



USC: An Open-Source Uzbek Speech Corpus and Initial Speech Recognition Experiments

Muhammadjon Musaev¹, Saida Mussakhojayeveva², Ilyos Khujayorov¹, Yerbolat Khassanov², Mannon Ochilov¹, and Huseyin Atakan Varol²(✉)

¹ Computer Systems, Tashkent University of Information Technology named after Muhammad Al-Khwarizmi, Tashkent, Uzbekistan

² Institute of Smart Systems and Artificial Intelligence (ISSAI), Nazarbayev University, Nur-Sultan, Kazakhstan

{saida.mussakhojayeveva,yerbolat.khassanov,ahvarol}@nu.edu.kz

Abstract. We present a freely available speech corpus for the Uzbek language and report preliminary automatic speech recognition (ASR) results using both the deep neural network hidden Markov model (DNN-HMM) and end-to-end (E2E) architectures. The Uzbek speech corpus (USC) comprises 958 different speakers with a total of 105 h of transcribed audio recordings. To the best of our knowledge, this is the first open-source Uzbek speech corpus dedicated to the ASR task. To ensure high quality, the USC has been manually checked by native speakers. We first describe the design and development procedures of the USC, and then explain the conducted ASR experiments in detail. The experimental results demonstrate promising results for the applicability of the USC for ASR. Specifically, 18.1% and 17.4% word error rates were achieved on the validation and test sets, respectively. To enable experiment reproducibility, we share the USC dataset, pre-trained models, and training recipes in our GitHub repository (https://github.com/IS2AI/Uzbek_ASR).

Keywords: Dataset · Speech recognition · End-to-end · Transformer · Conformer · LSTM · DNN-HMM · Uzbek · Low-resource

1 Introduction

In this paper, we present an open-source Uzbek speech corpus (USC) dedicated to advancing automatic speech recognition (ASR) research for the Uzbek language. Uzbek is the official language of Uzbekistan also spoken in other neighboring countries, such as Afghanistan, Kazakhstan, Kyrgyzstan, Tajikistan, and Turkmenistan. It is an agglutinative language spoken by over 35 million people worldwide [2], which makes it the second-most widely spoken language in the Turkic languages family. With the USC, we aim to promote the development and usage of the Uzbek language in speech-enabled applications, such as message dictation, voice search, voice command, and other voice-controlled smart devices. We also believe that the USC will help to facilitate the development of

assistive technologies in the Uzbek language for people with special needs (e.g., the hearing impaired).

Previously, several works have addressed Uzbek speech recognition [21,23]. However, to the best of our knowledge, there has been no work presenting an open-source Uzbek speech corpus of high quality and sufficient size for training robust speech recognition systems. As a result, there is no generally accepted common Uzbek dataset, and thus, each research group conducts experiments and reports results on their internal data. This hinders experiment reproducibility and performance benchmarking, which retards the further development of Uzbek ASR technologies.

To address this problem, we created the USC dataset containing around 105 h of transcribed audio recordings spoken by 958 speakers from different regions and age groups. The USC is primarily designed for the ASR task, however, it can also be used to aid other speech-related tasks, such as speech synthesis and speech translation. To the best of our knowledge, the USC is the first open-source Uzbek speech corpus available for both academic and commercial use under the Creative Commons Attribution 4.0 International License¹. We expect that the USC will be a valuable resource for the general speech research community and become the baseline dataset for Uzbek ASR research. Therefore, we invite other researchers to use our dataset and help to further explore it with us.

To demonstrate the reliability of the USC, we conducted initial ASR experiments using both the hybrid deep neural network hidden Markov model (DNN-HMM) and end-to-end (E2E) architectures. Additionally, we investigated the impact of neural language models (LMs) and data augmentation techniques on the Uzbek speech recognition performance. In our experiments, the best DNN-HMM ASR system achieved 18.8% and 23.5% word error rates (WER) on the validation and test sets, respectively. The best E2E ASR system achieved 18.1% and 17.4% WERs on the validation and test sets, respectively. These results showcase the high quality of audios and transcripts in the USC.

The main contribution of this work is two-fold:

- We developed the first open-source speech corpus for the Uzbek language.
- We conducted initial Uzbek speech recognition experiments using both the conventional DNN-HMM and recently proposed E2E architectures.

The rest of the paper is organized as follows: Sect. 2 reviews past works on Uzbek speech recognition and datasets. Section 3 extensively describes the USC dataset construction procedures. The speech recognition experiments and obtained results are presented in Sect. 4. Lastly, Sect. 5 concludes the paper and points out directions of future work.

¹ <https://creativecommons.org/licenses/by/4.0/>.

2 Related Work

Speech is the most natural means of communication between humans, and researchers have long dreamed of employing it for interacting with machines. As a result, ASR research has attracted a great deal of attention over the past few decades [36]. In particular, various ASR architectures [4, 5, 12, 13] and annotated datasets [8, 25, 31] for training have been introduced. Unfortunately, most of the datasets are developed for popular languages such as English, Spanish, and Mandarin whereas less popular languages do not get much attention. Consequently, the less popular languages face an acute shortage of research and development of ASR technologies [7].

To address the aforementioned problem, many datasets have been developed in less popular languages. For example, to advance speech processing research in Kazakhstan, researchers developed open-source Kazakh speech corpora for building speech recognition [17] and speech synthesis [24] applications. To enable speech research and increase accessibility of speech-enabled applications for illiterate users, Doumbouya et al. [9] released 150 h of transcribed audio data for West African languages. Similarly, several large-scale multilingual speech corpora construction projects were initiated, e.g., VoxForge [3], Babel [10], M-AILABS [32], and Common Voice [6]. However, these projects do not include the Uzbek language yet.

In the context of the Uzbek language, some works have previously attempted Uzbek speech recognition. For example, Musaev et al. [23] developed an ASR system for geographical entities using a dataset consisting of 3,500 utterances. Similarly, the authors of [21] developed a read speech recognition system using 10 h of transcribed audio. The works of [20] and [22] addressed spoken digit and voice command recognition systems under the limited vocabulary scenarios, respectively. It should be mentioned that the datasets used in these works were very limited and specialized for narrow application domains. Other existing Uzbek datasets are prohibitively expensive or publicly unavailable [1]. Therefore, the development of an open-source Uzbek speech corpus of sufficient size is of paramount importance.

3 The USC Dataset Construction

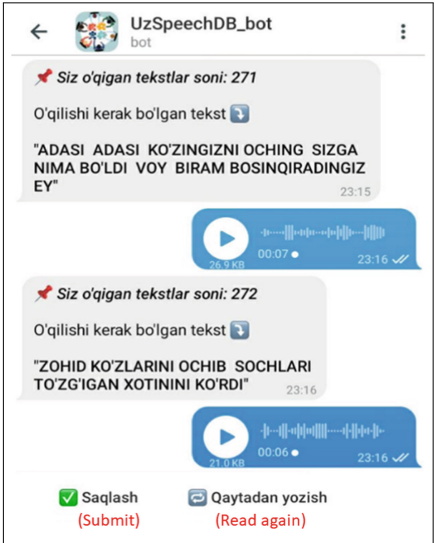
The Uzbek data collection project was conducted with the ethical approval of the Expert Committee consisting of members from the Tashkent University of Information Technology named after Muhammad Al-Khwarizmi. Each reader participated voluntarily and was informed of the data collection and use protocols. The dataset was collected by two means: crowdsourcing and audiobooks.

3.1 Crowdsourcing

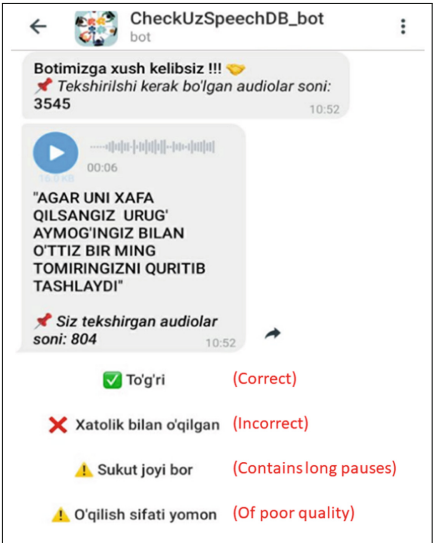
The crowdsourcing process consisted of three main stages—namely, text collection, text narration, and audio checking, which will be thoroughly described in the following sections.

Text Collection. We first collected Uzbek textual data from various sources including news portals, electronic books from modern Uzbek literature and national legislation database. The texts were collected automatically using web crawlers and they cover a wide range of topics such as politics, finance, entertainment, and law. In addition, we manually filtered the collected texts to eliminate defects peculiar to web crawlers and exclude non-Uzbek sentences and inappropriate content (e.g., user privacy and violence). We kept sentences containing borrowed words from other languages such as English. Lastly, we removed sentences containing numerals and sentences with more than 30 words. In total, over 100 thousand sentences were prepared for narration.

Text Narration. To narrate the collected sentences, we employed the Telegram [33] messaging platform, which is widely used in Uzbekistan. Specifically, we developed a Telegram bot that first presents a welcome message with instructions and then starts the narration process (see Fig. 1a). During the narration process, the bot sends a sentence to a reader and receives the corresponding audio recording. The bot allowed readers to listen to recorded audio and decide whether to submit or re-record it. In addition, the bot stored the reader IDs and other information including the age, gender, and geographical location. We attracted readers aged 18 or above by advertising the data collection project in social media, news, and open messaging communities on WhatsApp and Telegram.



(a) Data collection bot



(b) Data checking bot

Fig. 1. Examples of interaction with the Telegram bots during the (a) data collection and (b) data checking stages.

Audio Checking. To ensure the high quality of collected data, we developed an additional Telegram bot for checking the audio recordings. Different from the audio collection bot, the checker bot sends an audio recording and the corresponding sentence to an examiner (see Fig. 1b). As examiners, we recruited several volunteers among native Uzbek speakers. The examiners were instructed to inspect received audios with the sentences and mark them as “correct”, “incorrect”, “contains long pauses”, or “of poor quality”. Audio and sentence pairs marked as “correct” were added to the final speech corpus. For pairs marked as “incorrect”, the audio recording was removed, and the sentence was transferred to the audio collection bot for re-reading. For pairs marked as “contains long pauses” or “of poor quality”, we manually applied additional quality improvement procedures (e.g., trimming long pauses, splitting audio into several segments, and normalizing audio) and then added the pairs to the final speech corpus. To make our dataset close to the real-world scenarios, we kept utterances containing background noises.

3.2 Audiobooks

To collect data from audiobooks, we extracted freely available audiobooks narrated by 20 Uzbek audiobook narrators. From each book, we took only a 30-minute audio excerpt to balance the data contributed by each speaker. These excerpts were automatically segmented and aligned with the corresponding text by using the Aeneas Python library [30]. The generated segments were manually inspected and then added to the final speech corpus.

3.3 Dataset Statistics and Structure

The dataset statistics are reported in Table 1. In total, over 108,000 utterances were collected resulting in around 105 h of transcribed speech data. The utterance duration and length distributions are shown in Fig. 2. We split the dataset into training, validation, and test sets. The speakers in these sets are non-overlapping. For experiment reproducibility, we ask researchers planning to use our dataset to follow the provided splitting.

The USC dataset is structured as follows. We split the dataset into three folders corresponding to the training, validation, and test sets. Each folder contains

Table 1. The USC dataset specifications.

Category	Train	Valid	Test	Total
Duration (hours)	96.4	4.0	4.5	104.9
# Utterances	100,767	3,783	3,837	108,387
# Words	569.0k	22.5k	27.1k	618.6k
# Unique Words	59.5k	8.4k	10.5k	63.1k
# Speakers	879	41	38	958

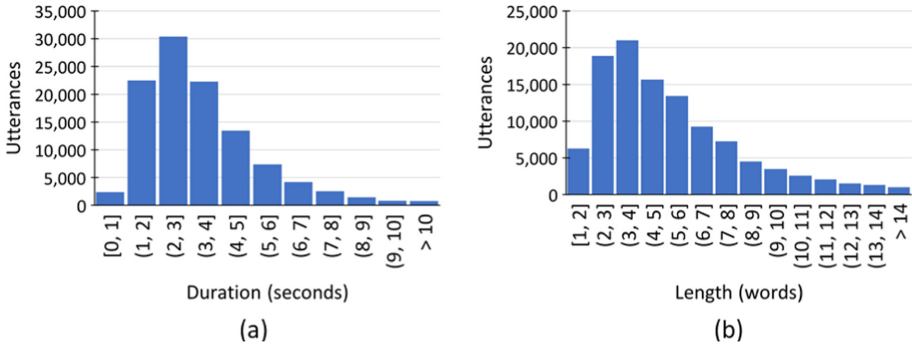


Fig. 2. Utterance (a) duration and (b) length distributions in the USC.

audio recordings and transcripts. The audio and corresponding transcription file names are the same, except that the audio recordings are stored as WAV files, whereas the transcriptions are stored as TXT files using the UTF-8 encoding. All the transcriptions are represented using the Uzbek Latin alphabet consisting of 29 letters and the apostrophe symbol.

4 Speech Recognition Experiments

We conducted speech recognition experiments to demonstrate the reliability of the USC dataset. We built both DNN-HMM and E2E speech recognition models using our dataset (see Sect. 3) and evaluated them using the character error rate (CER) and word error rate (WER) metrics. We did not use any external data and other available linguistic resources such as lexicon, pronunciation models, and vocabulary. Note that we left the detailed performance comparison of various ASR architectures for the Uzbek language as future work. Hence, in our experiments, we used the standard ASR architectures with the recommended specifications (i.e., number of encoder and decoder blocks, number of layers, layer dimensions, optimizer, initial learning rate, number of training epochs, and so on).

4.1 Experimental Setup

We trained all ASR models using the training set on a single V100 GPU running on the NVIDIA DGX-2 server. The hyper-parameters were tuned using the validation set, and the best-performing models were evaluated using the test set. The characteristics of the built DNN-HMM and E2E ASR systems are described in the following sections. For more information on the implementation details and hyper-parameter values, we refer the interested readers to our GitHub repository².

² <https://github.com/IS2AI/Uzbek-ASR>.

The DNN-HMM ASR. To build DNN-HMM ASR systems, we used the Kaldi framework [28] and followed the Wall Street Journal (WSJ) recipe. The acoustic model was constructed using the factorized time-delay neural networks (TDNN-F) [27] trained with the lattice-free maximum mutual information (LF-MMI) [29] training criterion. The inputs were Mel-frequency cepstral coefficients (MFCC) features with cepstral mean and variance normalization extracted every 10 ms over a 25 ms window. In addition, we applied data augmentation techniques based on the three-way speed perturbation [19] and spectral augmentation [26].

Due to the strong grapheme-to-phoneme relation in Uzbek, we employed a graphemic lexicon. The graphemic lexicon is comprised of 59.5k unique words extracted only from the training set. As a language model (LM), we used the Kneser-Ney smoothed 3-gram LM³ trained on the transcripts of the training set and with the vocabulary covering all words in the graphemic lexicon.

The E2E ASR. To build E2E ASR systems, we used the ESPnet framework [35] and followed the WSJ recipe. In particular, we built three types of E2E ASR architectures based on the 1) long short-term memory (LSTM) [16], 2) Transformer [34], and 3) Conformer [14] networks. All E2E ASR architectures were jointly trained with the connectionist temporal classification (CTC) [11] objective function under the multi-task learning framework [18]. The input speech features were represented as an 80-dimensional filterbank features with pitch computed every 10 ms over a 25 ms window. The output units were represented using 29 characters consisting of 26 letters⁴, the apostrophe symbol, and special tokens $\langle unk \rangle$ and $\langle space \rangle$. The batch size in all E2E ASR models was set to 64. To prevent overfitting, we applied data augmentation techniques based on speed perturbation and spectral augmentation. The results for the Transformer and Conformer based E2E ASR models are reported on the average model constructed using the last 10 checkpoints.

In addition, we built a character-level LSTM LM using the transcripts of the training set. The LSTM LM was constructed as a stack of two layers each with a memory cell size of 650. It was employed during the decoding stage using shallow fusion [15] for all the E2E architectures. For decoding, we set the beam size to 20 and the LSTM LM interpolation weight to 1 in all the E2E ASR models.

1) *E2E-LSTM ASR.* The LSTM-based E2E ASR was constructed using 3 encoder and 1 decoder blocks. Each encoder block consists of a bidirectional LSTM layer with 1,024 units per direction. The decoder block consists of a unidirectional LSTM layer with 1,024 units. The interpolation weight of the CTC objective was set to 0.5 and 0.3 for the training and decoding stages, respectively. The model was trained for 100 epochs using the Adadelta optimizer [37].

³ We trained several N-gram LMs with different orders and smoothing techniques and picked the one that obtained the best perplexity score on the validation set.

⁴ Note that the Uzbek Latin alphabet contains 29 letters, however, some of the letters are represented using digraphs (e.g., ng, sh, ch, o' and g'), which we broke down into smaller units and obtained 25 letters. The 26th letter is 'w' obtained from international words.

2) *E2E-Transformer ASR*. The Transformer-based E2E ASR was constructed using 12 encoder and 6 decoder blocks. We set the number of heads in the self-attention layer to 4 each with 256-dimension hidden states and the feed-forward network dimensions to 2,048. The interpolation weight for the CTC objective was set to 0.3 for both the training and decoding stages. The model was trained for 160 epochs using the Noam optimizer [34] with an initial learning rate of 10 and 25k warm-up steps. The dropout rate and label smoothing were set to 0.1.

3) *E2E-Conformer ASR*. The specifications of the Conformer-based E2E ASR are similar to the Transformer-based model. It was also constructed using 12 encoder and 6 decoder blocks with a similar number of attention heads and feed-forward network dimensions. However, the interpolation weight for the CTC objective was set to 0.2 and 0.3 for the training and decoding stages, respectively. The model was trained for 100 epochs using the Noam optimizer [34] with an initial learning rate of 5 and 25k warm-up steps. The dropout rate and label smoothing were set to 0.1.

4.2 Experiment Results

Table 2 presents the experiment results in terms of the CER and WER on the validation and test sets. All ASR models achieve competitive results. Specifically, the best result is achieved by the E2E-Conformer, followed by the E2E-Transformer, the DNN-HMM, and then the E2E-LSTM model. We observed

Table 2. The CER (%) and WER (%) results of different ASR models built using USC. The impact of language model (LM), speed perturbation (SP), and spectral augmentation (SA) are also reported.

Model	LM	SP	SA	Valid		Test	
				CER	WER	CER	WER
DNN-HMM	Yes	Yes	No	6.9	18.8	7.5	23.5
	Yes	Yes	Yes	6.9	19.9	8.1	24.9
E2E-LSTM	No	No	No	13.8	43.1	14.0	44.0
	Yes	No	No	14.9	30.0	14.3	31.4
	Yes	Yes	No	13.7	27.6	14.4	30.6
	Yes	Yes	Yes	12.6	24.9	12.0	27.0
E2E-Transformer	No	No	No	12.3	35.2	9.4	31.6
	Yes	No	No	11.7	25.7	8.7	23.9
	Yes	Yes	No	10.7	23.9	8.4	23.0
	Yes	Yes	Yes	9.9	21.4	7.6	21.0
E2E-Conformer	No	No	No	12.7	37.6	10.7	35.1
	Yes	No	No	11.5	27.5	9.7	26.3
	Yes	Yes	No	9.2	21.7	7.5	21.2
	Yes	Yes	Yes	7.8	18.1	5.8	17.4

that integrating LMs into E2E ASR is effective for the Uzbek language, where absolute WER improvements of 7.7%–12.6% are achieved on the test set. The application of speed perturbation to the E2E ASR models gains additional absolute WER improvements of 0.8%–5.1% on the test set. Spectral augmentation further improves the E2E ASR models by absolute WERs of 2.0%–3.8% on the test set, however, it does not improve the performance of the DNN-HMM model. Overall, the lowest WER results are 18.1% and 17.4% on the validation and test sets respectively, which were achieved by the E2E-Conformer. These results successfully demonstrate the utility of the USC dataset for training ASR models.

5 Conclusion

We developed an open-source Uzbek speech corpus containing around 105 h of transcribed audio recordings spoken by 958 speakers. The corpus was carefully checked by native speakers to ensure high quality. We believe that our corpus will further advance Uzbek speech processing research and become the primary dataset for comparing different ASR technologies among different research groups. In addition, we conducted preliminary ASR experiments using both the hybrid DNN-HMM and state-of-the-art E2E architectures. The best ASR model trained on our dataset achieved 18.1% and 17.4% WERs on the validation and test sets respectively, which demonstrates the reliability of the USC. In future work, we plan to further increase our dataset size and conduct additional ASR experiments.

References

1. Speechocean’s Uzbek speech corpus. <http://en.speechocean.com/datacenter/details/1847.html>. Accessed 21 May 2021
2. Uzbek language. https://en.wikipedia.org/wiki/Uzbek_language. Accessed 20 May 2021
3. Voxforge. <http://www.voxforge.org/>. Accessed 11 May 2021
4. Hinton, G., et al.: Deep neural networks for acoustic modeling in speech recognition: the shared views of four research groups. *IEEE Signal Process. Mag.* **29**(6), 82–97 (2012)
5. Abdel-Hamid, O., Mohamed, A., Jiang, H., Deng, L., Penn, G., Yu, D.: Convolutional neural networks for speech recognition. *IEEE ACM Trans. Audio Speech Lang. Process.* **22**(10), 1533–1545 (2014)
6. Ardila, R., et al.: Common voice: a massively-multilingual speech corpus. In: *LREC*, pp. 4218–4222. *ELRA* (2020)
7. Besacier, L., Barnard, E., Karpov, A., Schultz, T.: Automatic speech recognition for under-resourced languages: a survey. *Speech Commun.* **56**, 85–100 (2014)
8. Bu, H., Du, J., Na, X., Wu, B., Zheng, H.: AISHELL-1: an open-source Mandarin speech corpus and a speech recognition baseline. In: *Proceedings of the O-COCOSDA*, pp. 1–5. *IEEE* (2017)
9. Doumbouya, M., Einstein, L., Piech, C.: Using radio archives for low-resource speech recognition: towards an intelligent virtual assistant for illiterate users. In: *Proceedings of the AAAI Conference on Artificial Intelligence*, pp. 14757–14765. *AAAI Press* (2021)

10. Gales, M.J.F., Knill, K.M., Ragni, A., Rath, S.P.: Speech recognition and keyword spotting for low-resource languages: Babel project research at CUED. In: Proceedings of the Workshop on Spoken Language Technologies for Under-Resourced Languages (SLTU), pp. 16–23. ISCA (2014)
11. Graves, A., Fernández, S., Gomez, F.J., Schmidhuber, J.: Connectionist temporal classification: labelling unsegmented sequence data with recurrent neural networks. In: Proceedings of the International Conference on Machine Learning (ICML), Pittsburgh, Pennsylvania, USA, 25–29 June 2006, pp. 369–376 (2006)
12. Graves, A., Jaitly, N.: Towards end-to-end speech recognition with recurrent neural networks. In: Proceedings of the International Conference on Machine Learning (ICML), Beijing, China, 21–26 June 2014, vol. 32, pp. 1764–1772 (2014)
13. Graves, A., Mohamed, A., Hinton, G.E.: Speech recognition with deep recurrent neural networks. In: Proceedings of the IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP), pp. 6645–6649. IEEE (2013)
14. Gulati, A., et al.: Conformer: convolution-augmented transformer for speech recognition. In: Proceedings of INTERSPEECH, pp. 5036–5040. ISCA (2020)
15. Gülçehre, Ç., et al.: On using monolingual corpora in neural machine translation. CoRR abs/1503.03535 (2015). <http://arxiv.org/abs/1503.03535>
16. Hochreiter, S., Schmidhuber, J.: Long short-term memory. *Neural Comput.* **9**(8), 1735–1780 (1997)
17. Khassanov, Y., Mussakhojayeva, S., Mirzakhmetov, A., Adiyev, A., Nurpeiissov, M., Varol, H.A.: A crowdsourced open-source Kazakh speech corpus and initial speech recognition baseline. In: Proceedings of the Conference of the European Chapter of the Association for Computational Linguistics, pp. 697–706. Association for Computational Linguistics, April 2021
18. Kim, S., Hori, T., Watanabe, S.: Joint CTC-attention based end-to-end speech recognition using multi-task learning. In: Proceedings of the IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP), pp. 4835–4839 (2017)
19. Ko, T., Peddinti, V., Povey, D., Khudanpur, S.: Audio augmentation for speech recognition. In: Proceedings of INTERSPEECH, pp. 3586–3589 (2015)
20. Musaev, M., Khujayorov, I., Ochilov, M.: Image approach to speech recognition on CNN. In: Proceedings of the International Symposium on Computer Science and Intelligent Control (ISCSIC), pp. 57:1–57:6. ACM (2019)
21. Musaev, M., Khujayorov, I., Ochilov, M.: Development of integral model of speech recognition system for Uzbek language. In: Proceedings of the IEEE International Conference on Application of Information and Communication Technologies (AICT), pp. 1–6. IEEE (2020)
22. Musaev, M., Khujayorov, I., Ochilov, M.: The use of neural networks to improve the recognition accuracy of explosive and unvoiced phonemes in Uzbek language. In: Proceedings of the Information Communication Technologies Conference (ICTC), pp. 231–234. IEEE (2020)
23. Musaev, M., Khujayorov, I., Ochilov, M.: Automatic recognition of Uzbek speech based on integrated neural networks. In: Aliev, R.A., Yusupbekov, N.R., Kacprzyk, J., Pedrycz, W., Sadikoglu, F.M. (eds.) WCIS 2020. AISC, vol. 1323, pp. 215–223. Springer, Cham (2021). https://doi.org/10.1007/978-3-030-68004-6_28
24. Mussakhojayeva, S., Janaliyeva, A., Mirzakhmetov, A., Khassanov, Y., Varol, H.A.: KazakhTTS: an open-source Kazakh text-to-speech synthesis dataset. CoRR abs/2104.08459 (2021). <https://arxiv.org/abs/2104.08459>

25. Panayotov, V., Chen, G., Povey, D., Khudanpur, S.: Librispeech: an ASR corpus based on public domain audio books. In: Proceedings of the IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP), pp. 5206–5210. IEEE (2015)
26. Park, D.S., et al.: SpecAugment: a simple data augmentation method for automatic speech recognition. In: Proceedings of INTERSPEECH, pp. 2613–2617 (2019)
27. Povey, D., et al.: Semi-orthogonal low-rank matrix factorization for deep neural networks. In: Proceedings of INTERSPEECH, pp. 3743–3747 (2018)
28. Povey, D., et al.: The Kaldi speech recognition toolkit. In: Proceedings of the IEEE Workshop on Automatic Speech Recognition and Understanding (2011)
29. Povey, D., et al.: Purely sequence-trained neural networks for ASR based on lattice-free MMI. In: Proceedings of INTERSPEECH, pp. 2751–2755 (2016)
30. ReadBeyond: Aeneas. <https://www.readbeyond.it/aeneas/>
31. Hernandez, F., Nguyen, V., Ghannay, S., Tomashenko, N., Estève, Y.: TED-LIUM 3: twice as much data and corpus repartition for experiments on speaker adaptation. In: Karpov, A., Jokisch, O., Potapova, R. (eds.) SPECOM 2018. LNCS (LNAI), vol. 11096, pp. 198–208. Springer, Cham (2018). https://doi.org/10.1007/978-3-319-99579-3_21
32. Solak, I.: The M-AILABS speech dataset. <https://www.caito.de/2019/01/the-m-ailabs-speech-dataset/>. Accessed 11 May 2021
33. Telegram FZ LLC and Telegram Messenger Inc.: Telegram. <https://telegram.org>
34. Vaswani, A., et al.: Attention is all you need. In: Proceedings of the Annual Conference on Neural Information Processing Systems, 4–9 December 2017, Long Beach, CA, USA, pp. 5998–6008 (2017)
35. Watanabe, S., et al.: ESPnet: end-to-end speech processing toolkit. In: Proceedings of INTERSPEECH, Hyderabad, India, 2–6 September 2018, pp. 2207–2211 (2018)
36. Yu, D., Deng, L.: Automatic Speech Recognition. SCT, Springer, London (2015). <https://doi.org/10.1007/978-1-4471-5779-3>
37. Zeiler, M.D.: ADADELTA: an adaptive learning rate method. CoRR abs/1212.5701 (2012). <http://arxiv.org/abs/1212.5701>