathbf{Y}^*) - \bar{W}) \quad (3)$$

$\bar{W} = \sum_{\mathbf{Y}_j \in \text{n-best}} \hat{p}(\mathbf{Y}_j | \mathbf{X}) W(\mathbf{Y}_j, \mathbf{Y}^*)$ denotes the expected word errors or risk among all the hypotheses for training an utterance. Monte-Carlo approximation of the n-best list results in $\bar{W} \approx \sum_{\mathbf{Y}_j \in \text{n-best}} W(\mathbf{Y}_j, \mathbf{Y}^*)$ as the average number of word

errors. The subtraction of \bar{W} in (3) performs variance reduction [32, 14]. In [8], the authors applied \bar{W} directly in (2) also as a form of variance reduction, but computing \bar{W} is empirical and tricky [22].

3.2. O-1 Training Procedure

The O-1 training objective $\mathcal{L}(\mathbf{X}, \mathbf{Y}, \mathbf{Y}^*)$ is defined as:

$$\mathcal{L}_{O1} = -\log p(\mathbf{Y}_{oracle} | \mathbf{X}) \cdot (1 - \text{WER}(\mathbf{Y}_{oracle}, \mathbf{Y}^*)) + \log p(\mathbf{Y}_{1best} | \mathbf{X}) \cdot \text{WER}(\mathbf{Y}_{1best}, \mathbf{Y}^*) \quad (4)$$

The training steps to implement Equation (4) are enumerated below:

1. Beam-search decoding is performed during training to extract n-best hypotheses $\mathbf{Y}_{i=1:n}$.
2. Each hypothesis \mathbf{Y}_i is scored against the ground-truth or pseudo-label from the teacher to compute WER.
3. The hypothesis with the best WER serves as the oracle hypothesis \mathbf{Y}_{oracle} and the hypothesis with the top probability becomes the 1-best hypothesis \mathbf{Y}_{1best} . Both oracle and 1-best hypotheses are chosen from the beam while dropping the rest of the hypotheses.
4. $-\log p(\mathbf{Y}_{i,j} | \mathbf{X})$ is the RNN-T loss for $i = oracle$ hypothesis and $j = 1 - best$ hypothesis computed using (1).
5. Each of these two RNN-T losses is scaled by its corresponding WER as described in Equation (4). This idea of scaling objectives with the evaluation metric has also been explored for pre-training in [33].
6. When the scores of 1-best and oracle hypotheses are equal, $\mathcal{L}_{O1} \rightarrow 0$ in (4) and the auxiliary \mathcal{L}_{rnt} loss will overcome this zero loss scenario as given in Equation (5).

It can be seen that O-1 training simplifies EMBR training by eliminating the need to consider all the alternate hypotheses and eliminates the computation of \bar{W} . O-1 objective also does not require a softmax computation that EMBR requires (Equation (2)). The final EMBR and O-1 losses used in our experiments is a multitask training objective given by:

$$\mathcal{L}'_{embr} = \mathcal{L}_{embr} + \gamma \mathcal{L}_{rnt} \quad (5)$$

$$\mathcal{L}'_{O1} = \mathcal{L}_{O1} + \lambda \mathcal{L}_{rnt}$$

O-1 explicitly targets the hypothesis with less WER to be boosted without relying on an error threshold $W - \bar{W}$ as in Equation (3). O-1 performs discriminative training between oracle and 1-best scores. This formulation of O-1 naturally addresses the label bias problem as well; when the model is pushing the oracle hypothesis score to be on top of the beam, the 1-best hypothesis score has to come down to respect the discriminative aspect of O-1.

3.3. O-1 Unsupervised training

Hard distillation [34, 35] involves training a student model with the target labels obtained on-the-fly from the teacher model as in (6). Given the reduced computational complexity compared to EMBR, we choose to employ O-1 in an unsupervised setting using the distillation loss given in [34]. Thus, O-1 loss with hard distillation becomes:

$$\mathcal{L}_{distill} = \mathcal{L}_{rnt}(\mathbf{X}, \mathbf{Y}') \quad (6)$$

$$\mathcal{L}_{O1, distill} = \mathcal{L}_{O1}(\mathbf{X}, \mathbf{Y}, \mathbf{Y}') + \lambda \mathcal{L}_{distill},$$

where, \mathbf{Y}' denotes the pseudo-label generated by the teacher model during training.

4. Experiments

4.1. Datasets

We conduct experiments on both publicly available and in-house corpora. SpeechStew, described in [36], a publicly available and well-benchmarked corpus, composed of multiple datasets (≈ 5000 hours) spanning several acoustic conditions and domains. It comprises of 960 hrs of Librispeech, 100 hrs of AMI, 1.5k hrs of CommonVoice corpus, 50 hrs of English broadcast news, 2k hrs of Switchboard/Fisher, 450 hrs of TED-lium and 80 hrs of WSJ corpus. The in-house corpus consists of Google’s short Voice Search (VS) utterances in U.S. English. This *supervised* ASR training data is anonymized and hand-transcribed. The development and test sets are a small fraction of the training set held out for validation and evaluation. Also included in the training are nearly 300 thousand hours of multi-domain utterances spanning domains such as Dictation, YouTube and Telephony for US English [37]. The *unsupervised* dataset contains 500M short search utterances. While the public corpora allows us to benchmark against state-of-the-art algorithms, the in-house corpora allows us to study performance on tail distributions and evaluate the scalability of the algorithm to large-scale training data and streaming E2E setups.

4.2. Model Description

The model architecture used in this work is based on the RNN-T architecture first described in [31]. It consists of a conformer encoder [38] and an LSTM decoder. The encoder comprises of several blocks, with each block including a series of multi-headed self attention, depth-wise convolution and feed-forward layers. For our experiments with SpeechStew, we use two different configurations, namely a 100M and 600M parameter model with 17 and 24 conformer layers with 512 and 1024 encoder dimensions, respectively. Both models use relative attention. All models are trained on use 80-dimensional log-mel filter bank coefficients. We use SpecAugment [39] with time and frequency masking for all experiments. For SpeechStew, the number of frequency masking bins $F = 27$ and time bins $T = 10$ with a maximum mask ratio set to 0.05. For the in-house dataset, $F = 27$, $T = 50$ and mask ratio to 0.1. For experiments with the in-house corpus, we use a 120M parameter, streaming Hybrid Autoregressive Transducer (HAT) [40] architecture, comprising of a conformer encoder with 12 layers and a step decoder with V2 embedding as described in [16]. In addition to SpecAugment, we use multistyle training similar to [41]. The experiments on SpeechStew use 1024 word-piece targets and the in-house experiments on the Voice Search task use 4K word-piece targets derived using the algorithm presented. in [42].

4.3. O-1 Supervised Self Training

We use the following training strategy to study the impact of the proposed algorithm with supervised-only training data. A 100M RNN-T model is trained from scratch using the supervised, SpeechStew corpora. The 600M parameter model is finetuned from a wav2vec2 [43] encoder that is pretrained with Librilight data. The model is trained to convergence (approximately 300K iterations) and serves as the baseline model for subsequent EMBR and O-1 training (100K iterations). Based on empirical evidence we set the RNN-T loss scaling factor as

$\lambda = 0.1, \gamma = 0.1$ for the EMBR \mathcal{L}_{embr} and the O-1 training objectives \mathcal{L}_{O1} . In all our O-1 experiments, the RNN-T loss in Equation (1) is normalized with the target label sequence length before using for O-1 loss computation in (4). SpeechStew test sets [36] are used for evaluation.

4.4. O-1 Unsupervised Learning

In the unsupervised learning scenario, the teacher is a non-causal model with right context fixed to 64 and the student is a streaming model with zero right context. The teacher model is trained with the supervised data alone and the baseline student model is trained with the same supervised data and additional unsupervised data. For unsupervised learning, we use a teacher model’s transcripts as pseudo labels. We use the hard distillation (Equation (6)), to train both student and teacher models jointly. Typically CE training converges after 500K iterations. This serves as the CE baseline to continue with EMBR training (Equation (2)). For O-1 training (Equation (4)), the teacher and student are initialized in the same way the baseline, i.e., we do not require the additional CE training step that conventional EMBR training requires. As explained in Section 3, CE objective is included as an auxiliary loss in the O-1 objective. All experiments related to unsupervised learning are conducted on the in-house corpora, since vast amounts of untranscribed data is available. Each training batch contains equal amounts of transcribed and untranscribed data. Beam-search is done with a beam size of 8 for hypotheses generation in EMBR training and inference. In the case of O-1, beam size of 18 is set for SpeechStew and 8 for in-house dataset during training. O-1 inference is done using beam size of 8 to perform fair comparison with EMBR.

5. Results and Analysis

Table 1 shows the performance comparison of O-1 with the baseline and the EMBR models. EMBR shows inconsistent behavior in different testsets. For instance, EMBR degrades on AMI-IHM, Librispeech *test* and *testother*, TED-LIUM and WSJ-eval92. EMBR shows gains on AMI-SDM, CHiME and Commonvoice which are relatively harder tasks over rest of the datasets. O-1 provides consistent improvement in all testsets over both the baseline and the EMBR models with a maximum of 25% relative gain on WSJ-eval92. The primary reason for O-1’s improved performance is learning from the best hypothesis targets. This enables the model to achieve closer to oracle performance during training as shown in Figure 2C and gets reflected in evaluation. We also conducted the same experiment

Table 1: Recognition performance on Speechstew evaluation sets with 100M and 600M parameter models

#Params	% WER	AMI		CHiME	CV	Libri-test		TED	WSJ
		IHM	SDM			clean	other		
100M	Baseline	10.2	24.8	34.9	11.5	2.3	4.9	4.8	2.6
	EMBR	11.3	23.7	33.2	11.2	2.4	5.4	4.7	2.7
	O-1	9.8	23.6	31.6	11.1	2.1	4.6	4.4	2.0
600M	Baseline	9.6	23.8	33.5	11.5	1.8	3.6	5.6	1.4
	EMBR	9.1	22.7	29.7	10.2	1.9	3.6	4.8	1.4
	O-1	8.7	20.4	25.3	8.9	1.7	3.4	4.3	1.3

on a 600M parameter model to investigate if the performance of the new objective is complementary to the increase in model complexity. Table 1 shows that on average, O-1 achieves 8.4% rel. improvement over EMBR and 13% relative improvement over the state-of-the-art baseline system in [36]. O-1 shows relative improvements of up to 25% on WSJ, 15% on CHiME




Figure 2: [A] Average WER of all test sets in SpeechStew evaluated at different checkpoints trained with different beam sizes using the proposed O-1 objective and EMBR. [B] Plot showing the number of hypotheses required to reach the optimal performance in EMBR and O-1. Beam search is fixed to compute 8 hypotheses. EMBR is computed from top8 hypotheses. O-1’s WER is close to the oracle WER (1.67%) compared to EMBR on librispeech test clean. Model is trained on SpeechStew datasets. [C] 1-best WER vs oracle WER on Librispeech test-clean using O-1, baseline and EMBR

and greater than 10% on TED, AMI-SDM and CommonVoice.

5.1. Comparison with Scheduled Sampling

Analogous to our approach, scheduled sampling also allows the model to train from its own predictions. However, O-1 differs from scheduled sampling in two aspects: (1) The oracle hypothesis serves as the target score to be improved, while the 1-best hypothesis score is suppressed. In the case of scheduled sampling, the 1-best hypothesis is chosen as the target based on the selection ratio. (2) The oracle hypothesis score is scaled by the corresponding WER in O-1 which is not done in scheduled sampling. Empirical analysis is conducted on Librispeech test set to evaluate the behavior of O-1 and scheduled sampling approach. Here we perform scheduled sampling at the utterance level without any language model [26]. O-1 obtains 8% relative improvement over scheduled sampling approach on test-other and 12.5% relative gain on test.

5.2. Training duration, Performance and Beam size

During O-1 training all the samples in the beam are not involved in loss computation which reduces the exposure to erroneous samples. Batch size is not impacted by the number of samples in the beam during O-1 training. These two factors allow the model trained with O-1 to use a larger beam size unlike EMBR training. Figure 2A shows the effect of increasing the beam size during training. “O1 - 18” (O-1 trained with beam size of 18) shows substantial gains over EMBR. We also conducted an experiment by training using the oracle hypothesis only and this achieved a train time similar to the baseline model. Efficient inferencing on TPUs [44] results in maximum computation resting on the N-best hypothesis generation. On average, the O-1 processes 945 examples/sec and EMBR processes 754 examples/sec for the SpeechStew data with a total of 3.6 million examples, which makes it faster over EMBR training.

5.3. Are all hypothesis needed during training ?

Figure 2B shows the performance of EMBR and O-1 based on the number of hypotheses involved in loss computation. Here, we argue that the reason for EMBR attaining the best performance with 8 hypotheses is that on average the oracle hypothesis is between 5th and 8th position within the top 8 hypotheses. O-1 is capable of retrieving the oracle hypothesis with a single hypothesis used in its training, and attains best performance threshold with the inclusion of 1-best.

To further evaluate the performance difference between actual and oracle WER, the O-1 and EMBR models were evaluated

at each training checkpoint as shown in Figure 2C. O-1’s 1-best WER gets closer to the oracle WER as training progresses and matches the oracle performance at 100K iterations, indicating that use of all n-best hypotheses is unnecessary.

5.4. Results on O-1 Unsupervised Learning

The experiments on a large-scale in-house data set comprising of short utterances up to 4 seconds long, shows a similar trend to the public data sets. Two audio sources were used for this evaluation, real and TTS generated audio. The test sets include several named entities. For example, the RPN datasets contains proper nouns, place names, music genres, artists and other media related entities. Each test set contains 10K utterances. From Table 2 we can see that O-1 attains 8% and 9% relative improvement over EMBR on two of the four testsets while yielding modest improvements on the remaining test sets. This indicates that the O-1 objective handles streaming recognition well compared to EMBR and scales to large-scale data sets.

Table 2: Recognition performance on in-house data sets in U.S. English including VS and rare proper nouns (RPN-) from multiple application sources.

% WER	VS	RPN-M	RPN-N	RPN-P	RPN-Q	RPN-Y
Baseline	6.0	13.2	10.1	37.3	20.6	24.1
EMBR	5.8	13.1	9.9	36.9	20.3	24.0
O-1	5.8	12.4	9.6	35.5	18.5	22.2

6. Conclusion

We have proposed O-1, a novel self-training objective to learn from oracle hypothesis. Our experiments shows that the O-1 reduces the difference between 1-best and oracle performance by a significant margin in multiple evaluation sets. O-1 shows promise by performing well on tasks where EMBR fails to improve. O-1 enjoys additional benefit of optimized training time over EMBR. In addition the results show that O-1 shows gains in both supervised and unsupervised setting under various test conditions. O-1 achieves 13%, 15% and 25% relative improvement over EMBR on AMI-HM, Librispeech-testother and WSJ-eval92 testsets respectively. In a real-time, streaming setting, O-1 provides 9% relative improvement over EMBR on a large scale in-house test set. O-1 can be further improved in the future using multiple hypotheses from the teacher model and applying advanced distillation approaches.

7. References

- [1] Colin Wei, Kendrick Shen, et al., “Theoretical analysis of self-training with deep networks on unlabeled data,” in *ICLR*, 2021.
- [2] Yves Grandvalet and Yoshua Bengio, “Semi-supervised learning by entropy minimization,” *Advances in NeurIPS*, vol. 17, 2004.
- [3] Marc’Aurelio Ranzato, Sumit Chopra, et al., “Sequence level training with recurrent neural networks,” in *ICLR*, 2016.
- [4] Ronald J Williams and David Zipser, “A learning algorithm for continually running fully recurrent neural networks,” *Neural computation*, vol. 1, no. 2, pp. 270–280, 1989.
- [5] Samy Bengio, Oriol Vinyals, et al., “Scheduled sampling for sequence prediction with recurrent neural networks,” *Advances in NeurIPS*, vol. 28, 2015.
- [6] Dzmitry Bahdanau, Philemon Brakel, et al., “An actor-critic algorithm for sequence prediction,” *ArXiv*, vol. abs/1607.07086, 2016.
- [7] Ronan Collobert, Awni Hannun, et al., “A fully differentiable beam search decoder,” in *Proceedings of the 36th ICML*, Kamalika Chaudhuri and Ruslan Salakhutdinov, Eds. 09–15 Jun 2019, vol. 97 of *PMLR*, pp. 1341–1350, PMLR.
- [8] Rohit Prabhavalkar, Tara N Sainath, et al., “Minimum word error rate training for attention-based sequence-to-sequence models,” in *IEEE ICASSP*. IEEE, 2018, pp. 4839–4843.
- [9] Liang Lu, Zhong Meng, et al., “On Minimum Word Error Rate Training of the Hybrid Autoregressive Transducer,” in *Proc. Interspeech 2021*, 2021, pp. 3435–3439.
- [10] Daniel Povey, *Discriminative training for large vocabulary speech recognition*, Ph.D. thesis, University of Cambridge, 2005.
- [11] Mathias Müller and Rico Sennrich, “Understanding the properties of minimum Bayes risk decoding in neural machine translation,” in *International Joint Conference on NLP (Volume 1: Long Papers)*, Online, Aug. 2021.
- [12] Zhiyun Lu, Yanwei Pan, et al., “Input length matters: Improving rnn-t and mwer training for long-form telephony speech recognition,” *arXiv preprint arXiv:2110.03841*, 2021.
- [13] Vaibhava Goel and William J Byrne, “Minimum bayes-risk automatic speech recognition,” *Computer Speech & Language*, vol. 14, no. 2, pp. 115–135, 2000.
- [14] Matt Shannon, “Optimizing expected word error rate via sampling for speech recognition,” *arXiv preprint arXiv:1706.02776*, 2017.
- [15] Murali Karthick Baskar, Lukáš Burget, et al., “Promising accurate prefix boosting for sequence-to-sequence asr,” in *IEEE ICASSP*, 2019, pp. 5646–5650.
- [16] Wang Weiran, Tongzhou Chen, et al., “Improving Rare Word Recognition with LM-aware MWER Training,” in *Proc. Interspeech 2022*, 2022, pp. 1031–1035.
- [17] Karel Veselý, Arnab Ghoshal, et al., “Sequence-discriminative training of deep neural networks,” in *Interspeech*, 2013, vol. 2013, pp. 2345–2349.
- [18] Daniel Povey and P Woodland, “Improved discriminative training techniques for large vocabulary continuous speech recognition,” in *ICASSP*. IEEE, 2001, vol. 1, pp. 45–48.
- [19] Brian Kingsbury, “Lattice-based optimization of sequence classification criteria for neural-network acoustic modeling,” in *2009 IEEE ICASSP*. IEEE, 2009, pp. 3761–3764.
- [20] Daniel Povey and Brian Kingsbury, “Evaluation of proposed modifications to mpe for large scale discriminative training,” in *2007 IEEE ICASSP’07*. IEEE, 2007, vol. 4, pp. IV–321.
- [21] Mohammad ZeinElddeen, Kartik Audhkhasi, et al., “Robust knowledge distillation from rnn-t models with noisy training labels using full-sum loss,” in *IEEE ICASSP*, 2023, pp. 1–5.
- [22] Steven J Rennie, Etienne Marcheret, et al., “Self-critical sequence training for image captioning,” in *Proceedings of the IEEE conference on CVPR*, 2017, pp. 7008–7024.
- [23] Tianxing He, Jingzhao Zhang, et al., “Exposure bias versus self-recovery: Are distortions really incremental for autoregressive text generation?,” *arXiv preprint arXiv:1905.10617*, 2019.
- [24] Abhishek Arun, Barry Haddow, et al., “A unified approach to minimum risk training and decoding,” in *Proceedings on Statistical Machine Translation and MetricsMATR*, 2010, pp. 365–374.
- [25] Sam Wiseman and Alexander M. Rush, “Sequence-to-sequence learning as beam-search optimization,” in *Proceedings of the 2016 Conference on EMNLP*, Austin, Texas, Nov. 2016, pp. 1296–1306, Association for Computational Linguistics.
- [26] Xiaodong Cui, Brian Kingsbury, et al., “Reducing Exposure Bias in Training Recurrent Neural Network Transducers,” in *Proc. Interspeech 2021*, 2021, pp. 1802–1806.
- [27] Dzmitry Bahdanau, Jan Chorowski, et al., “End-to-end attention-based large vocabulary speech recognition,” in *IEEE ICASSP*. IEEE, 2016, pp. 4945–4949.
- [28] Jinxi Guo, Gautam Tiwari, et al., “Efficient minimum word error rate training of rnn-transducer for end-to-end speech recognition,” in *Interspeech 2020*, 2020.
- [29] V. Doumptiotis, S. Tsakalidis, et al., “Discriminative training for segmental minimum bayes risk decoding,” in *ICASSP Proceedings*, 2003, vol. 1, pp. I–I.
- [30] Sara Sabour, William Chan, et al., “Optimal completion distillation for sequence learning,” in *ICLR*, 2019.
- [31] Alex Graves, “Sequence transduction with recurrent neural networks,” *arXiv preprint arXiv:1211.3711*, 2012.
- [32] Richard S Sutton, David McAllester, et al., “Policy gradient methods for reinforcement learning with function approximation,” *NeurIPS*, vol. 12, 1999.
- [33] Murali Karthick Baskar, Andrew Rosenberg, et al., “Ask2mask: Guided data selection for masked speech modeling,” *IEEE JSTSP*, vol. 16, no. 6, pp. 1357–1366, 2022.
- [34] Daniel S. Park, Yu Zhang, et al., “Improved Noisy Student Training for Automatic Speech Recognition,” in *Proc. Interspeech 2020*, 2020, pp. 2817–2821.
- [35] Dongseong Hwang, Khe Chai Sim, et al., “Comparison of soft and hard target rnn-t distillation for large-scale asr,” *arXiv preprint arXiv:2210.05793*, 2022.
- [36] William Chan, Daniel Park, et al., “Speechstew: Simply mix all available speech recognition data to train one large neural network,” *arXiv preprint arXiv:2104.02133*, 2021.
- [37] Arun Narayanan, Rohit Prabhavalkar, et al., “Recognizing long-form speech using streaming end-to-end models,” in *2019 IEEE ASRU*. IEEE, 2019, pp. 920–927.
- [38] Anmol Gulati, James Qin, et al., “Conformer: Convolution-augmented transformer for speech recognition,” *arXiv preprint arXiv:2005.08100*, 2020.
- [39] Daniel S. Park, William Chan, et al., “SpecAugment: A Simple Data Augmentation Method for Automatic Speech Recognition,” in *Proc. Interspeech 2019*, 2019, pp. 2613–2617.
- [40] Ehsan Variani, David Rybach, et al., “Hybrid autoregressive transducer (hat),” in *IEEE ICASSP*. IEEE, 2020, pp. 6139–6143.
- [41] Chanwoo Kim, Ananya Misra, et al., “Generation of large-scale simulated utterances in virtual rooms to train deep-neural networks for far-field speech recognition in google home,” 2017.
- [42] Taku Kudo and John Richardson, “Sentencepiece: A simple and language independent subword tokenizer and detokenizer for neural text processing,” *arXiv preprint arXiv:1808.06226*, 2018.
- [43] Alexei Baevski, Yuhao Zhou, et al., “wav2vec 2.0: A framework for self-supervised learning of speech representations,” in *Advances in NeurIPS*, 2020.
- [44] Norman P Jouppi, Doe Hyun Yoon, et al., “Ten lessons from three generations shaped googleâs tpuv4i: Industrial product,” in *2021 ACM/IEEE ISCA*. IEEE, 2021, pp. 1–14.