

Supplementary Resources and Analysis for Automatic Speech Recognition Systems Trained on the Loquacious Dataset

Nick Rossenbach*[†], Robin Schmitt*[†], Tina Raissi*,
Simon Berger*[†], Larissa Kleppel*, Ralf Schlüter*[†]

*RWTH Aachen University, [†]AppTek.ai

Aachen, Germany

{lastname}@ml.rwth-aachen.de

Abstract

The recently published Loquacious dataset aims to be a replacement for established English automatic speech recognition (ASR) datasets such as LibriSpeech or TED-Lium. The main goal of the Loquacious dataset is to provide properly defined training and test partitions across many acoustic and language domains, with an open license suitable for both academia and industry. To further promote the benchmarking and usability of this new dataset, we present additional resources in the form of n -gram language models (LMs), a grapheme-to-phoneme (G2P) model and pronunciation lexica, with open and public access. Utilizing those additional resources we show experimental results across a wide range of ASR architectures with different label units and topologies. Our initial experimental results indicate that the Loquacious dataset offers a valuable case study for a variety of common challenges in ASR.

Keywords: Loquacious, language model, pronunciation lexicon, speech recognition, reference baselines

1. Introduction

In order to compare ASR systems across scientific publications, it is necessary to clearly define the training and testing conditions and data used. For this purpose, many benchmarks and corresponding datasets have been published in the past. Notable datasets for academic use covering English language are: Switchboard (Godfrey et al., 1992), LibriSpeech (Panayotov et al., 2015) or TED-Lium (Hernandez et al., 2018). Especially LibriSpeech, with more than 1700 citations in 2024 alone¹, is dominating ASR research. However, due to the increase in computation power, more and more research is performed without any constraints on the used data. This generally yields better results, but limits scientific comparability. On the other hand, existing datasets only provide a smaller amount of data, and are often not realistic in the diversity of audio and text conditions. Also, some datasets are not directly accessible, were removed from public access after publication², or have restricted licenses.

The Loquacious dataset (Parcollet et al., 2025) was published with the idea in mind to replace older datasets, offering multiple benefits. The important highlights are:

- A diverse selection of data sources containing both read and spontaneous speech
- Three different training subsets, consisting of 250, 2.5k, and 25k hours, with 250 and 2.5k having uniformly balanced data sources
- Defined dev and test sets for every data source

- A strong and consistent text normalization
- Open licensing for academia and industry

Such aspects make the Loquacious dataset a promising replacement for older standard datasets, with the 2.5k hours subset having a similar magnitude as the 960 hours of LibriSpeech. Moreover, due to similar text normalization, it is straightforward to use the dataset with existing LibriSpeech-oriented pipelines.

As the original release of Loquacious only included training and test data, we aim to extend the available material with count-based LMs, pronunciation lexica, and pre-trained G2P models. Such resources exist for LibriSpeech³ and other datasets, and are relevant for the comparison of ASR pipelines that require such resources. The introduction of a pronunciation lexicon allows for explicit decoder representations via lexical trees or finite state machines. These representations can integrate word-level and lexical information together with model specific label topologies (Mohri et al., 2002). This type of constraint could benefit the ASR performance, for example under domain mismatch conditions (Raissi et al., 2025). Moreover, the alignment of models using a blank-free label topology with phonemic units and silence modeling offers the advantage of providing word and phoneme boundaries unambiguously. Prior works also show that phoneme-based models are able to provide alignments with more accurate time information relative to the evidence in the signal, compared to the byte-pair encoding (BPE) based models (Rousso et al., 2024; Jiang et al., 2023).

¹according to Google Scholar, 15th Oct. 2025

²The TED terms of use disallow machine learning usage, existing datasets based on TED were taken offline

³<https://openslr.org/11/>

1.1. Contributions

With this work, we benchmark several ASR architectures with different label topologies using BPE and phoneme label units. We show the effect of different decoding methods, with open or closed vocabulary and optional use of an LM. For this purpose, we extend the Loquacious dataset with a pronunciation lexicon, a Sequitur (Bisani and Ney, 2008) G2P model and pre-trained n -gram count-based LMs. All resources are available as part of the official Loquacious repository on HuggingFace⁴. The scripts for creating the resources are available on Github⁵.

We provide results using lexicon-constrained and 4-gram LM decoding across many different ASR architectures for the 250 hours and 2.5k hours subsets of Loquacious. In addition, we include additional ablation studies regarding the pronunciation lexicon and the effect of pronunciation variants. As the original publication for Loquacious mostly contained results without data augmentation, we investigate the effect of SpecAugment (Park et al., 2019) and speed perturbation (Ko et al., 2015). Finally, we discuss additional insights we gained about this new dataset while working with it.

2. Language Modeling and Lexicon

Despite the prominence of neural (large) language models in ASR, count-based LMs still have their advantages as they are fast to train and need often only negligible computational resources. Just recently Nvidia explored GPU support for count-based LMs (Bataev et al., 2025). Count-based LMs are suitable for resource-efficient analysis of how well training data is matching the test data.

2.1. Vocabulary

Word-level LMs require a fixed vocabulary. We aim for a size that is similar to that of the official lexicon for LibriSpeech, which includes 200k words. In order to avoid an arbitrary cutoff, we choose the following process to determine the vocabulary:

1. We use the CMU Pronunciation Dictionary (CMUdict)⁶ 0.7b as the starting point for our vocab, which contains pronunciations for about 124k words, with 134k total pronunciations.
2. We select every word from the CMUdict that is appearing at least once in the full text of

the Loquacious *train-large* training data. This results in 95k words.

3. We then select every word that appears at least 4 times in the *train-large* training data, which gives us 197k words.
4. The union of both results in 216k words which we use as target vocabulary.

2.2. Pronunciation Lexicon

Characters or sub-word methods such as BPE (Sennrich et al., 2016) or SentencePieces via unigram LM (Kudo, 2018) nowadays are common for target label representations in ASR. To additionally allow for phoneme representations, we create phoneme pronunciations for all words in the vocabulary file as follows:

1. For all words that are part of the CMUdict, we take the existing entries (with one or more pronunciations) and remove stress markers.
2. We train a 5th-order G2P model with Sequitur using the full CMUdict as training data. For higher orders, we see no further improvement on our held-out cross-validation set.
3. We generate two pronunciation variants using Sequitur for all words which are not in the CMUdict.

For the last step, we either take only the most probable variant, or we select the second variant based on a probability threshold which is set to either 60% or 80%. We only consider a second pronunciation variant generated if the re-normalized posterior probability of the first variant is below this threshold.

2.3. Count-based Language Models

We use the KenLM toolkit (Heafield, 2011) to train ARPA-format count-based language models. We create LMs in three different settings:

1. 3rd-order with singleton pruning of 3-grams.
2. 4th-order with singleton pruning of 3-grams and 4-grams.
3. 4th-order without any pruning.

The perplexities and out-of-vocabulary (OOV) rates can be found in Table 1. As expected, every test-set has different perplexities, indicating both the complexity of the text and how much fitting training data is available within the full Loquacious dataset.

3. Automatic Speech Recognition

The original Loquacious publication only listed results for ASR systems that combine attention encoder-decoder (AED) (Chan et al., 2016) and connectionist temporal classification (CTC) (Graves et al., 2006). We use a set of different systems

⁴<https://huggingface.co/datasets/speechbrain/LoquaciousSet>

⁵<https://github.com/rwth-i6/LoquaciousAdditionalResources>

⁶<http://www.speech.cs.cmu.edu/cgi-bin/cmudict>

Table 1: Perplexities and OOV percentage of the different count-based LMs on the respective dev sets. All LMs are restricted to the 216k words vocabulary.

Language Model	n -gram count	Loquacious		Commonvoice		LibriSpeech		VoxPopuli		Yodas	
		OOV [%]	PPL	OOV [%]	PPL	OOV [%]	PPL	OOV [%]	PPL	OOV [%]	PPL
3-gram pruned	36M	0.58	222	1.25	361	0.58	231	0.13	154	0.24	201
4-gram pruned	54M	0.58	202	1.25	327	0.58	211	0.13	139	0.24	187
+ LibriSpeech	201M	0.58	196	1.27	352	0.58	149	0.13	173	0.24	243
4-gram unpruned	246M	0.58	193	1.25	311	0.58	195	0.13	135	0.24	182

with different training pipelines to verify our lexicon and n -gram LMs, as well as showing open-vocabulary results and other ablation studies. ASR systems using CTC loss only have an encoder with independent label outputs. Such systems tend to profit the most from additional LMs as they have no explicit language context modeling. For comparison, we include recurrent neural network transducer (RNN-T)-based (Graves, 2012) ASR systems. They provide a decoder structure for explicit language context modeling, so that additional language model information might be less relevant. To include a setup which properly supports pronunciation variants, we also include a factored hybrid (FH) model (Raissi et al., 2020). FH is a modern alternative to the classic hybrid neural network hidden Markov model (HMM) (Boullard and Morgan, 1993), being trained from scratch and thus avoiding both initial alignments from Gaussian mixture HMMs and decision-tree-based phone clustering. The diphone FH is trained similarly to CTC but offers a blank-free label topology with phonemic units and diphone context modeling. Finally, we include an AED system to be able to compare to the original Loquacious publication, even if in our case it does not support word level count-LM decoding.

3.1. General Settings

The ASR systems in this work are based on different software and settings, covering a broader range of different pipelines. All neural models are trained using RETURNN (Zeyer et al., 2018) or variants of it⁷. The FH system pipeline uses a TensorFlow backend, while all other models are based on PyTorch. All experiment pipelines are based on Sisyphus (Peter et al., 2018). Prefix-tree-based decoding is performed using Flashlight (Kahn et al., 2022) via the *torchaudio* interface for CTC and RASR (Rybach et al., 2011) for (m)RNN-T and FH. All systems use a 12-layer Conformer (Gulati et al., 2020) encoder with a hidden size of 512. We use 80-dimensional log-mel spectrograms with a window size of 25ms and a shift of 10ms including SpecAugment and speed perturbation, with

⁷<https://github.com/JackTemaki/MiniReturnn>

exception of FH that uses only SpecAugment. The features are down-sampled with two convolutional layers by a factor of six for RNN-T and AED and four for the other systems. We train each system for 100 epochs on the 250 hours *train.small* subset and for 40 epochs on the 2.5k hours *train.medium* subset. The training pipelines include different variants of Adam and one-cycle-learning-rate (OCLR) scheduling. We evaluate each model on the full dev and test sets, and in many cases report the word error rates (WER) individually for the 4 different test sets. For any system using an LM, we perform tuning of the LM and label prior correction scales on a fixed random subset of the full dev set. When using phonemes, we use an additional label set for end-of-word phonemes.

3.2. CTC

The CTC system has a total parameter count of 77M. For the BPE-based systems, three different decoding modes are used:

1. Simple greedy decoding by picking the position-wise maximum label
2. Time-synchronous beam-search.
3. Time-synchronous beam-search with lexical prefix tree and optional 4-gram LM.

Phoneme labels are only supported when using prefix-tree-based decoding.

3.3. Full-Context RNN-T

The standard RNN-T system is built analogously to the CTC system. The prediction network consists of a single LSTM layer with a hidden dimension of 512. The joiner network is a single linear mapping with a ReLU activation function and an output dimension of 640, yielding a total parameter count of 80M. We use two different decoding modes for the BPE version:

1. Time-synchronous beam-search.
2. Time-synchronous beam-search with lexical prefix tree and optional 4-gram LM including zero-encoder internal language model correction (Variani et al., 2020; Meng et al., 2021).

Table 2: Recognition results for BPE ASR systems that work without additional LMs or lexicon. Abbreviations for test sets: Commonvoice (CV), LibriSpeech (LS), VoxPopuli (VP), Yodas (YD).

Data	Architecture	Decoding	WER [%]					
			Loquacious		CV	LS	VP	YD
			dev	test	test			
<i>train.small</i> 250h	CTC	Greedy	17.2	18.6	30.1	16.7	12.5	21.1
	RNN-T		15.2	16.1	26.8	14.1	11.0	17.9
		Beam-Search	14.9	15.9	26.3	13.9	10.9	17.6
	AED	Greedy	15.3	16.7	28.0	14.7	10.8	20.0
		Beam-Search	14.7	16.0	26.8	14.2	10.4	18.0
<i>train.medium</i> 2500h	CTC	Greedy	10.6	11.5	19.3	9.0	8.1	15.4
	RNN-T		9.0	10.3	16.1	7.2	7.0	21.9
		Beam-Search	8.4	9.3	15.4	7.0	6.9	12.4
	AED	Greedy	8.7	10.0	16.5	7.5	7.0	15.1
		Beam-Search	8.4	9.5	15.9	7.3	7.0	13.1

3.4. Context-1 Monotonic RNN-T

The monotonic transducer (Tripathi et al., 2019) uses a similar encoder as the CTC and RNN-T models. It has a context limited to just one history label which is embedded with an embedding size of 256 and then forwarded through a feed-forward prediction network with 2 layers of size 640 and tanh activation. The joiner network consists of a single layer of size 1024, also with tanh activation. In total, this leads to a parameter count of around 79M. In addition to BPE labels the context-1 transducer also supports phoneme labels. The decoding options are identical to the full-context transducer.

3.5. Attention Encoder-Decoder

Our AED system uses the same features and encoder as our CTC system. For the decoder, we use a 6-layer Transformer with a dimension of 512. The model has 101M parameters and uses the same number of layers and hidden sizes as the 100M parameter model from the original Loquacious publication. The model is trained using cross-entropy loss with label smoothing and two CTC auxiliary losses on the 4th and 8th encoder layers. Decoding is done without an external LM using a custom implementation of label-synchronous beam search. This is different from the original Loquacious publication, which used joint AED/CTC decoding.

3.6. Factored Hybrid

Our diphone FH has an overall number of 75M parameters. Decoding is performed using time-synchronous beam search based on dynamic programming and lexical prefix trees. It is important to note that consistent use of pronunciation variants is only well-defined for this model, as it does not include a blank label. This is due to the fact that

the context dependency modeled on the static decoder structure, as done in the classic weighted finite state framework, is problematic when a blank label is introduced. In Section 4.5, we show the effect of adding pronunciation variants in decoding only.

4. Experimental Results

4.1. Baselines

We first evaluate CTC, RNN-T and AED models without the introduced lexicon and LMs for *train.small* and *train.medium*. Table 2 shows the results for each of the architectures with BPE labels and with greedy decoding or optimal beam search. Here, we do not include results of the original Loquacious publication, as they did not use any data augmentation, which is essential for good results on smaller datasets. See Section 4.6 for experiments on the effect of regularization techniques and a fair comparison to the original publication. We tested 128, 256, 512, and 1k labels for CTC and RNN-T and 1k, 2k, and 10k for AED. For CTC and RNN-T, 128 was always slightly better than larger sizes, while for AED 1k is optimal for *train.small* and 10k for *train.medium*. From the comparison of different ASR systems on the different test sets, it is clearly visible that it is important to have a task with various testing conditions and different training data scales. For systems trained on *train.medium*, the results on LibriSpeech and VoxPopuli are very homogeneous for AED and RNN-T. For Yodas and Commonvoice, there are larger differences. When switching from *small* to *medium*, doing beam search became more important for both RNN-T and AED on Yodas, but less relevant on LibriSpeech and VoxPopuli.

4.2. Count-based Language Model

We use the 216k words vocabulary from Section 2.1 and the pruned 4-gram LM from Section 2.3 to further improve on the baseline results as shown in Table 3. As expected using a lexical tree-based search with a language model yields much better results for models trained on *train.small*. Surprisingly, just performing vocabulary restricted search without a language model already improved the BPE-based model as well. We see such improvements for all systems trained on *train.small* as well as for CTC on *train.medium*. But even the context-dependent transducer models still showed slight improvements on the *train.medium*. When adding the 4-gram LM, all models largely improve. We also tested for a shift in results when we optimize the LM-scale and prior scale for a CTC not on the combined dev set, but on each of the subsets. For CTC trained on *train.small*, we observed that different scales yield an improvement from 17.2 to 16.6 on the Yodas test set.

Table 4 shows BPE-CTC ASR results using the four LMs, for which the statistics were presented in Table 1. Despite the difference in LM perplexities, we do not see a substantial difference in the total Loquacious WER, but only for the Yodas test set. When adding the official LibriSpeech LM data⁸, we only see improvements on the LibriSpeech test set. The other sets remain the same, or in case of Yodas, even slightly degrade. Given that the LM with LibriSpeech data is 20 times larger (6GB instead of 300MB) than the 3-gram, it does not really provide much additional benefit.

4.3. Phoneme-Lexicon

Using the CTC system, we compare the performance when replacing BPE labels with phonemes using the created pronunciation lexicon from Section 2.2. Table 5 shows the results of the phoneme-based CTC systems compared to the BPE labels. It can be seen that when training on *train.small* only the phonetic-based CTC system slightly outperforms its BPE-based counterpart. On *train.medium*, the phonetic-based system variant is slightly worse compared to its BPE counterpart for all architectures. Nevertheless, the performance gap is small enough so that the phonetic-based systems can be used for scenarios where this brings advantage, as e.g. for applications where custom pronunciation lexicon entries are desired, or for using systems with more explicit alignment modeling such as FH.

4.4. Restricted Vocabulary

To verify that it is actually necessary to expand the vocabulary and lexicon beyond what is given in

the CMUdict, we perform recognition with a lexicon and LM that was restricted to words in the CMUdict. Table 6 shows the results for the CTC model trained on *train.small*. There is a noticeable degradation, showing the justification for expanding the original CMUdict lexicon via G2P.

4.5. Pronunciation Variants

The CTC pipeline does not support pronunciation variants, so we use the FH pipeline to check if pronunciation variants generated by the G2P process are relevant or not. The lexica are defined as described in Section 2.2, so one without and two with additional G2P generated pronunciation variants. We train two models, one without seeing any pronunciation variants, and one without variants for words that are not in CMUdict. The recognition process either includes the lexicon that was used for training, or the augmented lexica with different probability masses. Table 7 shows that adding pronunciations during decoding helps, or does not lead to significant degradation. The ratio of words with G2P generated pronunciations in *dev.all* is only 0.6%, and only about 50% of those get a second pronunciation, so the potential of improving accuracy by adding more variants is rather limited. Moreover, the model trained on single pronunciations generally performs better. We believe this might be due to the additional alignment paths corresponding to the variants, which might affect the convergence of the model trained from scratch.

4.6. Regularization and Overfitting Behavior

The original Loquacious publication reported only results without data augmentation, with the exception of their largest system trained on *train.large*. We tested the effect of SpecAugment and speed perturbation, which are standard augmentations that were used for past LibriSpeech, TED-Lium and Switchboard systems. Table 8 shows the results for the AED and for the CTC system with 4-gram LM when disabling SpecAugment and speed perturbation. It is widely assumed that data augmentation is essential when working with rather small data sizes of a few hundred or thousand hours. We also observed substantial improvements for *train.medium*.

5. Further Analysis and Statistics

In this section, we want to share some of the insights we gained while working with Loquacious, which are meant to show that it presents interesting challenges which are not present in more artificial datasets like LibriSpeech.

⁸<https://www.openslr.org/11/>

Table 3: Recognition results for different BPE ASR systems to compare the effect of vocabulary restricted search and the addition of the pruned 4-gram LM. Abbreviations for test sets: Commonvoice (CV), LibriSpeech (LS), VoxPopuli (VP), Yodas (YD).

Data	Architecture	Vocab	LM	WER [%]					
				Loquacious		CV	LS	VP	YD
				dev	test	test			
<i>train.small</i> 250h	CTC	open	no	17.2	18.6	30.1	16.7	12.5	21.1
		closed	no	16.2	17.3	28.2	15.4	11.7	19.9
			yes	12.6	13.8	21.8	11.6	10.2	17.2
	RNN-T	open	no	14.9	15.9	26.3	13.9	10.9	17.6
		closed	no	14.6	15.5	25.7	13.6	10.7	17.2
			yes	13.1	14.1	22.9	11.9	10.1	16.5
	mRNN-T	open	no	15.4	16.5	27.0	14.6	11.2	18.6
		closed	no	14.7	15.7	25.6	13.8	10.8	18.2
			yes	12.2	13.4	21.5	11.1	9.7	16.7
<i>train.medium</i> 2500h	CTC	open	no	10.6	11.5	19.3	9.0	8.1	15.4
		closed	no	10.0	10.8	18.2	8.4	7.7	14.4
			yes	8.4	9.2	14.6	6.8	7.4	12.4
	RNN-T	open	no	8.5	9.3	15.4	7.1	7.0	12.6
		closed	no	8.4	9.2	15.3	7.0	6.9	12.5
			yes	8.0	8.8	14.0	6.4	6.8	14.0
	mRNN-T	open	no	9.1	10.2	16.3	7.3	7.3	18.8
		closed	no	8.9	10.0	15.8	7.1	7.2	18.7
			yes	8.0	8.8	14.3	6.4	7.1	12.0

Table 4: Recognition results for different n -gram LMs with a BPE-CTC on *train.medium*.

LM	WER[%]			
	Loquacious		LS	YD
	dev	test	test	
3-gram pruned	8.4	9.2	6.8	12.8
4-gram pruned	8.4	9.2	6.8	12.4
+ LibriSpeech	8.3	9.1	6.5	12.9
4-gram unpruned	8.3	9.1	6.8	12.0

5.1. Detailed Word Error Rates

Tables 9 and 10 show the detailed WERs for the LibriSpeech and Yodas dev sets for different models and label units trained on the *train.small* and *train.medium* sets, respectively. For AED, we report results for different vocabulary sizes because we observed much worse performance with the 10K vocabulary on *train.small*. For LibriSpeech, we observe that for a given model the amount of deletions and insertions lie in a similar range. Furthermore, these numbers are also similar across models. For Yodas however, we observe different error patterns for different configurations. For *train.small* (Table 9), all models have more deletions than insertions. For *train.medium* (Table 10), AED and RNN-T have more insertions than deletions, while CTC has more deletions than inser-

tions.

In order to better understand the error patterns, we manually compared the model outputs to the references. Regarding the insertion behavior of our models, our experience on LibriSpeech shows that they sometimes transcribe single words as multiple words (e.g. “await” → “a weight”) or, in case of AED, attend to the same audio segment multiple times, resulting in repeated words. On the Yodas dev set, we additionally observe “oscillations” (Frieske and Shi, 2024), i.e. cases where most of the model output is correct, but with the addition of a repeating n -gram. Most frequently, models insert a name at the beginning of an utterance (e.g. “would involve...” → “Matheus Aaron would involve...”), which is related to the name appearing in training transcripts without being present in the audio (see Section 5.2). Interestingly, we observe this issue for all models, except for a variant of FH for which we used chunked Viterbi training with a fixed target alignment.

Concerning the case of high deletions, we observe that all models omit large sections of some utterances in varying degrees. This behavior is most extreme in case of the CTC models, where almost whole utterances are sometimes omitted. Interestingly, the sequences where this happens are mostly disjoint between the AED model on the one side and the CTC and RNN-T models on the other side. Upon manual inspection, the affected utterances do not seem to be less intelligible than

Table 5: Recognition results for different ASR systems comparing BPE to phoneme performance. Results of FH are excluding speed perturbation. All results include decoding with the pruned 4-gram LM. Abbreviations for test sets: Commonvoice (CV), LibriSpeech (LS), VoxPopuli (VP), Yodas (YD).

Data	Architecture	Label	WER [%]						
			Loquacious		CV	LS	VP	YD	
			dev	test	test				
<i>train.small</i> 250h	CTC	BPE	12.6	13.8	21.8	11.6	10.2	17.2	
		Phon	12.3	13.4	21.0	11.2	10.3	15.9	
	mRNN-T	BPE	12.2	13.4	21.5	11.1	9.7	16.7	
		Phon	12.5	13.6	21.6	11.2	10.4	16.6	
	FH	Phon	BPE	12.5	13.5	22.0	11.2	10.2	15.4
			Phon	12.5	13.5	22.0	11.2	10.2	15.4
<i>train.medium</i> 2500h	CTC	BPE	8.4	9.2	14.6	6.8	7.4	12.4	
		Phon	8.8	9.8	15.2	7.4	7.8	13.7	
	mRNN-T	BPE	8.0	8.8	14.3	6.4	7.1	12.0	
		Phon	8.5	9.2	14.5	6.8	7.7	12.1	
	FH	Phon	BPE	9.0	9.7	15.2	7.4	7.9	13.0
			Phon	9.0	9.7	15.2	7.4	7.9	13.0

Table 6: Recognition results and OOV rates with the BPE-CTC using a restricted lexicon on *train.small*.

Test-Set	Lexicon			
	CMUdict		Full	
	OOV [%]	WER [%]	OOV [%]	WER [%]
Loq.	1.85	14.5	0.47	13.8
CV	2.53	22.9	0.99	21.8
LS	2.00	12.4	0.50	11.6
VP	1.33	10.6	0.11	10.2
YD	1.46	17.3	0.29	17.2

Table 7: Absolute errors for two FH models trained on *train.small* using different lexica. We show the effect of addition of pronunciation variants during recognition.

Pronunciation Variant		CV	LS	VP	YD
Train	Recog				
None	None	5257	5449	4340	867
	60% mass	5234	5439	4347	868
	80% mass	5190	5426	4336	898
CMUdict	CMUdict	5325	5567	4403	882
	60% mass	5324	5557	4403	875
	80% mass	5305	5551	4391	871

other utterances. Furthermore, all models have issues when speakers are spelling words letter by letter, in which case however the AED model succeeds more often than the other models.

Lastly, it is interesting to observe that with increased training data the “weaknesses” of the AED and phoneme CTC models, i.e. insertions and deletions, become more pronounced on Yodas but are halved for LibriSpeech (compare Table 9 and Table 10). For all models, except the BPE CTC,

Table 8: Recognition results in WER (%) on the *train.small* subset for removing SpecAugment and speed perturbation. *Original* refers to the 100M system in the original Loquacious publication.

Architecture	Spec-Aug.	Speed Pert.	Loquacious	
			dev	test
<i>Original</i> AED	no	no	22.3	23.8
			21.4	23.1
	yes	yes	15.1	16.6
			14.7	16.0
CTC + LM	no	no	19.1	20.6
	yes	no	12.9	13.9
		yes	12.6	13.8

the main driver for the reduction in WER on Yodas when switching to *train.medium* is the reduction in substitutions. For the BPE CTC, we even observe a small overall degradation with more training data, for which we don’t have an explanation yet.

5.2. Challenges of the Dataset

As noted in the original Loquacious publication, the dataset contains clean, read speech, as well as more noisy, spontaneous speech, and everything in between. Even though the authors employed measures like language identification to filter out erroneous segments, there are remaining aspects that make Loquacious more challenging compared to other academic tasks. In addition to the training data having transcriptions that do not match the audio in content or language, five of the Commonvoice test set utterances contained no speech even though they had non-empty transcriptions. We manually pickled 30 examples of *train.small* that a CTC model transcribed as empty. We found

Table 9: Detailed WERs in terms of substitutions, insertions, and deletions on the LibriSpeech and Yodas dev sets for different models trained on *train.small*. CTC, mRNN-T, and FH results are with 4-gram LM. In case of phonemes, we use end-of-word augmentation.

Architecture	Labels		WER [%]							
			LibriSpeech				Yodas			
	Tokens	Number	Sub	Del	Ins	Σ	Sub	Del	Ins	Σ
AED	BPE	1K	10.9	1.1	1.7	13.7	9.9	4.5	3.9	18.3
CTC		128	8.9	1.3	1.1	11.3	8.3	6.4	2.1	16.8
		Phonemes	80	8.4	1.0	1.3	10.6	8.1	5.1	2.4
RNN-T	BPE	128	11.3	1.3	1.3	13.9	9.6	5.4	3.5	18.5
mRNN-T			8.4	1.3	1.0	10.7	7.0	6.7	2.2	15.9
FH	Phonemes	80	8.3	1.3	1.2	10.8	7.9	6.1	2.3	16.3
			8.4	1.6	1.2	11.2	8.2	5.1	2.5	15.8

Table 10: Detailed WERs on the LibriSpeech and Yodas dev sets for different models trained on *train.medium*. CTC results are with 4-gram LM. In case of phonemes, we use end-of-word augmentation.

Architecture	Labels		WER [%]							
			LibriSpeech				Yodas			
	Tokens	Number	Sub	Del	Ins	Σ	Sub	Del	Ins	Σ
AED	BPE	10K	5.6	0.5	0.8	6.9	5.4	3.0	4.8	13.2
CTC		128	7.6	0.7	0.7	9.0	8.1	6.6	2.4	17.1
		Phonemes	80	5.7	0.6	0.8	7.0	5.8	6.0	2.3
RNN-T	BPE	128	5.7	0.6	0.7	7.0	5.5	3.3	3.8	12.6

that all of them contained speech, but 27 of them did not match the provided transcription at all and 3 turned out to be the Indonesian translations of the English speech. Lastly, we found that some transcripts in the training sets contain names at the beginning of the utterance that are not spoken. For example, the name “Matheus Aaron” is prepended to 10 utterances in *train.small* and 74 utterances in *train.medium*. We manually confirmed the absence of this name within the audio of the 10 utterances in *train.small*. Such aspects make the Loquacious dataset more realistic to work with, while still being well defined and normalized.

6. Future Work

Given the short time since the release of the Loquacious dataset, we focused only on the 250 hours and 2.5k hours subsets. This is fine in order to investigate the Loquacious dataset as a replacement for common smaller academic tasks such as LibriSpeech or TED-Lium. Still, it is important to expand such studies to the full 25k hours dataset. For this task, in order to investigate improvements with LM integration, much larger text corpora are needed which have to be normalized and standardized to fit the Loquacious task. Also, the current systems might improve with further optimization and tuning specifically for the Loquacious task.

7. Conclusion

In this work, we presented additional resources related to ASR training and evaluation with the Loquacious dataset. The following resources are published online under a permissive license:

- Vocabulary file containing 216k words
- Lexicon files containing 223k (no G2P variants) and 280k (with G2P variants) pronunciations
- Sequitur G2P model file
- A set of different ARPA count-based LMs trained with KenLM:
 - pruned 3-gram ARPA LM
 - pruned 4-gram ARPA LM
 - un-pruned 4-gram ARPA LM

The creation of such resources allowed us to conduct experiments with a large variety of different settings related to ASR architectures, decoding methods and label types. We found that on both the 250 hours *train.small* and 2.5k hours *train.medium* subsets, the use of a simple count-based LM substantially improves recognition results. More surprisingly, even just using a word vocabulary constrained search would yield improvements. We compared BPE to phoneme representations, showing that phonemes only outperform BPE for a CTC model trained on the 250 hours subset. Still, the phoneme-based systems had a comparable perfor-

mance, and for certain research or applications it might be beneficial to also explore such systems. We discussed some aspects that make Loquacious more interesting for research compared to earlier academic datasets or other alternatives. Finally, we presented first results with multiple architectures that can be used as reference for reasonable ASR performance on the *train.small* and *train.medium* sub-tasks.

8. Limitations

The biggest limitation of this work is that it excludes the 25k hours *train.large* subset of the Loquacious dataset. Working on such data requires a much larger amount of technical preparation, model sizes and training times, which were out of the scope for this work. In particular, we wanted to look at Loquacious as a replacement for smaller academic tasks, covering a variety of conditions instead of aiming for a single best result. What is excluded regarding the variety are further experiments regarding the model size, as all models are in a similar range of parameters between roughly 80M and 100M. While for the BPE-models the BPE sizes are approximately tuned to be optimal, there are some edge cases such as character representations which we did not consider yet. We were not able to consistently publish all the training recipes, which is planned for a further publication. Nevertheless, the continuous progress of our Loquacious pipelines can be found in our public pipeline repository⁹.

9. Acknowledgements

This work was partially supported by NeuroSys Phase 2, which as part of the initiative “Clusters4Future” is funded by the Federal Ministry of Research, Technology and Space (BMFTR) (funding IDs 03ZU2106DA and 03ZU2106DD), and by the project RESCALE within the program *AI Light-house Projects for the Environment, Climate, Nature and Resources* funded by the Federal Ministry for the Environment, Nature Conservation, Nuclear Safety and Consumer Protection (BMUV), funding ID: 67KI32006A. The authors gratefully acknowledge the computing time provided to them at the NHR Center NHR4CES at RWTH Aachen University (project number: p0023999). This is funded by the Federal Ministry of Education and Research, and the state governments participating on the basis of the resolutions of the GWK for national high performance computing at universities. We thank Titouan Parcollet for allowing us to host our resources as part of the official Loquacious release,

⁹https://github.com/rwth-i6/i6_experiments

and thank Mattia Di Gangi, Florian Lux and Moritz Gunz for critical proofreading.

10. Bibliographical References

- Vladimir Bataev, Andrei Andrusenko, Lilit Grigoryan, Aleksandr Laptev, Vitaly Lavrukhin, and Boris Ginsburg. 2025. [NGPU-LM: GPU-Accelerated N-Gram Language Model for Context-Biasing in Greedy ASR Decoding](#). In *Proc. Interspeech*, pages 644–648.
- Maximilian Bisani and Hermann Ney. 2008. [Joint-sequence models for grapheme-to-phoneme conversion](#). *Speech Communication*, 50(5):434–451.
- Hervé A. Bourlard and Nelson Morgan. 1993. *Connectionist Speech Recognition: a Hybrid Approach*. Kluwer Academic Publishers, Norwell, MA.
- William Chan, Navdeep Jaitly, Quoc Le, and Oriol Vinyals. 2016. [Listen, attend and spell: A neural network for large vocabulary conversational speech recognition](#). In *Proc. IEEE ICASSP*, pages 4960–4964.
- Rita Frieske and Bertram E. Shi. 2024. [Hallucinations in neural automatic speech recognition: Identifying errors and hallucinatory models](#). Preprint arXiv: 2401.01572.
- J.J. Godfrey, E.C. Holliman, and J. McDaniel. 1992. [Switchboard: telephone speech corpus for research and development](#). In *Proc. IEEE ICASSP*, volume 1, pages 517–520 vol.1.
- Alex Graves. 2012. [Sequence transduction with recurrent neural networks](#). In *RL 2012 - ICML Workshop on Representation Learning*.
- Alex Graves, Santiago Fernández, Faustino Gomez, and Jürgen Schmidhuber. 2006. [Connectionist temporal classification: labelling unsegmented sequence data with recurrent neural networks](#). ICML '06, page 369–376, New York, NY, USA.
- Anmol Gulati, James Qin, Chung-Cheng Chiu, Niki Parmar, Yu Zhang, Jiahui Yu, Wei Han, Shibo Wang, Zhengdong Zhang, Yonghui Wu, and Ruoming Pang. 2020. [Conformer: Convolution-augmented transformer for speech recognition](#). In *Proc. Interspeech*, pages 5036–5040.
- Kenneth Heafield. 2011. [KenLM: Faster and smaller language model queries](#). In *Proceedings of the Sixth Workshop on Statistical Machine Translation*, pages 187–197, Edinburgh, Scotland. Proc. ACL.

- François Hernandez, Vincent Nguyen, Sahar Ghannay, Natalia A. Tomashenko, and Yannick Estève. 2018. [TED-LIUM 3: Twice as much data and corpus repartition for experiments on speaker adaptation](#). In *Speech and Computer - 20th International Conference, SPECOM 2018, Leipzig, Germany, September 18-22, 2018, Proceedings*, volume 11096 of *Lecture Notes in Computer Science*, pages 198–208. Springer.
- Dongcheng Jiang, Chao Zhang, and Philip C Woodland. 2023. [A neural time alignment module for end-to-end automatic speech recognition](#). In *Proc. Interspeech*, pages 1374–1378.
- Jacob Kahn, Vineel Pratap, Tatiana Likhomanenko, Qiantong Xu, Awni Hannun, Jeff Cai, Paden Tomasello, Ann Lee, Edouard Grave, Gilad Aviodov, Benoit Steiner, Vitaliy Liptchinsky, Gabriel Synnaeve, and Ronan Collobert. 2022. [Flashlight: Enabling innovation in tools for machine learning](#). Preprint arXiv: 2201.12465.
- Tom Ko, Vijayaditya Peddinti, Daniel Povey, and Sanjeev Khudanpur. 2015. [Audio augmentation for speech recognition](#). In *Proc. Interspeech*, pages 3586–3589.
- Taku Kudo. 2018. [Subword regularization: Improving neural network translation models with multiple subword candidates](#). In *Proc. ACL*, pages 66–75, Melbourne, Australia.
- Zhong Meng, Sarangarajan Parthasarathy, Eric Sun, Yashesh Gaur, Naoyuki Kanda, Liang Lu, Xie Chen, Rui Zhao, Jinyu Li, and Yifan Gong. 2021. [Internal language model estimation for domain-adaptive end-to-end speech recognition](#). In *Proc. IEEE SLT*, pages 243–250.
- Mehryar Mohri, Fernando Pereira, and Michael Riley. 2002. Weighted finite-state transducers in speech recognition. *Computer Speech & Language*, 16(1):69–88.
- Vassil Panayotov, Guoguo Chen, Daniel Povey, and Sanjeev Khudanpur. 2015. [Librispeech: An asr corpus based on public domain audio books](#). In *Proc. IEEE ICASSP*, pages 5206–5210.
- Titouan Parcollet, Yuan Tseng, Shucong Zhang, and Rogier C. van Dalen. 2025. [Loquacious Set: 25,000 Hours of Transcribed and Diverse English Speech Recognition Data for Research and Commercial Use](#). In *Proc. Interspeech*, pages 4053–4057.
- Daniel S. Park, William Chan, Yu Zhang, Chung-Cheng Chiu, Barret Zoph, Ekin D. Cubuk, and Quoc V. Le. 2019. [SpecAugment: A simple data augmentation method for automatic speech recognition](#). Preprint arXiv: 1904.08779.
- Jan-Thorsten Peter, Eugen Beck, and Hermann Ney. 2018. [Sisyphus, a workflow manager designed for machine translation and automatic speech recognition](#). In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing, EMNLP*, pages 84–89.
- Tina Raissi, Eugen Beck, Ralf Schlüter, and Hermann Ney. 2020. [Context-dependent acoustic modeling without explicit phone clustering](#). In *Proc. Interspeech*, pages 4377–4381.
- Tina Raissi, Nick Rossenbach, and Ralf Schlüter. 2025. [Analysis of Domain Shift across ASR Architectures via TTS-Enabled Separation of Target Domain and Acoustic Conditions](#). In *Proc. IEEE ASRU*.
- Rotem Rousso, Eyal Cohen, Joseph Keshet, and Eleanor Chodroff. 2024. [Tradition or Innovation: A Comparison of Modern ASR Methods for Forced Alignment](#). In *Proc. Interspeech*, pages 1525–1529.
- David Rybach, Stefan Hahn, Patrick Lehnen, David Nolden, Martin Sundermeyer, Zoltan Tüske, Simon Wiesler, Ralf Schlüter, and Hermann Ney. 2011. RASR-the RWTH Aachen university open source speech recognition toolkit. In *Proc. IEEE ASRU*.
- Rico Sennrich, Barry Haddow, and Alexandra Birch. 2016. [Neural machine translation of rare words with subword units](#). In *Proc. ACL*, pages 1715–1725.
- Anshuman Tripathi, Han Lu, Hasim Sak, and Hagen Soltau. 2019. [Monotonic recurrent neural network transducer and decoding strategies](#). In *Proc. IEEE ASRU*, pages 944–948.
- Ehsan Variani, David Rybach, Cyril Allauzen, and Michael Riley. 2020. [Hybrid autoregressive transducer \(hat\)](#). In *Proc. IEEE ICASSP*, pages 6139–6143.
- Albert Zeyer, Tamer Alkhouli, and Hermann Ney. 2018. [RETURNN as a Generic Flexible Neural Toolkit with Application to Translation and Speech Recognition](#). In *Proc. ACL*, pages 128–133, Melbourne, Australia.