



HAL
open science

End-to-end Joint Rich and Normalized ASR with a limited amount of rich training data

Can Cui, Imran Ahamad Sheikh, Mostafa Sadeghi, Emmanuel Vincent

► **To cite this version:**

Can Cui, Imran Ahamad Sheikh, Mostafa Sadeghi, Emmanuel Vincent. End-to-end Joint Rich and Normalized ASR with a limited amount of rich training data. 2023. hal-04304642

HAL Id: hal-04304642

<https://inria.hal.science/hal-04304642>

Preprint submitted on 24 Nov 2023

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.



Distributed under a Creative Commons Attribution 4.0 International License

END-TO-END JOINT RICH AND NORMALIZED ASR WITH A LIMITED AMOUNT OF RICH TRAINING DATA

Can Cui^{1,2}, Imran Sheikh², Mostafa Sadeghi¹, Emmanuel Vincent¹

¹Université de Lorraine, CNRS, Inria, LORIA, F-54000 Nancy, France

²Vivoka, Metz, France

ABSTRACT

Joint rich and normalized automatic speech recognition (ASR), that produces transcriptions both with and without punctuation and capitalization, remains a challenge. End-to-end (E2E) ASR models offer both convenience and the ability to perform such joint transcription of speech. Training such models requires paired speech and rich text data, which is not widely available. In this paper, we compare two different approaches to train a stateless Transducer-based E2E joint rich and normalized ASR system, ready for streaming applications, with a limited amount of rich labeled data. The first approach uses a language model to generate pseudo-rich transcriptions of normalized training data. The second approach uses a single decoder conditioned on the type of the output. The first approach leads to E2E rich ASR which perform better on out-of-domain data, with up to 9% relative reduction in errors. The second approach demonstrates the feasibility of an E2E joint rich and normalized ASR system using as low as 5% rich training data with moderate (2.42% absolute) increase in errors.

Index Terms— ASR, rich transcriptions, RNN-T, limited data, streaming ready

1. INTRODUCTION

Transcribing speech into rich text with punctuation and capitalization has been an active area of research in automatic speech recognition (ASR) [1–3]. Rich transcriptions are essential in applications that expect human-readability of the transcribed text as well in applications that use natural language processing (NLP) systems that rely on punctuation and case information. Joint rich and normalized ASR, that produces transcriptions both with and without punctuation and capitalization, is highly desirable because (a) it simplifies model deployment and maintenance, and (b) it extends compatibility with NLP models which either exploit or discard punctuation and case information.

The conventional approach to rich ASR involves post-processing the output of a normalized ASR system with a case and punctuation restoration model [1–4]. For instance,

the authors in [3] used a modified Recurrent Neural Network-Transducer (RNN-T) ASR model with stateless prediction network [5] and a fine-tuned ELECTRA-Small language model (LM) [6]. While this approach does not require rich transcription as ASR training labels, it increases the total model size and inference time. In addition, it does not exploit acoustic information to transcribe punctuation.

The second approach employs an E2E ASR model to directly generate rich transcriptions [7, 8]. However, this can lead to poorer word recognition, hence poorer normalized ASR performance. To mitigate this, the study in [9] introduced an auxiliary CTC loss for transcribing non-punctuated text at the output of an intermediate layer of the E2E model. More recently, the study in [10] proposed three autoregressive decoders that separately output spoken, written, and combined transcriptions. While both of these methods leverage normalized transcriptions to enhance the accuracy of rich transcriptions, the challenges related to lack of training data and increased model size persist.

This paper aims for an E2E joint rich and normalized ASR system which is (a) accurate at rich as well as normalized transcription tasks, (b) trainable with a limited amount of rich labeled data, and (c) ready for streaming applications. We introduce and compare two complementary approaches to train a stateless Transducer-based E2E joint rich and normalized ASR model. The first training approach uses an LM to generate pseudo-rich transcriptions of normalized training data. However, such LMs may not be accurate enough or available for certain domains [11] and/or languages. We propose a second training approach to address such scenarios. This approach uses a single decoder which is conditioned on the type of the output. The rest of the paper is organized as follows. Section 2 presents the preliminaries. Section 3 presents our training approaches. Section 4 presents our experimental evaluation, followed by the conclusion in Section 5.

2. PRELIMINARIES

2.1. Stateless Transducer-based Rich ASR

An E2E rich ASR system [9] directly transcribes speech into a rich transcription. Given an acoustic feature sequence $X \in$




Fig. 1. (a) Stateless Transducer-based Rich ASR, (b) 2-Decoder Joint Normalized-Rich ASR, (c) the proposed Conditioned Predictor ASR, (d) input layers in the Original and Conditioned Predictors.

$\mathbb{R}^{L \times A}$ where L is the sequence length and A the feature dimension, the training objective is to maximize the probability

$$P(Y^R|X) = \prod_{s=1}^S P(y_s^R | y_{[1:s-1]}^R, X), \quad (1)$$

by generating a sequence $Y^R \in \mathbb{R}^{S^R}$, where R stands for rich transcription, and S^R represents the rich sequence length. The loss function can be written as

$$\mathcal{L}^R = - \sum_{Y^R, X} \log P(Y^R|X). \quad (2)$$

A stateless Transducer-based E2E ASR system, which uses an RNN-T framework [12] with a stateless prediction network [5], can be extended to the rich transcription task (see Fig. 1(a)). This brings the high accuracy and naturally streaming recognition capability of RNN-T to rich transcription.

2.2. 2-Decoder Joint Normalized-Rich ASR

Inspired by [9, 10], we have designed an ASR system with two sets of Predictors and Joiners (see Fig. 1(b)). Using the output of the same Encoder, these two Decoders generate the rich and normalized transcriptions. Their respective training objectives can be expressed as¹

$$P(Y^N|X) = \prod_{s=1}^{S^N} P(y_s^N | y_{[1:s-1]}^N, X), \quad (3)$$

$$P(Y^R|X) = \prod_{s=1}^{S^R} P(y_s^R | y_{[1:s-1]}^R, X),$$

where N stands for normalized transcription, and S^N , S^R represent the normalized, rich sequence length. Therefore, the joint loss function is defined as

$$\begin{aligned} \mathcal{L}_{2\text{-decoder}} &= \mathcal{L}^N + \mathcal{L}^R \\ &= - \sum_{Y^N, X} \log P(Y^N|X) - \sum_{Y^R, X} \log P(Y^R|X). \end{aligned} \quad (4)$$

¹These equations are kept consistent with (1) and do not account for the stateless and/or streaming operation of the Transducer model.

3. PROPOSED METHODS

3.1. Training Rich ASR using Pseudo-rich Transcriptions

Automatically generated pseudo-rich transcriptions can be used to train rich ASR models in the absence of human-labeled rich ASR training data. The use of pseudo-rich transcriptions to train rich ASR models remains unexplored in prior works. The work in [8] proposed a semi-supervised learning method in which a rich ASR system is first trained on a small amount of human-labelled rich training data and then used to automatically generate pseudo-rich transcripts of a speech corpus. The focus of this approach is rich transcription of speech phenomena such as fillers, laughter and coughs, etc. and hence automatically transcribing the training speech corpus is necessary. In contrast, our rich transcription task requires case and punctuation prediction. Case and punctuation enhanced rich transcriptions can be directly obtained from normalized transcriptions using state-of-the-art case and punctuation restoration models [1–4]. Interestingly, this approach can be applied in scenarios where there is no rich labelled training data at all.

3.2. Conditioned Predictor ASR

A drawback of the pseudo-rich transcription driven rich ASR training approach, presented in Section 3.1, is that errors made by the case and punctuation prediction model may propagate into the rich ASR model. Moreover, such case and punctuation restoration models may not be accurate enough, or even available, for certain domains [11] and/or languages. Hence, we propose a second training approach to address such scenarios.

Our approach effectively utilizes the small amount of rich training data and the large amount of normalized training data by using a single conditioned Predictor to handle rich and normalized transcriptions. As shown in Fig. 1 (c), the Predictor is conditioned on the transcription mode ID input which specifies whether we want a normalized (N) or rich (R) output. The mode ID input is inspired by the language ID information used by E2E multilingual ASR models [13, 14]. As shown in

Fig. 1 (d), the token embeddings are concatenated with the mode embedding before feeding them to the following Predictor layers. The Conditioned Predictor ASR system can use a loss function similar to Equation (4). However, only a portion of the input samples will have a rich transcription reference and will use the corresponding loss.

4. EXPERIMENTS

4.1. Dataset and metrics

4.1.1. Dataset

We use the LibriSpeech corpus [15] to conduct our experiments. Following the standard split, our training set includes train-100, train-360 and train-500 subsets. Similarly, the validation sets consist of dev-clean and dev-other, while the test sets are test-clean and test-other. We retrieve the rich transcripts of the LibriSpeech corpus from the original Project Gutenberg texts.² In order to evaluate the model’s performance on real unseen data we use the independent headset microphone (IHM) recordings of the AMI meeting corpus [16] test set along with their original rich transcription.

4.1.2. Metrics

The regular Word Error Rate (WER) metric used in ASR evaluation is insufficient to assess the performance in rich transcription tasks. Traditionally, Precision, Recall, and F1 score metrics have been used to evaluate performance of rich transcription [3, 4]. However, these prior works present separate performance figures for each of the punctuation marks and lack a standardized metric for casing. In order to have an understanding of the overall performance of prediction of punctuation marks as well as casing, and to facilitate the comparison of error rates across words, cases, and punctuation marks, we introduce the concepts of Case-only Error Rate (CaseER), Punctuation-only Error Rate (PuncER) and Case-Punctuation-aware WER (CP-WER) alongside WER. Their calculation methods are as follows:

$$\begin{aligned} \text{CaseER} &= \frac{E_{np-c} - E_{np-nc}}{N_c}, \\ \text{PuncER} &= \frac{E_{p-nc} - E_{np-nc}}{N_p}, \\ \text{CP-WER} &= \frac{E_{p-c}}{N_{p-c}}, \quad \text{WER} = \frac{E_{np-nc}}{N_{np-nc}} \end{aligned} \quad (5)$$

where E represents the total number of substitution, deletion and insertion errors, N represents the total number of words and/or punctuation marks. Suffixes c and p stand for presence of case and punctuation marks whereas suffices nc and np stand for absence of case and punctuation marks, respectively. Table 1 shows an example of these metrics computed on a sample reference and prediction.

²available at <https://www.openslr.org/12/>

Table 1. Example of evaluation metrics computed for reference *Hi, I am Chloe.* and hypothesis *hey I am chloe.*

WER	CaseER	PuncER	CP-WER
1/4 = 0.25	(2 - 1)/2 = 0.5	(2 - 1)/2 = 0.5	3/6 = 0.5

4.2. Model and training setup

We use 80-dimensional log Mel filterbank features as input to all ASR models. Our tokenizer is a SentencePiece model [17] with a vocabulary size of 500. All the ASR models (E2E ASR, 2-decoder ASR, Conditioned Predictor ASR) were trained using the recipe for the pruned stateless-Transducer based on the Zipformer [18] Encoder in the icefall toolkit.³ For Conditioned Predictor ASR, the Token Embedding and Mode Embedding have 500 and 12 dimensions, respectively. All the models were trained for 40 epochs.

We use the Rpunct [19] model to generate rich transcriptions from normalized transcriptions. This model is derived from the BERT [20] LM after fine-tuning it for English punctuation and case restoration tasks. Given lower-cased text, the model can restore upper-casing, period (“.”), exclamation (“!”), question mark (“?”), comma (“,”), colon (“:”), semi-colon (“;”), apostrophe (“’”) and dash (“-”). The generated pseudo-rich transcriptions may differ from the expected rich transcriptions, as illustrated by some sample Librispeech sentences in Table 2.

Table 2. Examples of paired rich transcriptions and pseudo-rich transcriptions from LibriSpeech dataset.

Rich transcriptions	Pseudo-rich transcriptions
See, I could even get in myself!	See I could even get in myself.
A man said to the universe: S ir, I exist!	A man said to the universe s ir, I exist.
A red-haired mountain of a man, with an apparently inexhaustible store of energy.	A red haired mountain of a man with an apparently inexhaustible store of energy.

4.3. Evaluation results

4.3.1. Effectiveness of E2E models

Table 3 presents the (%) error rates for different rich ASR systems trained using the conventional and proposed approaches. The conventional cascaded system, which uses a normalized output ASR followed by a case and punctuation restoration LM, is included as a baseline. Comparison of error rates of the cascaded system and E2E models trained using rich data shows that the latter can lead to significant reductions in CP-WER. For instance, when using the 2-Decoder ASR (R + N) we obtain a relative CP-WER reduction of up to 24% (from 11.7% to 8.85%).

³https://github.com/k2-fsa/icefall/blob/master/egs/librispeech/ASR/pruned_transducer_stateless7

Table 3. ASR evaluation results in term of (%) Word Error Rate (WER), Case-only Error Rate (CaseER), Punctuation-only Error Rate (PuncER) and Case-Punctuation-aware WER (CP-WER). All ASR are trained on LibriSpeech train-960 set with different type of transcriptions (Trans.): normalized (N), rich (R) or pseudo-rich (R’). The 2-Decoder ASR and the Conditioned Predictor ASR use both rich (R or R’) and normalized transcriptions of the entire train-960 set.

Type of ASR	# Param.	Trans.	LibriSpeech test-clean				LibriSpeech test-other				AMI			
			WER	CaseER	PuncER	CP-WER	WER	CaseER	PuncER	CP-WER	WER	CaseER	PuncER	CP-WER
Cascaded (ASR+LM)	180 M	N	2.18	35.60	44.33	11.70	5.03	34.82	47.69	14.91	37.46	45.92	74.84	45.99
Rich only	70 M	R	2.45	28.21	30.39	9.12	5.73	25.14	30.51	11.94	39.59	38.57	97.02	50.61
	70 M	R’	2.35	31.42	44.43	11.60	5.49	30.14	47.39	14.94	39.01	37.39	71.82	46.33
2-Decoder	71 M	R + N	2.33	26.99	29.99	8.85	5.45	24.76	31.05	11.76	39.12	39.44	96.48	50.18
	71 M	R’ + N	2.28	31.68	45.10	11.67	5.38	29.53	47.81	14.88	39.33	33.64	65.63	45.43
Conditioned Predictor	70 M	R + N	2.35	35.61	29.14	9.32	5.49	30.77	29.62	12.01	40.70	38.40	96.40	51.47
	70 M	R’ + N	2.25	41.86	44.60	12.30	5.24	33.99	46.84	15.42	38.77	35.62	85.65	45.24

Table 4. (%) error rates of Conditioned Predictor ASR on LibriSpeech test-clean when trained with different proportions of rich and normalized (R + N) training data.

(R : N) Proportion	WER	CaseER	PuncER	CP-WER
0.50 : 0.50	2.38	34.86	29.33	9.32
0.35 : 0.65	2.37	45.44	28.99	9.95
0.20 : 0.80	2.34	56.37	29.96	10.83
0.05 : 0.95	2.56	58.58	31.59	11.46

When comparing the 2-Decoder ASR with the proposed Conditioned Predictor ASR on the LibriSpeech test set, the latter exhibits a decrease in PuncWER. For example, when using the same R + N training set, the Conditioned Predictor ASR achieves a relative decrease of 3% on test-clean and 5% on test-other in PuncER compared to the former. But 2-Decoder ASR have a better performance in term of case generation, by reducing relatively 24% the CaseER on test-clean.

4.3.2. Rich versus pseudo-rich transcriptions for training

Comparison of different E2E ASR trained using rich transcriptions versus pseudo-rich transcriptions shows that use of pseudo-rich training transcriptions yields higher error rates on in-domain LibriSpeech test-clean and test-other sets. In contrast, the use of pseudo-rich training transcriptions yields lower error rates on the out-of-domain AMI test set. For instance, the 2-Decoder ASR exhibits a CP-WER increase of 2.82% (from 8.85% to 11.67%) and 3.12% (from 11.76% to 14.88%) absolute on LibriSpeech test-clean and test-other sets, respectively, when using pseudo-rich training transcriptions (R’+N) instead of the rich training transcriptions (R+N). However, the same model exhibits a CP-WER reduction of 4.75% (from 50.18% to 45.43%) absolute on the AMI test set when using pseudo-rich training transcriptions.

4.3.3. Conditioned Predictor ASR with limited rich data

Unlike the 2-Decoder ASR, the Conditioned Predictor ASR can be trained with different amount of rich and normalized

training data. To evaluate the effectiveness of this capability, we trained the Conditioned Predictor ASR with different proportions of rich and normalized (R + N) training data. Table 4 presents the error rates obtained on LibriSpeech test-clean set. It can be observed that the CP-WER and PuncER performances of this model increase by small margins even the amount of rich training data is reduced drastically. Using only 5% of rich training data can still achieve a CP-WER of 11.74%, with only 2.42% absolute increase. This shows that the Predictor-Conditioned ASR can effectively utilize normalized data to enhance the generation of rich labels. Therefore, this model can serve scenarios having limited or severely limited rich training data.

5. CONCLUSION

In this paper, we introduced two approaches to train a stateless Transducer-based E2E joint rich and normalized ASR with a limited amount of rich labelled data. The first approach uses a language model to generate pseudo-rich transcriptions of normalized training data. Experimental results demonstrate that, compared to the model trained on rich data, the model trained on pseudo-rich data achieves a maximum relative reduction in CP-WER of up to 9% on on out-of-domain data. The second approach uses a single conditioned Predictor. The second approach demonstrates the feasibility of an E2E joint rich and normalized ASR system using as low as 5% rich training data with moderate (2.42% absolute) increase in errors.

6. ACKNOWLEDGMENTS

Experiments presented in this paper were carried out using the Grid’5000 testbed, supported by a scientific interest group hosted by Inria and including CNRS, RENATER and several Universities as well as other organizations (see <https://www.grid5000.fr>).

7. REFERENCES

- [1] Jiangyan Yi, Jianhua Tao, Ye Bai, Zhengkun Tian, and Cunhang Fan, “Adversarial transfer learning for punctuation restoration,” *arXiv e-prints*, pp. arXiv–2004, 2020.
- [2] Junwei Liao, Sefik Eskimez, Liyang Lu, Yu Shi, Ming Gong, Linjun Shou, Hong Qu, and Michael Zeng, “Improving readability for automatic speech recognition transcription,” *ACM Transactions on Asian and Low-Resource Language Information Processing*, vol. 22, no. 5, pp. 1–23, 2023.
- [3] Martin Poláček, Petr Červa, Jindřich Žďánký, and Lenka Weingartová, “Online Punctuation Restoration using ELECTRA Model for streaming ASR Systems,” in *Proc. INTERSPEECH 2023*, 2023, pp. 446–450.
- [4] Qiushi Huang, Tom Ko, H Lilian Tang, Xubo Liu, and Bo Wu, “Token-level supervised contrastive learning for punctuation restoration,” in *INTERSPEECH 2021*, 2021, pp. 2012–2016.
- [5] Mohammadreza Ghodsi, Xiaofeng Liu, James Apfel, Rodrigo Cabrera, and Eugene Weinstein, “RNN-transducer with stateless prediction network,” in *IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2020, pp. 7049–7053.
- [6] Kevin Clark, Minh-Thang Luong, Quoc V. Le, and Christopher D. Manning, “ELECTRA: Pre-training text encoders as discriminators rather than generators,” in *ICLR*, 2020.
- [7] Hayato Futami, Emiru Tsunoo, Kentaro Shibata, Yosuke Kashiwagi, Takao Okuda, Siddhant Arora, and Shinji Watanabe, “Streaming joint speech recognition and disfluency detection,” in *IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2023, pp. 1–5.
- [8] Tomohiro Tanaka, Ryo Masumura, Mana Ihuri, Akihiko Takashima, Shota Orihashi, and Naoki Makishima, “End-to-end rich transcription-style automatic speech recognition with semi-supervised learning,” in *INTERSPEECH 2021*, 2021, pp. 4458–4462.
- [9] Jumon Nozaki, Tatsuya Kawahara, Kenkichi Ishizuka, and Taiichi Hashimoto, “End-to-end speech-to-punctuated-text recognition,” in *INTERSPEECH 2022*, 2022, pp. 1811–1815.
- [10] Mana Ihuri, Hiroshi Sato, Tomohiro Tanaka, Ryo Masumura, Saki Mizuno, and Nobukatsu Hojo, “Transcribing Speech as Spoken and Written Dual Text Using an Autoregressive Model,” in *Proc. INTERSPEECH 2023*, 2023, pp. 461–465.
- [11] Monica Sunkara, Srikanth Ronanki, Kalpit Dixit, Sravan Bodapati, and Katrin Kirchhoff, “Robust prediction of punctuation and truecasing for medical ASR,” in *Proceedings of the First Workshop on Natural Language Processing for Medical Conversations*, Online, July 2020, pp. 53–62, Association for Computational Linguistics.
- [12] Alex Graves, “Sequence transduction with recurrent neural networks,” *arXiv e-prints*, pp. arXiv–1211, 2012.
- [13] Jinyu Li et al., “Recent advances in end-to-end automatic speech recognition,” *APSIPA Transactions on Signal and Information Processing*, vol. 11, no. 1, 2022.
- [14] Shubham Toshniwal, Tara N Sainath, Ron J Weiss, Bo Li, Pedro Moreno, Eugene Weinstein, and Kanishka Rao, “Multilingual speech recognition with a single end-to-end model,” in *IEEE international conference on acoustics, speech and signal processing (ICASSP)*, 2018, pp. 4904–4908.
- [15] Vassil Panayotov, Guoguo Chen, Daniel Povey, and Sanjeev Khudanpur, “Librispeech: An ASR corpus based on public domain audio books,” in *IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2015, pp. 5206–5210.
- [16] Jean Carletta, Simone Ashby, Sebastien Bourban, Mike Flynn, Mael Guillemot, Thomas Hain, Jaroslav Kadlec, Vasilis Karaiskos, Wessel Kraaij, Melissa Kronenthal, et al., “The AMI meeting corpus: A pre-announcement,” in *International Workshop on Machine Learning for Multimodal Interaction*, 2005, pp. 28–39.
- [17] Taku Kudo and John Richardson, “Sentencepiece: A simple and language independent subword tokenizer and detokenizer for neural text processing,” in *2018 Conference on Empirical Methods in Natural Language Processing: System Demonstrations*, 2018, pp. 66–71.
- [18] Daniel Povey, “zipformer,” https://github.com/k2-fsa/icefall/blob/master/egs/librispeech/ASR/pruned_transducer_stateless7/zipformer.py, 2023, GitHub repository.
- [19] Daulet Nurmanbetov, “rpunct,” <https://github.com/Felflare/rpunct.git>, 2021, GitHub repository.
- [20] Jacob Devlin Ming-Wei Chang Kenton and Lee Kristina Toutanova, “Bert: Pre-training of deep bidirectional transformers for language understanding,” in *Proceedings of naacL-HLT*, 2019, vol. 1, p. 2.