

HYSTOC: OBTAINING WORD CONFIDENCES FOR FUSION OF END-TO-END ASR SYSTEMS

Karel Beneš, Martin Kocour, Lukáš Burget

Brno University of Technology, Czechia

ABSTRACT

End-to-end (e2e) systems have recently gained wide popularity in automatic speech recognition. However, these systems do generally not provide well-calibrated word-level confidences. In this paper, we propose Hystoc, a simple method for obtaining word-level confidences from hypothesis-level scores. Hystoc is an iterative alignment procedure which turns hypotheses from an n-best output of the ASR system into a confusion network. Eventually, word-level confidences are obtained as posterior probabilities in the individual bins of the confusion network. We show that Hystoc provides confidences that correlate well with the accuracy of the ASR hypothesis. Furthermore, we show that utilizing Hystoc in fusion of multiple e2e ASR systems increases the gains from the fusion by up to 1 % WER absolute on Spanish RTVE2020 dataset. Finally, we experiment with using Hystoc for direct fusion of n-best outputs from multiple systems, but we only achieve minor gains when fusing very similar systems.

Index Terms— confidences measures, system fusion, end-to-end systems, automatic speech recognition

1. INTRODUCTION

The rising popularity of end-to-end (e2e) systems comes largely from their seemingly simple structure. While the underlying neural architectures vary greatly [1, 2, 3], the decoding process is generally very simple. Oftentimes, it consists of iteratively finding the highest-scoring entry in a softmax layer (listen-attend-spell [1], LAS, and recurrent neural transducers [4], RNN-T), possibly even in parallel (connectionist temporal classification [5], CTC). These approaches are straightforward compared to a complicated beam-search decoding through a recognition network [6] in the hybrid DNN/HMM systems.

However, this simplicity comes at a cost: The outputs of e2e systems rarely provide localized transcription variants or their scores. It may be possible to extract some sort of localized probability information; however, the outputs of modern large classification networks are known to be miscalibrated in general [7, 8]. The nature of this localized information may vary in its very nature (e.g., the input-synchronous CTC assumes a specific relation between neural network outputs and final transcription, and it operates on a very different time-scale than label-synchronous LAS system) or in granularity (graphemic / sub-word / word level outputs), thus preventing direct use of this information for confidence estimation on a unified level, e.g. words.

The lack of word-level confidences limits the utility of e2e system, as word-level confidences provide several benefits: (1) They can be utilized by downstream applications [9], (2) they can be used in the scope of semi-supervised learning to filter out uncertain parts

of the machine annotated data [10], or (3) they can be used for finer fusion of multiple ASR systems [11].

In this work, we propose to estimate word-level confidences and reap the benefits they provide by finding the common denominator of e2e systems: They are trained to directly model the probability $P(\mathbf{y}|\mathbf{x})$ of token sequence \mathbf{y} being the correct transcription of the input audio \mathbf{x} . Therefore, we assume that when a list of N best transcription variants (hypotheses) is produced by an e2e system, a score $s_i \propto \log P(\mathbf{y} = \mathbf{h}_i|\mathbf{x})$ is available for each hypothesis \mathbf{h}_i . Since this score s_i is independent of the internal tokenization of \mathbf{h}_i by the ASR system, any desired re-tokenization of the hypothesis can be taken to bridge the gap between different e2e systems and allowing for a unified approach.

In this paper, we propose Hystoc¹ – a simple alignment-based procedure for obtaining word-level confidences. We show that the confidences obtained with Hystoc are surprisingly well calibrated and that they significantly improve gains from fusion of multiple e2e systems. Finally, we extend Hystoc to directly fuse individual hypotheses from different systems.




Fig. 1: A confusion network representing strings ABC, AB, and AC with probabilities 0.7, 0.2, and 0.1 respectively. Note that while each ϵ -transition corresponds to a single hypothesis omitting the corresponding letter, all of the letter transitions actually aggregate multiple hypotheses.

2. DISTILLING WORD-LEVEL CONFIDENCES

In this section, we describe how Hystoc² restores word-level confidences from hypothesis level scores. Prior to the actual computation, observe that all hypotheses h_i can be re-tokenized into a desired target form, e.g. into whitespace-delimited words for European languages.

Hystoc is an iterative alignment procedure. Specifically, it aligns individual hypotheses h_i into a confusion network [12] (see Fig. 1) in the order of decreasing hypothesis scores s_i . When aligning h_i into the confusion network, Hystoc finds the highest-scoring path h^* through the confusion network and finds a Levenshtein alignment between h_i and h^* . Thus, the bins in the confusion network are determined into which the individual tokens from h_i shall be incorporated, possibly introducing new bins in case the alignment contains insertions. The score of each transition corresponding to a

¹<https://github.com/BUTSpeechFIT/hystoc>

²HYpothesis Scores to TOken Confidences

token from h_i is then increased³ by s_i/T . Here, s_i is the score of the hypothesis h_i and T is *temperature* — a hyperparameter of Hystoc.

Finally, we normalize the scores in each bin by applying softmax. The confidence of any word is then its probability within the corresponding bin.

If the outputs of the ASR system are to be rescored by any sort of external language model (LM), it is sufficient to incorporate the LM probabilities into the scores s_i of the hypotheses prior to the application of Hystoc.

2.1. Related work

Rover — a standard tool for fusing outputs from multiple ASR systems [11] — operates similar to Hystoc. However, Rover already expects word-level confidences in the input. When an insertion/deletion mismatch is detected, Rover introduces an ϵ -transition with a fixed confidence C_ϵ . Hystoc only assumes hypothesis level scores at the input and derives the scores of possible ϵ -transitions individually.

There is a multitude of methods for estimating word-level confidences for e2e ASR systems [13], but these typically include evaluating multiple models or training whole new models for confidence estimation. Hystoc operates on the outputs of a single system.

Already in 1996, Wessel et al. have observed that word-level confidences derived from N -best list are well correlated to the error rates of hybrid GMM-HMM systems [14]. Hystoc updates this observation for neural e2e systems and utilizes this information for effective system fusion.

Kišš has recently shown that hypothesis level scores are well correlated with segment-level error rate for character-level CTC systems in OCR and serve as an efficient filtration criterion for semi-supervised training [15]. Hystoc derives token-level confidences for a tokenization independent of the underlying neural model.

In parallel to our work, Huang et al. have shown that the same technique for computing confidences outperformed more sophisticated techniques in context of keyword spotting [16].

2.2. Hystoc fusion

The primary application aim of this work is to enable fusion of generic end-to-end systems, via the standard Rover tool. Nevertheless, the core procedure of Hystoc is unaware of the source of the individual hypotheses that are aligned together. Therefore, we can apply Hystoc confidence estimation to accumulation of hypotheses across systems, under the very mild condition that they are run on the same VAD segmentation. This way, a system fusion is performed. We hypothesise that this comes with the added benefit of also considering all the other hypotheses as opposed to only considering a 1-best output from each system.

We propose three strategies of incremental alignment:

Direct, where we take every hypothesis h_i^s along with its associated score s_i^s as-is and treat them as coming from a single system. This assumes that the scores coming from all systems have comparable dynamic range.

Normalized, where we first normalize the scores for each system s' separately so that $\sum_i \exp(s_i^{s'}) = 1$. Then, we again drop the notion of individual systems and align the hypotheses into the confusion network in the order given by their scores. This way, we

³In log-domain, i.e., by log-add-exp.

Table 1: Performance of the 1-best output of each individual system, measured by % WER.

System		w/o LM	w/ LM
ConfA	A	17.3	16.7
ConfB	B	17.7	17.1
RNN-T	R	21.7	21.2
Kaldi	K	22.1	20.0

roughly balance the scale of scores across systems, but allow more confident systems to impact the confusion network more by having their top hypotheses aligned into it first.

Normalized round-robin addresses this aspect and first introduces the top-1 hypothesis from each system, then the second best from each system and so on.

3. EXPERIMENTS

We demonstrate the quality of calibration of Hystoc confidence scores as well as their suitability for performing system fusion on the Albayzin 2022 challenge.

This speech recognition challenge comes with a collection of three audio databases: RTVE2018 and RTVE2020, consisting of recordings from various Spanish TV shows broadcasted between 2015 and 2019, and RTVE2020, comprised of audio material from historical recordings, popular broadcasted TV shows and fictional shows [17]. We followed the original data splits and used the RTVE2022’s dev partition with 2.5 hours of audio for development and RTVE2020’s test partition with 39 hours for cross-validation of acoustic models. The remaining 738 hours were used for training. Most of the provided training data are not human-revised, so we followed a data cleaning procedure from a previous work [18, Section 2.1], resulting in 512 hours of training recordings and 41 minutes of well-annotated dev data. The cross-validation set remained untouched in our experiments.

3.1. Single systems

Since the main evaluation method of this work is system fusion, we work with multiple ASR systems operating on a shared segmentation of the audio data. In total, we consider four different ASR models:

ConfA and **ConfB** XLS-R Conformer [19] systems, using XLS-R wav2vec features as input and predicting 1500 byte pair encoding (BPE) units. The Conformer model was trained from scratch on the training data, while the XLS-R feature extractor remained frozen.

RNN-T A recurrent transducer [20] based on CRDN encoder and GRU prediction network. Operates on 80-dimensional f-banks and predicts a vocabulary of 1000 BPE tokens.

Kaldi A hybrid TDNN-f/HMM system trained with the LF-MMI objective in Kaldi. A 3-gram language model is used for decoding, its optimal weight is 0.7.

Furthermore, we train an external language model (LM) based on the LSTM architecture [21], with two layers of 1500 units each. The LM operates on a 20k BPE vocabulary. We pretrained it on Spanish News Crawl monolingual texts prepared for WMT’13 challenge⁴ and fine-tuned it to the transcripts of the training data. The

⁴<https://www.statmt.org/wmt13/translation-task.html>

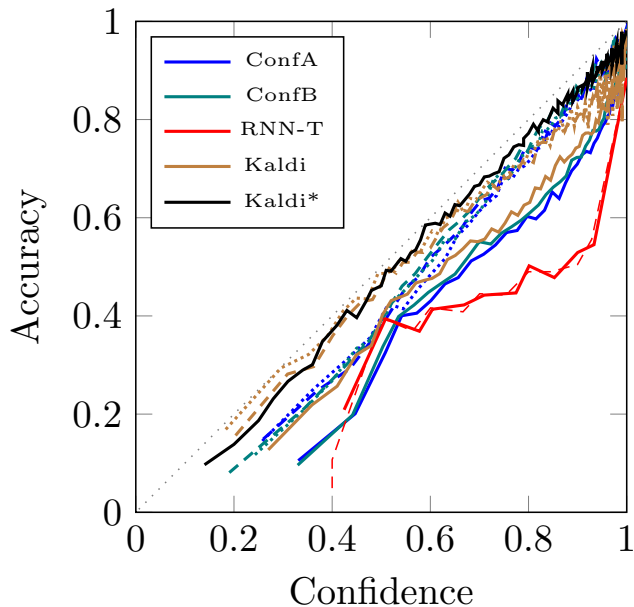


Fig. 2: Confidence as a predictor of accuracy. Dashed line corresponds to using temperature $T = 3.0$, dotted to $T = 10.0$. Kaldi* is the set of confidences coming directly from the original Kaldi lattices, i.e. without their reduction to 100-best. Each coordinate is computed from 2500 words.

LM is utilized by rescored the 100-best outputs from each system. When rescored the Kaldi system, we perform a log-linear interpolation between the scores from the LSTM-LM and the original 3-gram LM. When rescored the e2e systems, we add the LM probability to the ASR score in the log-domain and introduce a *token insertion bonus*, i.e. an increase in log-probability of a hypothesis proportional to its length as measured in number of LM tokens. The optimal values are in range $\langle 0.2, 0.3 \rangle$ for LM weight and $\langle 5.5, 6.5 \rangle$ for token insertion bonus. For more details, refer to the system description paper [22].

The performance of our ASR systems, both with and without LM rescored, is summarized in Table 1. Note that rescored of N -best lists improved WER of each system by 0.7% absolute on average.

3.2. Quality of confidences

With the ASR systems in place, we explore the quality of the confidence as a predictor of the word accuracy. To this end, we process the dev data with an ASR system and for each token, we record its confidence and whether it aligns correctly to the ground truth. Then, we sort the tokens by confidence and group them into batches of 2500 tokens. For each of these batches, we plot its median confidence and the proportion of correctly aligned tokens in the batch. Note that this analysis correctly takes into account substitution and insertion errors, but it does not capture deletions.

The results of this analysis are shown⁵ in Figure 2. With the exception of the RNN-T model, all confidences are in a close linear relationship with the word accuracies in the corresponding cohorts. With an elevated temperature $T = 3.0$, the dependency is

⁵For systems w/o LM rescored. The trends are very similar for systems with the rescored.

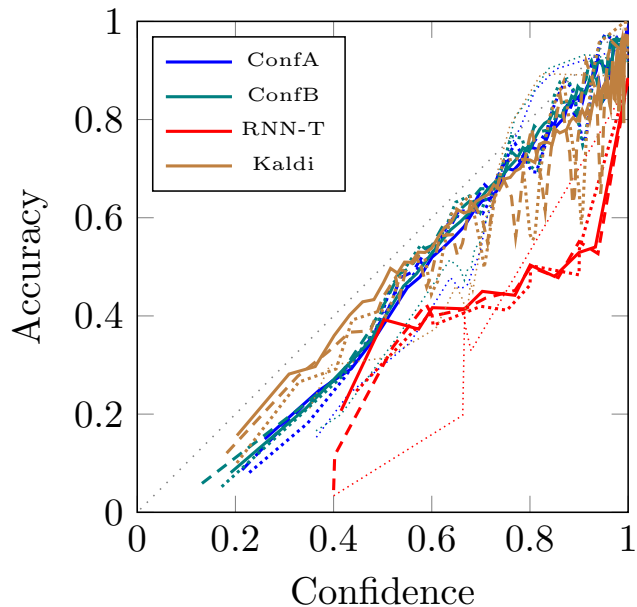


Fig. 3: Confidence quality as a function of number of top scoring hypotheses considered by Hystoc. Thinly dotted, dotted, dashed and full lines correspond respectively to 3-, 10-, 17- and 100-best outputs from each system.

even closer to the optimum. Increasing the temperature further to $T = 10.0$ brings no additional gain.

Increasing the temperature means giving more credibility to hypotheses which were originally scored as less likely, so we investigated the WER of the best path through these smoothed confusion networks. The change in WER caused by increasing T from 1.0 to 3.0 was smaller than 0.1% absolute and it was even negative in majority of cases. The one outlier was the ConfA system rescored with the external LM, where the WER increased by 0.2% absolute. Overall, we conclude that it is safe to increase the temperature if one desires more precise estimates of word accuracy.

The confidences obtained from the hybrid Kaldi system are better than traditionally expected (see e.g. [10]), which is likely the consequence of using the sequence-discriminative LF-MMI criterion. Comparing the confidences obtained from the directly from the full lattices (black curve) with those computed from the 100-best lists, we see that the degradation is not too sharp and is largely mitigated by increasing the temperature.

Building on this observation, we check how quality of the confidences changes when we further reduce the number of N -best outputs fed to Hystoc. Figure 3 shows that while applying Hystoc to 3-best is insufficient, the confidences are already reasonably reliable with as little as 10-best outputs from an ASR. In this case, the confidences of the e2e systems are actually better behaved than those from Kaldi outputs. Again, with the exception of the RNN-T system.

3.3. Advantage of confidences in Rover fusion

Observing that the obtained per-word confidences are reasonable predictors of word accuracy, we proceed to leverage them in fusion of multiple systems. To this end, we run the standard Rover fusion on transcriptions with Hystoc confidences. As the natural baseline,

Table 2: Performance of Rover fusion as a function of temperature T used to obtain the per-system confidences; results reported as % WER. Note that with $T = 0$, only the best path is considered from each system and the resulting confidences are then all equal to 1.0.

LM	Systems	T		
		0.0	1.0	3.0
no	A	17.3	17.3	17.3
	A + B	17.4	16.8	16.8
	A + B + K	16.7	16.2	16.2
	A + B + R	16.9	16.5	16.5
	A + B + K + R	16.1	15.9	16.0
	yes	A	16.7	16.7
yes	A + B	17.2	16.5	16.5
	A + B + K	16.4	15.4	15.7
	A + B + R	16.4	15.8	16.0
	A + B + K + R	15.8	15.3	15.6

we take the single-best hypothesis from each system and assign confidence 1.0 to each word in it. Formally, we achieve this effect by setting the temperature to $T = 0$. Motivated by the effect of temperature observed in the previous section, we introduce one more setup into this experiment, where we increase the temperature of each system correspondingly to 3.0. We tune the Rover fusion hyperparameters, i.e. confidence interpolation coefficient α and ϵ -confidence C_ϵ , in a range of $(0, 1)$.

The results of this experiment are summarized in Table 2. Providing the Hystoc confidences is beneficial in all cases, though the scale of the improvement varies. In general, the improvement is roughly on the scale of having one more system included in the fusion, i.e., the fusion of N systems with the Hystoc confidences is about as accurate as fusion of $N + 1$ systems without confidences. This is in line with the observations made in the Rover paper [11] for hybrid systems. Interestingly, the improvement from Hystoc confidences is better pronounced for systems with external LM rescoring.

While the optimal values of α were consistently around 0.55 for $T = 1.0$, they increased to about 0.75 for $T = 3.0$. This suggests that the fusion mechanism generally searches for a smoothed but informative confidence distributions.

3.4. Replacing Rover fusion with Hystoc

As described in Section 2.2, we continue to explore the possibility of fusing ASR outputs directly on level of aligning individual hypotheses. We fuse the same sets of models as in the previous experiment, making the results directly comparable. As the elevated temperature does not consistently improve the results, we use $T = 1.0$. With Hystoc fusion, there are no more hyperparameters to tune.

The results are shown in Table 3 and overall, they are considerably worse than in the Rover fusion. As long as we do not include the RNN-T, the fusion is at least slightly better than using the single best system. When RNN-T is included, the WERs deteriorate dramatically, esp. when LM rescoring is not applied. While we do not have a solid explanation why this happens for RNN-T in particular, it is in line with the observation that the confidences derived from RNN-T are not very well calibrated. This suggests that the confidence-accuracy analysis is a reasonable predictor of the usefulness of a system in the fusion.

Table 3: Performance of Hystoc fusion with different schemes of accumulating individual hypotheses, reported as % WER. “Norm-RR” stands for normalized round-robin.

LM	Systems	Scheme		
		Direct	Normalized	Norm-RR
no	A	17.3	17.3	17.3
	A + B	17.0	17.0	17.2
	A + B + K	17.0	17.1	17.3
	A + B + R	21.5	16.9	16.9
	A + B + K + R	21.6	19.6	19.5
	yes	A	16.7	16.7
yes	A + B	16.8	16.3	16.4
	A + B + K	16.7	16.5	16.5
	A + B + R	18.7	16.1	16.1
	A + B + K + R	18.9	19.2	19.2

One notable exception is the fusion of the two conformer systems with LM, where the Hystoc fusion yields slightly better results. While not significant on its own, this result does provide some hope for further improvements of fusion of similar systems.

4. CONCLUSION

In this paper, we proposed Hystoc, a straight-forward technique for computing word-level confidences from N -best lists produced by any end-to-end system. Hystoc is fully independent of both the tokenization and the exact objective formulation used by the underlying ASR model, making it a very generic tool.

We have shown that the confidences estimated by Hystoc are on-par with confidences obtained from lattices from a WFST-based decoder, and with as few as 10-best, the confidence estimation is reasonably stable. Practically, we have shown that Hystoc confidences bring a significant gain into the fusion of ASR systems, yielding 0.2 to 1% absolute improvement over a confidence-free voting scheme. Finally, we experimented with a direct Hystoc fusion of N -best lists from different systems. Here, the results are negative, we only obtained minor gains when fusing a specific combination of models, while for most other setups, the performance deteriorated.

We see two lines of further work on Hystoc: First, when aligning the individual hypotheses, Hystoc drops the notion of sequences immediately and treats tokens individually, allowing “hypothesis hopping” at any position. It could be beneficial to devise a more fine-grained approach where it would only be possible to switch between hypotheses in positions of high disagreement. Second, the iterative nature of hypothesis alignment in Hystoc introduces hard decisions early in the process of constructing the confusion network, possibly limiting their quality. We expect that doing a full multi-sequence alignment might mitigate this issue — however its direct computation is prohibitively expensive.

5. ACKNOWLEDGEMENTS

The work was supported by Czech Ministry of Culture NAKI III project JARIN (DH23P03OVV010) and Czech National Science Foundation (GACR) project NEUREM3 No. 19-26934X. Computing on IT4I supercomputer was supported by the Czech Ministry of Education, Youth and Sports through the e-INFRA CZ (ID:90140).

6. REFERENCES

- [1] William Chan, Navdeep Jaitly, Quoc Le, and Oriol Vinyals, “Listen, attend and spell: A neural network for large vocabulary conversational speech recognition,” in *2016 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2016, pp. 4960–4964.
- [2] Jason Li, Vitaly Lavrukhin, Boris Ginsburg, Ryan Leary, Oleksii Kuchaiev, Jonathan Cohen, Huyen Nguyen, and Ravi Gadde, “Jasper: An end-to-end convolutional neural acoustic model,” 09 2019, pp. 71–75.
- [3] Anmol Gulati, James Qin, Chung Cheng Chiu, Niki Parmar, Yu Zhang, Jiahui Yu, Wei Han, Shibo Wang, Zhengdong Zhang, Yonghui Wu, and Ruoming Pang, “Conformer: Convolution-augmented transformer for speech recognition,” *arXiv preprint arXiv:2005.08100*, 2020.
- [4] Alex Graves, “Sequence transduction with recurrent neural networks,” in *2012 Representation Learning Workshop, ICML*, 2012.
- [5] Alex Graves, Santiago Fernández, Faustino Gomez, and Jürgen Schmidhuber, “Connectionist temporal classification: Labelling unsegmented sequence data with recurrent neural networks,” in *Proceedings of the 23rd International Conference on Machine Learning*, New York, NY, USA, 2006, ICML ’06, p. 369–376, Association for Computing Machinery.
- [6] Mehryar Mohri, Fernando Pereira, and Michael Riley, “Weighted finite-state transducers in speech recognition,” *Computer Speech & Language*, vol. 16, pp. 69–88, 01 2002.
- [7] Yaniv Ovadia, Emily Fertig, Jie Ren, Zachary Nado, D. Sculley, Sebastian Nowozin, Joshua Dillon, Balaji Lakshminarayanan, and Jasper Snoek, “Can you trust your model’s uncertainty? evaluating predictive uncertainty under dataset shift,” in *Advances in Neural Information Processing Systems*, H. Wallach, H. Larochelle, A. Beygelzimer, F. d’Alché-Buc, E. Fox, and R. Garnett, Eds. 2019, vol. 32, Curran Associates, Inc.
- [8] Chuan Guo, Geoff Pleiss, Yu Sun, and Kilian Q. Weinberger, “On calibration of modern neural networks,” in *Proceedings of the 34th International Conference on Machine Learning - Volume 70*. 2017, ICML’17, p. 1321–1330, JMLR.org.
- [9] Matthias Sperber, Graham Neubig, Jan Niehues, and Alex Waibel, “Neural lattice-to-sequence models for uncertain inputs,” in *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing*, Copenhagen, Denmark, Sept. 2017, pp. 1380–1389, Association for Computational Linguistics.
- [10] Karel Veselý, Lukáš Burget, and Jan Černocký, “Semi-supervised dnn training with word selection for asr,” in *Proceedings of Interspeech 2017*. 2017, vol. 2017, pp. 3687–3691, International Speech Communication Association.
- [11] Jonathan G. Fiscus, “A post-processing system to yield reduced word error rates: Recognizer output voting error reduction (ROVER),” in *Proceedings of the IEEE Automatic Speech Recognition and Understanding Workshop (ASRU)*, Santa Barbara, 1997.
- [12] B. Bangalore, G. Bordel, and G. Riccardi, “Computing consensus translation from multiple machine translation systems,” in *IEEE Workshop on Automatic Speech Recognition and Understanding*, 2001. ASRU ’01., 2001, pp. 351–354.
- [13] Dan Oneata, Alexandru Caranica, Adriana Stan, and Horia Cucu, “An evaluation of word-level confidence estimation for end-to-end automatic speech recognition,” 01 2021, pp. 258–265.
- [14] Frank Wessel, R. Schluter, Klaus Macherey, and Hermann Ney, “Confidence measures for large vocabulary continuous speech recognition,” *Speech and Audio Processing, IEEE Transactions on*, vol. 9, pp. 288 – 298, 04 2001.
- [15] Martin Kišš, Karel Beneš, and Michal Hradiš, “At-st: Self-training adaptation strategy for ocr in domains with limited transcriptions,” in *Lladós J., Lopresti D., Uchida S. (eds) Document Analysis and Recognition - ICDAR 2021*. 2021, vol. 12824 of *Lecture Notes in Computer Science*, pp. 463–477, Springer Nature Switzerland AG.
- [16] Ruizhe Huang, Matthew Wiesner, Leibny Paola Garcia-Perera, Dan Povey, Jan Trmal, and Sanjeev Khudanpur, “Building keyword search system from end-to-end asr systems,” in *ICASSP 2023 - 2023 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2023, pp. 1–5.
- [17] Eduardo Lleida, Alfonso Ortega, Antonio Miguel, Virginia Bazán, Carmen Pérez, Manuel Gómez, and Alberto de Prada, “Rtve 2018, 2020 and 2022 database description,” .
- [18] Martin Kocour, Guillermo Cámara, Jordi Luque, David Bonet, Mireia Farrús, Martin Karafiát, Karel Veselý, and Jan Černocký, “BCN2BRNO: ASR System Fusion for Albayzin 2020 Speech to Text Challenge,” in *Proc. IberSPEECH 2021*, 2021, pp. 113–117.
- [19] Pengcheng Guo, Florian Boyer, Xuankai Chang, Tomoki Hayashi, Yosuke Higuchi, Hirofumi Inaguma, Naoyuki Kamo, Chenda Li, Daniel Garcia-Romero, Jiatong Shi, Jing Shi, Shinji Watanabe, Kun Wei, Wangyou Zhang, and Yuekai Zhang, “Recent developments on espnet toolkit boosted by conformer,” 2020.
- [20] Alex Graves, “Sequence transduction with recurrent neural networks,” *CoRR*, vol. abs/1211.3711, 2012.
- [21] Martin Sundermeyer, Ralf Schlüter, and Hermann Ney, “LSTM neural networks for language modeling,” in *Proc. Interspeech 2012*, 2012, pp. 194–197.
- [22] Martin Kocour, Jahnvi Umesh, Martin Karafiát, Ján Švec, Fernando Lopez, Karel Beneš, Mireia Sánchez Diez, Igor Szóke, Jordi Luque, Karel Veselý, Lukáš Burget, and Jan Černocký, “Bcn2brno: Asr system fusion for albayzin 2022 speech to text challenge,” in *Proceedings of IberSpeech 2022*. 2022, pp. 276–280, International Speech Communication Association.