Gamli - Icelandic Oral History Corpus: Design, Collection and Evaluation

Anonymous Author
Affiliation / Address line 1
Affiliation / Address line 2
Affiliation / Address line 3
email@domain

Anonymouser Author
Affiliation / Address line 1
Affiliation / Address line 2
Affiliation / Address line 3
email@domain

Anonymousest Author
Affiliation / Address line 1
Affiliation / Address line 2
Affiliation / Address line 3
email@domain

Abstract

This paper presents *Gamli*, an ASR corpus for Icelandic oral histories, the first of its kind for this language, derived from the Ísmús ethnographic collection. Corpora for oral histories differ in various ways from corpora for general ASR, namely they contain spontaneous speech, multiple speakers per channel, noisy environments, the effects of historic recording equipment, and typically a large proportion of elderly speakers. Gamli contains 188 hours of aligned speech and transcripts, split into a training set and a test set. We describe our approach for creating the transcripts, through both Optical Character Recognition of previous transcripts and post-editing of ASR output. We also describe our approach for aligning, segmenting, and filtering the corpus and finally training a Kaldi ASR system, which achieves 22.1% word error rate (WER) on the Gamli test set, a substantial improvement from 53.4% word error rate from a baseline general ASR system for Icelandic.

1 Introduction

Icelandic open-licensed speech corpora have in recent years grown in volume and numbers, there are now *Talrómur* (Sigurgeirsson et al., 2021), *Málrómur* (Steingrímsson et al., 2017), *Samrómur* (Mollberg et al., 2020) and the Althingi's Parliamentary Speeches corpus (Helgadóttir et al., 2017b; Nikulásdóttir et al., 2018) to name a few. However both historical speech and older speakers are underrepresented in these corpora. For instance, regarding older speakers, in Samrómur, the largest open-licensed ASR corpus for Icelandic (2233 hours in the latest release (Hedström et al., 2022)), only 4.8% of speakers are over 60 years old.

Gamli, the oral history speech corpus presented in this paper differs from that in many ways. Firstly, it contains, predominantly, spontaneous speech in the form of interviews, secondly, it has a very high ratio of older speakers (94.8% of speakers are over 60 years old), thirdly, background noise is common as well as noise artefacts from historical recording equipment and lastly, historic dialects (word choice and accent) are much more prevalent than in existing corpora.

The corpus contains 188 hours of aligned speech and transcripts split into a training set and a test set. This data, based on valuable historical 20th century recordings stored at the Department of Ethnology and Folklore at The Árni Magnússon Institute for Icelandic Studies, is therefore an important addition to the existing Icelandic speech corpora.[1]

The custom ASR system presented in this paper along with the corpus will in due course be used to automatically transcribe all of the ethnographic audio recordings stored at the institute. The transcripts will then be made available on the online portal *Ísmús*[2] and paired with the respective recording.

2 Related Work

For many years, ASR systems have been trained on unaligned transcriptions (Panayotov et al., 2015) and even approximate transcriptions of spontaneous speech (Jang and Hauptmann, 1999). In the case of Icelandic ASR for spontaneous speech, there has been an ongoing project (Helgadóttir et al., 2017b), (Helgadóttir et al., 2017a) to align and filter Icelandic parliamentary transcripts for ASR in order to reduce the manual work

[1] The corpus is available under an open license at `https://anonymo.us/gamli`

[2] `www.ismus.is`

involved in transcribing parliamentary proceedings. Creating the corpora involves text normalization, time-alignment, and filtering utterances.

While ASR for oral histories is new for Icelandic, it is already being used in other languages. For example, the first large project was the MALACH project (Psutka et al., 2002) in 2002, where ASR transcriptions were used for indexing oral history archives and making them more searchable. However, some authors still consider oral history speech recognition an open problem (Picheny et al., 2019; Gref et al., 2020) and a recent study (Gref et al., 2022) found that human word error rate was 8.7% on a German oral history corpus (taking into account case-sensitivity and annotation of hesitations). Whereas (Lippmann, 1997) found a human word error rate of less than 4% on the Switchboard corpus of spontaneous telephony speech and less than 0.4% on the Wall Street Journal corpus of clear read speech. This suggests that the minimum possible word error rate for ASR might be much higher on oral histories than it is for cleaner speech corpora.

One other factor that makes oral history ASR an interesting challenge is the particularly high ratio of older speakers. It has been noted by (Vipperla et al., 2008) that for general ASR models, WER correlates strongly with age, even throughout a single speakers lifetime. This could be caused by multiple changes in aging voices, such as slower speaking rate, changes in F0 (decrease for males and increase for females), increase in jitter and shimmer (all from (Vipperla et al., 2008)), some of which could be mitigated by increasing the number of older speakers in the training set. However, other changes might not be so easily solved, such as a reduction of tongue and jaw strength and an increase in breathiness (all from (Vipperla et al., 2008)) which could reduce articulatory precision.

3 Origin of the corpus

The ethnography collection of the Department of Ethnology and Folklore at The Árni Magnússon Institute for Icelandic Studies contains more than 2,300 hours of audio recordings of oral heritage and traditions, with a little less than 2,500 interviewees. The oldest material are recordings made on wax cylinders in the early 20th century and the collection is continually expanding with new material being added every year.

The bulk of the collection, however, consists of recordings from the 1960's and 1970's, mainly the work of three collectors. Their focus was to gather ethnographic material from the whole country, first and foremost from older generations — the majority of the informants were born before or around the turn of the 20th century,

This resulted in an extensive collection of legends and fairy tales, accounts of beliefs and customs, poems, hymns, nursery rhymes, Icelandic ballads (*rímur*), occasional verses and more, with the material being variously spoken, sung or chanted. Apart from recited verse and that which is sung or chanted the speech is spontaneous. Accompanying the recordings is detailed metadata on the speaker, time and location of recording, as well as various other parameters such as genre (for different kinds of verse or prose material, e.g. poems or nursery rhymes, fairy tales or legends etc.), mode of performance (sung, chanted, spoken), key words, content (short summary, description), taletypes and motifs (in folktales and legends).

3.1 Speaker distribution in the collection

In their work the collectors mainly relied on a snowball method of sorts, asking speakers to point them to other possible informants, as well as contacting teachers or clergy to enquire about interesting subjects in their region. Speaker profession is often listed in the metadata, but there is no information about education, and most of the speakers were common people, i.e. workers, farmers, fishermen, housewives etc., with little formal education.

Gender was probably not a decisive factor at the outset and the total ratio is 57.6% male speakers and 42.4% female, i.e. based on the number of speakers. However, if audio length for each gender is included the difference increases quite a bit, i.e. 1504 hours (65%) for men vs. 821 hours (35%) for women.

As mentioned, the data in the collection also stands out in that that the age of the speakers is higher than in other existing Icelandic corpora. The oldest speaker in the collection was 105 years old at the time of recording in 1954 and the oldest speaker in the collection, with regards to date of birth, was born in 1827, and recorded in 1904 (not included in the *Gamli* corpus). In fact, 72.4% of the speakers are older than 63 and 31.4% are 71-80 years old. In *Gamli* this ratio is substantially higher, as detailed in Section 4.

3.2 Regional features in pronunciation

The speakers in the collection are from all over the country and therefore reflect the various regional differences in pronunciation much better than recently recorded speech corpora such as *Samrómur*, due to the fact that these regional features either have already more or less disappeared or are gradually disappearing. Amongst these features is for example the "hard" pronunciation of /p, t, k/ (still a distinct feature) and voiced pronunciation of /l, m, n/ before /p, t, k/ in North-Iceland, *rn*-, *rl*-pronunciation in South-East-Iceland, monophthongs before /ng, nk/ in the North-West etc.

While these features are not tagged in any way in the *Gamli* corpus, the ASR system trained on the corpus seems to prove well on these features, with possibly the exception of labial or velar stops before [ð], such as [hapðɪ] instead of [havðɪ] for *hafði* or [lakðɪ] instead of [laɣðɪ] for *lagði*. We have, however, not inspected this systematically, so it needs further looking into to state the precision with any certainty.

3.3 Recording procedure

Most of the recordings were made at the speakers' homes, in many cases in elderly homes, and carried out by the interviewer. It was not uncommon that other people, e.g. children, spouses etc. were present during the recording sessions, but they were in most cases not meant to play a part in the recording. Because of this, and for various other reasons, some background noise and disturbances occur in the recordings, e.g. children playing, traffic sounds, phones ringing etc., but these are generally not prominent.

Much of the recordings were recorded using high quality reel-to-reel tape recording devices, although some were done by amateurs who weren't as well equipped, whereas a part of the recordings are from the recording studios of The Icelandic National Broadcasting Service (Þorsteinsdóttir, 2013).

The digitalization of these recordings began in the late 1990's and continued into the early 2000's with the recordings being converted into WAV format as well as compressed MP3s for online use.

4 Corpus content

Gamli contains 188 hours of transcribed audio broken down into

1. $\sim$ 145 hours from optical character recognition (OCR) of previous transcriptions in various formats

2. $\sim$ 43 hours of new transcriptions (post-edited from ASR output)

The 145 hours include $\sim$ 8 hours defined as a test set, which was manually reviewed and corrected and annotated with speaker ID and time alignments in the annotation tool *ELAN*. The test set contains recordings with 10 speakers, 5 women (239 minutes) and 5 men (219 minutes), plus the interviewers (4 men) and serves for evaluating the system's performance.

A validation set has not been defined for the corpus as the acoustic model training in *Kaldi* (Povey et al., 2011) used a random sample of the training corpus for validation.

Data split	Hours	Male speakers	Female speakers	Total speakers
Training	180	115	85	200
Test	8	5	5	10

Table 1: Data splits in Gamli

4.1 Speaker distribution in the corpus

The corpus contains 210 unique speakers, 90 women and 120 men (plus the interviewers: 13 men and 1 woman). At the outset we aimed to have the gender ratio as equal as possible in the acoustic training data, but with three men surpassing 20 hours of speech each (with one topping at 29 hours) and accounting for more than one third of the entire data, that picture became quite distorted. As a result the gender bias in the corpus is even greater than in the collection itself, which is unfortunate, but simply reflects the data that was at hand, i.e. 73.5% vs. 26.5%, cf. Section 4.2.

The age ranges from 38 to 99, but most of the speakers are 60+ (94.8%), as shown in Figure 1, and the average age of the speakers is 77 years. This ratio is unprecedented in all existing corpora for Icelandic speech (cf. 4.8% in Samrómur as referred to in Section 1) and makes Gamli an important addition to that collection.

4.2 Corpus compilation

As mentioned, the largest part of the corpus, about 145 hours, stems from OCR of transcriptions at the Department of Ethnology and Folklore at The Árni Magnússon Institute for Icelandic Studies. These transcripts that were generated over several

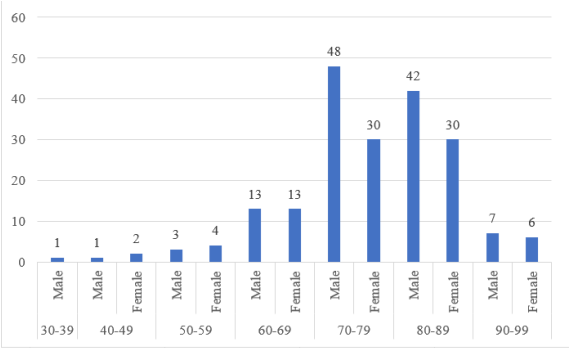

Figure 1: Age distribution of unique speakers in the training set

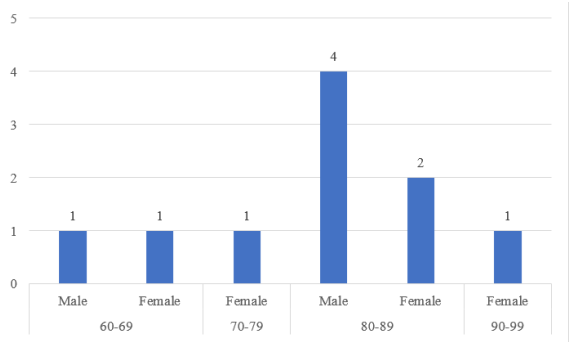

Figure 2: Age distribution of unique speakers in the test set

decades are not all in the same format (e.g. typewritten, dot printed, printed Word documents) and therefore needed first to be processed, i.e. scanned and OCRed (the results of which varied depending on the format). These transcripts were then catalogued and paired with the respective recordings.

Once this ready data had been processed the first ASR output was produced and manually corrected. During that process it became evident that some of the recordings were ill suited at this stage as they often contained poetry, nursery rhymes and in some cases singing, where the ASR system could not be expected to do well as the focus was on spontaneous speech, where it performed much better (cf. Section 6).

As a result, we made use of the detailed metadata search parameters in the Ísmús portal in order to filter the best in-domain data for further training. We mainly relied on the so-called *form* parameter (*genre*) to try to exclude everything but spontaneous speech. This gave much better results and resulted in the 43 hours of post-edited data mentioned in Section 4.

4.3 Normalizing, aligning, segmenting and filtering the transcripts for ASR training

A large part of the transcripts did not have time alignments and some had OCR spelling errors. Therefore, we had to process the utterances before using them to train the acoustic model. To do this, we first normalized all sentences using the Regina normalizer developed in (Sigurðardóttir, 2021) before aligning the transcripts to the audio and segmenting them. This step also removes sections with out-of-vocabulary words, which should account for errors stemming from the OCR.

We then filtered those segments, removing any that were deemed unintelligible to an intermediate ASR system. For this, a biased language model is applied to the segment, using words that appear in the utterance's transcription. It then removes segments where the system could not decode the words which appeared in the transcript. This is an iterative process, whereby an acoustic model is used to filter the training data, then that data is used to train a new acoustic model, which can then be used to re-align and re-filter the training data. These segmenting and filtering steps were all done with the Kaldi scripts (*Segment long utterances nn3*)[3] and (*Clean and segment data nn3*).[4]

5 Models (and out-of-domain data)

We trained a hybrid ASR system in Kaldi. That is, the language model and acoustic model were trained separately as opposed to an end-to-end system. For the acoustic and language models in the custom ASR system, we expanded the training sets with various out-of-domain data, which will be described in the following sections.

5.1 Acoustic Model

An acoustic model learns to map audio to a sequence of phonemes. The acoustic model is a TDNN (time-delayed neural network) chain model trained in Kaldi. It was trained on the in-domain data described above, but also on various out-of-domain data, which included the following datasets:

[3] https://github.com/kaldi-asr/kaldi/blob/master/egs/wsj/s5/steps/cleanup/segment_long_utterances_nnet3.sh

[4] https://github.com/kaldi-asr/kaldi/blob/master/egs/wsj/s5/steps/cleanup/clean_and_segment_data_nnet3.sh

1. Althingi's Parliamentary Speeches.[5] A corpus of 514.5 hours of recorded speech from the Icelandic parliament (Helgadóttir et al., 2017a)

2. 114.6 hours of speech from the first Samrómur release,[6] leaving out children.

3. 173.1 hours of unverified Samrómur data,[7] containing only speech with 50+ year old men and 60+ year old women.

4. 228.2 hours of the RÚV TV unknown speakers dataset.[8]

Data augmentation was also used to triple the entire training set. We added artificial noise and reverberation. For noisy data sets, e.g. call-center data sets, this is said to give better results than speed perturbations (Ko et al., 2017) and as was described earlier, background noise and disturbances are not uncommon in the data.

5.2 Language Model

A language model is necessary for outputting coherent texts, it learns a probability distribution for word sequences from a training corpus. The language model is an n-gram language model; 3-gram for decoding and 4-gram for rescoring. It was trained on in-domain data from the Gamli training set described in 4.2, both already existing ones and those resulting from the proofread ASR output. The out-of-domain data stems from the following sources:

1. The Icelandic Gigaword Corpus (IGC) (Steingrímsson et al., 2018). We use word forms from the 2022 version of the IGC.[9]

2. Ethnographic data from the National Museum of Iceland in *Sarpur*.[10]

3. Audio file descriptions from Ísmús [11] for their content.

4. Place name data from the Icelandic Place Name Collection.[12]

5.3 Vocabulary and Pronunciation Dictionary

The pronunciation dictionary maps words to sequences of phonemes. For the vocabulary we used:

1. All the word forms from *The Database of Icelandic Morphology* (Bjarnadóttir et al., 2019).

2. OOV words from audio file descriptions in Ísmús.

3. Vocabulary from the training set (only the data that was manually transcribed and not the OCR data); manually checked and added where appropriate.

4. OOV words from *Sarpur*; (manually checked and added where appropriate).

To get the phonemic transcriptions of each word a G2P model based on the *Icelandic Pronunciation Dictionary for Language Technology*[13] was used.

6 Evaluation

To assess the final ASR system's performance on the test set, we use Samrómur TDNN model as a baseline. This is a baseline model from a well-known dataset of read Icelandic speech. While the ASR baseline system, Samrómur achieved 53.4% WER on the Gamli test set, the final ASR system performed much better, achieving 22.1% WER on the same set, as shown in Table 2. This compares the two overall systems, each including their own acoustic model, language model, and vocabulary.

To investigate the differences in the two systems, we also compare the performance when taking demographic information into account in Figure 3. As stated earlier, the test set contains 10 speakers and a total of 8 hours of audio.

There appears to be a possible slight correlation between age and WER for the baseline system but not for the final system. Though it should be noted that the test set has too few data points to draw any significant conclusions. There is one outlier in the test set for both systems, an 85 year old man

[5] Available at: http://hdl.handle.net/20.500.12537/277
[6] Available at: http://hdl.handle.net/20.500.12537/189
[7] Available at: http://hdl.handle.net/20.500.12537/265
[8] Available at: http://hdl.handle.net/20.500.12537/191
[9] http://hdl.handle.net/20.500.12537/254
[10] https://sarpur.is/
[11] https://ismus.is/
[12] nafnid.is
[13] Available at: http://hdl.handle.net/20.500.12537/99

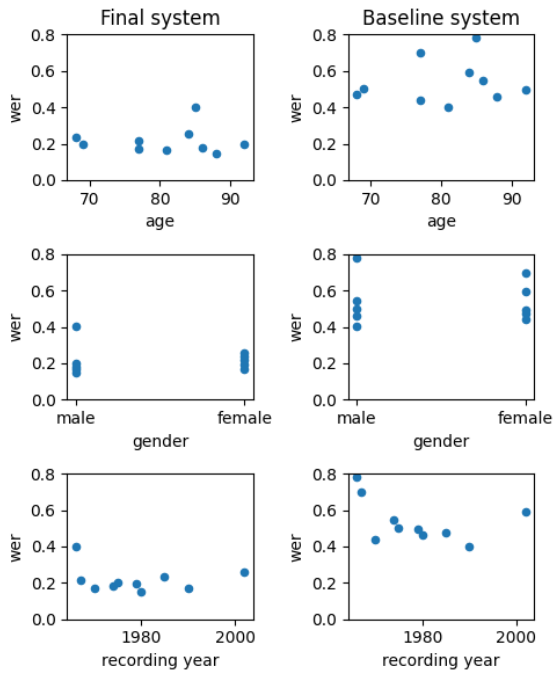

Figure 3: WER on the Gamli test set for the 10 unique speakers in the test set based on demographic information

recorded in 1966, upon manual inspection of the audio it seems the speaker has particularly slurred speech and there is some noise from the recording equipment.

	WER	OOV-rate total words	OOV-rate unique words
Baseline (Samrómur)	53.4%	1.1%	6.8%
Final	22.1%	0.5%	3.1%

Table 2: ASR performance on the Gamli oral history test set

7 Conclusion and Future Work

In this paper we have presented *Gamli*, a corpus suitable for training speech recognition systems, we have aligned and segmented Icelandic oral histories from manual transcriptions (both OCR from typewritten transcripts and post-edited from ASR output), and filtered out unintelligible segments.

We have described the compilation of the corpus, which has been published under an open license, the origins of the data and evaluation of an ASR system trained on the corpus. We have shown that using the corpus along with other relevant datasets can substantially lower WER for historical speech data, from 53.4% from a baseline model to 22.1%. We also draw the conclusion that it could be combined with other ASR training sets which lack in data from older speakers in order to reduce the word error rate for such speakers.

Our final ASR system will be used to automatically transcribe the entire ethnographic audio data stored in Ísmús, i.e. 2,300 hours of audio. We expect the outcome of that process to be in line with the results presented in this paper, with verse, nursery rhymes, singing etc. still remaining a challenge for the customised model, but accuracy for spontaneous speech to be more reliant on audio quality and clarity of speech. Where the quality of these two factors is high, we expect the system to perform well.

Even though the WER may differ substantially for some files, the general outcome will nonetheless be a somewhat readable version of the Ísmús ethnographic collection. That output can subsequently be used in a number of ways: making the data in Ísmús more accessible for the user, both laymen and researchers, indecing the archives for search queries (useful for longer audio files where the description can not do the entire content justice), and as a hypothesis transcript for post-editing of more transcripts.

The *Gamli* corpus itself should provide an interesting challenge to ASR researchers interested in spontaneous speech, older speakers, noisy audio, historical recordings and historical dialects.

References

Kristín Bjarnadóttir, Kristín Ingibjörg Hlynsdóttir, and Steinþór Steingrímsson. 2019. DIM: The Database of Icelandic Morphology. In *Proceedings of the 22nd Nordic Conference on Computational Linguistics, Turku, Finland*.

Michael Gref, Nike Matthiesen, Sreenivasa Hikkal Venugopala, Shalaka Satheesh, Aswinkumar Vijayananth, Duc Bach Ha, Sven Behnke, and Joachim Köhler. 2022. https://doi.org/10.48550/ARXIV.2201.06868 A study on the ambiguity in human annotation of german oral history interviews for perceived emotion recognition and sentiment analysis.

Michael Gref, Oliver Walter, Christoph Schmidt, Sven Behnke, and Joachim Köhler. 2020. Multi-staged cross-lingual acoustic model adaption for robust speech recognition in real-world applications–a case study on german oral history interviews. *arXiv preprint arXiv:2005.12562*.

Staffan Hedström, Judy Y. Fong, Ragnheiður Þórhallsdóttir, David Erik Mollberg, Smári Freyr Guðmundsson, Ólafur Helgi Jónsson, Sunneva Þorsteinsdóttir, Eydís Huld Magnúsdóttir, and Jon Gudnason. 2022. http://hdl.handle.net/20.500.12537/265 Samromur unverified 22.07. CLARIN-IS.

Inga Rún Helgadóttir, Róbert Kjaran, Anna Björk Nikulásdóttir, and Jón Guðnason. 2017a. http://hdl.handle.net/20.500.12537/277 Althingi's parliamentary speeches. CLARIN-IS.

Inga Rún Helgadóttir, Róbert Kjaran, Anna Björk Nikulásdóttir, and Jón Guðnason. 2017b. Building an asr corpus using althingi's parliamentary speeches. In *Interspeech*.

Photina Jaeyun Jang and Alexander G Hauptmann. 1999. Improving acoustic models with captioned multimedia speech. In *Proceedings IEEE International Conference on Multimedia Computing and Systems*, volume 2, pages 767–771. IEEE.

Tom Ko, Vijayaditya Peddinti, Daniel Povey, Michael L Seltzer, and Sanjeev Khudanpur. 2017. A study on data augmentation of reverberant speech for robust speech recognition. In *2017 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 5220–5224. IEEE.

Richard P. Lippmann. 1997. https://doi.org/https://doi.org/10.1016/S0167-6393(97)00021-6 Speech recognition by machines and humans. *Speech Communication*, 22(1):1–15.

David Erik Mollberg, Ólafur Helgi Jónsson, Sunneva THorsteinsdóttir, Steinþór Steingrímsson, Eydís Huld Magnúsdóttir, and Jón Guðnason. 2020. Samrómur: Crowd-sourcing data collection for icelandic speech recognition. In *International Conference on Language Resources and Evaluation*.

Anna B Nikulásdóttir, Inga R Helgadóttir, Matthías Pétursson, and Jón Guðnason. 2018. Open asr for icelandic: Resources and a baseline system. In *Proc. LREC*, volume 2018.

Vassil Panayotov, Guoguo Chen, Daniel Povey, and Sanjeev Khudanpur. 2015. https://doi.org/10.1109/ICASSP.2015.7178964 Librispeech: An asr corpus based on public domain audio books. In *2015 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 5206–5210.

Michael Picheny, Zóltan Tüske, Brian Kingsbury, Kartik Audhkhasi, Xiaodong Cui, and George Saon. 2019. Challenging the boundaries of speech recognition: The malach corpus. *arXiv preprint arXiv:1908.03455*.

Daniel Povey, Arnab Ghoshal, Gilles Boulianne, Lukas Burget, Ondrej Glembek, Nagendra Goel, Mirko Hannemann, Petr Motlicek, Yanmin Qian, Petr Schwarz, Jan Silovsky, Georg Stemmer, and Karel Vesely. 2011. The kaldi speech recognition toolkit. In *IEEE 2011 Workshop on Automatic Speech Recognition and Understanding*. IEEE Signal Processing Society. IEEE Catalog No.: CFP11SRW-USB.

Josef Psutka, Pavel Ircing, Josef V Psutka, Vlasta Radová, William J Byrne, Jan Hajič, Samuel Gustman, and Bhuvana Ramabhadran. 2002. Automatic transcription of czech language oral history in the malach project: Resources and initial experiments. In *Text, Speech and Dialogue: 5th International Conference, TSD 2002 Brno, Czech Republic, September 9–12, 2002 Proceedings 5*, pages 253–260. Springer.

Helga Svala Sigurðardóttir. 2021. http://hdl.handle.net/20.500.12537/158 Text normalization corpus 21.10 (2021-10-25). CLARIN-IS.

Atli Sigurgeirsson, Þorsteinn Gunnarsson, Gunnar Örnólfsson, Eydís Magnúsdóttir, Ragnheiður Þórhallsdóttir, Stefán Jónsson, and Jón Guðnason. 2021. https://aclanthology.org/2021.nodalida-main.50 Talrómur: A large Icelandic TTS corpus. In *Proceedings of the 23rd Nordic Conference on Computational Linguistics (NoDaLiDa)*, pages 440–444, Reykjavik, Iceland (Online). Linköping University Electronic Press, Sweden.

Steinþór Steingrímsson, Jón Guðnason, Sigrún Helgadóttir, and Eiríkur Rögnvaldsson. 2017. https://aclanthology.org/W17-0229 Málrómur: A manually verified corpus of recorded Icelandic speech. In *Proceedings of the 21st Nordic Conference on Computational Linguistics*, pages 237–240, Gothenburg, Sweden. Association for Computational Linguistics.

Steinþór Steingrímsson, Sigrún Helgadóttir, Eiríkur Rögnvaldsson, Starkaður Barkarson, and Jón Guðnason. 2018. https://aclanthology.org/L18-1690 Risamálheild: A very large Icelandic text corpus. In *Proceedings of the Eleventh International Conference on Language Resources and Evaluation (LREC 2018)*, Miyazaki, Japan. European Language Resources Association (ELRA).

Ravichander Vipperla, Steve Renals, and Joe Frankel. 2008. Longitudinal study of asr performance on ageing voices.

Rósa Þorsteinsdóttir. 2013. Ísmús (íslenskur músík- og menningararfur): An open-access database. *The Retrospective Methods Network Newsletter*, 7:97–101.