

Grammar Based Identification Of Speaker Role For Improving ATCO And Pilot ASR

Amrutha Prasad^{1,2}, Juan Zuluaga-Gomez^{1,3}, Petr Motlicek¹, Oliver Ohneiser⁴, Hartmut Helmke⁴, Saeed Sarfjoo¹, Iuliia Nigmatulina¹

¹Idiap Research Institute, Martigny, Switzerland

²Brno University of Technology, Brno, Czechia

³Ecole Polytechnique Federale de Lausanne (EPFL), Switzerland

⁴German Aerospace Center (DLR), Institute of Flight Guidance, Braunschweig, Germany

{amrutha.prasad, juan-pablo.zuluaga, petr.motlicek}@idiap.ch,
{oliver.ohneiser, hartmut.helmke}@dlr.de, {saeed.sarfjoo, iuliia.nigmatulina}@idiap.ch

Abstract

Assistant Based Speech Recognition (ABSR) for air traffic control is generally trained by pooling both Air Traffic Controller (ATCO) and pilot data. In practice, this is motivated by the fact that the proportion of pilot data is lesser compared to ATCO while their standard language of communication is similar. However, due to data imbalance of ATCO and pilot and their varying acoustic conditions, the ASR performance is usually significantly better for ATCOs than pilots. In this paper, we propose to (1) split the ATCO and pilot data using an automatic approach exploiting ASR transcripts, and (2) consider ATCO and pilot ASR as two separate tasks for Acoustic Model (AM) training. For speaker role classification of ATCO and pilot data, a hypothesized ASR transcript is generated with a seed model, subsequently used to classify the speaker role based on the knowledge extracted from grammar defined by International Civil Aviation Organization (ICAO). This approach provides an average speaker role identification accuracy of 83% for ATCO and pilot. Finally, we show that training AMs separately for each task, or using a multitask approach is well suited for this data compared to AM trained by pooling all data.

Index Terms: assistant based speech recognition, air traffic management, multitask acoustic model, speaker classification

1. Introduction

Past research [1, 2] as part of the MALORCA¹ project focused on i) improving ABSR accuracy for ATCOs, ii) reducing workload for ATCOs [3], and iii) increasing efficiency [4] of ATCOs. As part of an ongoing HAAWAI² project, we aim to research and develop a reliable and adaptable solution to automatically transcribe voice commands issued by both ATCOs and pilots.

An error resilient and accurate ASR system is critical in the ATC domain. Current state-of-the-art technologies require large amounts of data to train ASR systems. Goal of another ongoing project called ATCO2³ is to collect large set of voice recordings of ATCOs and pilots (with a minimum effort) for the aforementioned purpose. In order to train ASR for this

task, ATCO and pilot speech recordings are usually pooled together [5, 6, 1] despite having a significant variability in the data distribution (acoustic and grammatical conditions) and the number of speakers in the data. As a result of the variability in the data distribution, ASR performance is significantly different if applied on ATCO or pilot speech (i.e. ATCO's speech is easier to recognize). Our baseline system trained by pooling all data reveals that the absolute difference in WER for ATCO and pilot is 9.7% (ATCO WER: 36.1%, Pilot WER: 45.8%).

In this paper, we hypothesize that instead of developing the ASR as a single task, ATCO and pilot ASR can be considered as two separate tasks [7]. Specifically, this paper investigates a multitask approach to train AMs to be integrated in ASR for ATCO and pilot. An obvious first step is to automatically split the ATC speech communications into two tasks (i.e. obtaining these speaker labels manually on a large dataset would be expensive and time consuming). A common approach is to use speaker diarization to classify the speakers in the audio [8, 9]. Although the ATCO speech is often cleaner than the pilot (as the former communicates from a controlled acoustic environment), the speech recordings collected in ATCO2 project using Very High Frequency (VHF) receivers are noisy for both ATCO and pilot channels. In such a case, the speaker diarization system may fail to assign speaker labels (ATCO or pilot) accurately. Thus, a speaker diarization system cannot be easily deployed to obtain accurate speaker labels.

The vital aspect in the air traffic management (ATM) environment is the communication between a controller and pilot. For the smooth travel of the aircraft this communication is well defined with a standard phraseology by ICAO [10]. Another approach to obtain the speaker class is through leveraging the 'ICAO' grammar to classify an utterance as one of the classes on the text level. Once the speaker labels (ATCO and pilot) are available for the large data, AMs can be trained for both controllers and pilots through different approaches. In this study, we show that due to the poor acoustic conditions training a single AM by pooling all data does not provide the best performance for pilots even if the speech is constrained by grammar. To obtain better performance accuracy, AM should be trained separately for ATCO and pilot data or considered as different tasks by using a multitask approach.

Section 2 provides a brief overview of a multitask ASR. The datasets used are described in Section 3 followed by Section 4 that describes speaker role classification with text. Section 5 explains the experimental setup and the results obtained which are followed by the conclusion in Section 6.

¹Machine Learning Of speech Recognition models for Controller Assistance: <http://www.malorca-project.de/wp/>

²Highly Advanced Air Traffic Controller Workstation with Artificial Intelligence Integration: <https://www.hawaii.de>

³Automatic collection and processing of voice data from air-traffic communications <https://www.atco2.org/>




Figure 1: Pipeline for gathering ATCO-pilot speech data with VHF receivers. Speech segments that do not match air-surveillance data (i.e. prior knowledge) are discarded.

2. Related work

Previous research [11, 12, 13, 14, 15] has shown that to compensate for limited data available in low-resourced languages, multilingual systems are an effective way to train ASR systems. In such a system, the output layer could be a separate layer for each language, or a single layer shared between all languages [15]. The Kaldi [16] toolkit provides state-of-the-art techniques to train AMs, specifically Lattice-Free Maximum Mutual Information (LF-MMI) [17]. Recently, [11] showed that multilingual AM can be trained with LF-MMI [17]. In MMI training, the cost function is given as:

$$\mathcal{F}_{\text{MMI}} = \sum_{u=1}^U \log \frac{p(\mathbf{x}^{(u)} | \mathcal{M}_{\mathbf{w}^{(u)}}(\boldsymbol{\theta})) p(\mathbf{w}^{(u)})}{p(\mathbf{x}^{(u)} | \mathcal{M}_{\text{den}}(\boldsymbol{\theta}))}, \quad (1)$$

where $\mathbf{x}^{(u)}$ is an input sequence for an utterance u , U is a set of all utterances in the training data, $\mathcal{M}_{\mathbf{w}^{(u)}}$ corresponds to a numerator graph specific to a word sequence in transcription, \mathcal{M}_{den} is a denominator graph modelling all possible sequences which is usually a phone LM, $\boldsymbol{\theta}$ is a model parameter and $p(\mathbf{w}^{(u)})$ is a language model probability for an utterance.

However, in multitask training with separate output layers, the cost function from Equation 1 is computed for each task depending on the number of tasks. For T tasks, the output cost function for each task t depends only on the utterances of that task:

$$\mathcal{F}_{\text{MMI}}^{(t)} = \sum_{u=1}^{U_t} \log \frac{p(\mathbf{x}^{(u)} | \mathcal{M}_{\mathbf{w}^{(u)}}^t(\boldsymbol{\theta})) p(\mathbf{w}^{(u)})}{p(\mathbf{x}^{(u)} | \mathcal{M}_{\text{den}}^t(\boldsymbol{\theta}))}, \quad (2)$$

where U_t is the number of utterances in a minibatch for a task t , $\boldsymbol{\theta}$ contains the shared and task-dependent parameters, $\mathcal{M}_{\mathbf{w}^{(u)}}^t$ and $\mathcal{M}_{\text{den}}^t$ are task-specific numerator and denominator graphs, respectively. For a task t , a denominator graph is built using the task-specific phone. For each minibatch, the gradient of each task output layer is computed and updated.

The overall cost-function is then given as a weighted sum of all task-dependent cost-functions defined in Equation 3.

$$\mathcal{F}_{\text{MMI}} = \sum_{t=1}^T \alpha_t \mathcal{F}_{\text{MMI}}^t, \quad (3)$$

where α_t is a task-dependent weight.

Although language and phone sets are the same for ATCO and pilots, due to the variation in the acoustic conditions, we consider them as different tasks and propose to use a multitask approach to train AMs. We hypothesize that using a multitask approach can lead to better ASR performance for both ATCOs and pilots compared to a single AM trained by combining all data.

3. Datasets

The following subsections provide an overview of the data used in this paper.

3.1. Collection and pre-processing of VHF data

3.1.1. Data collection

To obtain ATC voice communications the following two sources are considered: (i) open-source speech like LiveATC⁴, and ii) speech collected with our own setup of VHF receivers. In addition to speech data, the time-aligned metadata available is used to obtain the contextual information (e.g. call sign list for each utterance) from the OpenSky Network⁵ (OSN). This process yielded 377 hours of speech data from Prague (LKPR) and Brno (LKTB) airports from August 2020 until January 2021 for ATCO2 project.

3.1.2. Data pre-processing

Figure 1 shows the pipeline used for preparing the VHF database. First, a seed ASR system is used to produce the transcripts for the 377 hours of collected data. The seed model is a ‘hybrid’ speech-to-text recognizer based on Kaldi [16] trained with LF-MMI [17]. The neural network has six convolutional layers followed by nine Factorized Time-Delay Neural Network (TDNN-F) [18].

A list of call signs retrieved from OSN is in ICAO format. The ICAO format is composed of three characters airline code (e.g. TVS) followed by the call sign number which consists of digits and an additional character combination, e.g. TVS84J. In order to use this prior knowledge, this format is transformed into its ‘expanded version’. Several variants exist for a given call sign. As illustrated in Figure 1, the call sign TVS84J can be pronounced as ‘skytravel eight four juliett’ or instead each letter can be spelt out ‘tango victor sierra eight four juliett’.

Then, an ensemble of call signs with its variants are created. Finally, string matching of this expanded call sign list is applied to the automatic transcripts. The utterances in which one of the call signs is found are stored. This pre-processing reduced the data from 377 hours to 66 hours.

3.2. Related ATC datasets available for training

In addition to the above data collection, ATCO2 has brought together several air-traffic command-related databases [1, 19, 20, 21, 22, 23] from different publicly available open data sources. The full set of databases span approximately 140 hours of speech data that are strongly related in both phraseology and structure seen in ATCO-pilot communications [5, 6]. These databases were additionally augmented by adding noises that match LiveATC audio channels, doubling the size of training data. Since each of the seven databases had different annotation ontologies (annotation procedure, rules, and symbols), the transcripts had to be standardized and normalized [21, 24].

4. Speaker role classification with text

As described in Section 1, to develop a reliable and better performing ASR for both controllers and pilots, respective labelled speech data are required. However, in most cases, e.g.

⁴LiveATC.net is a streaming audio network consisting of local receivers tuned to aircraft communications: <https://www.liveatc.net/>

⁵OpenSky Network: provides open access of real-world air traffic control data to the public




Figure 2: Speaker role identification based on grammar for VHF data. The same procedure is applied to other datasets used in this paper.

such as in ATCO2 project, although large amounts of data are collected, they do not contain speaker labels. The first task is therefore to split the speech recordings into two classes: ATCO and pilot. To accomplish this, we extract the information based on the ICAO grammar to identify the speaker’s role.

ICAO defines a separate grammar for ATCOs and pilots to enable clear communication. For instance, there are certain phrases/commands that an ATCO should use in a specific order. This knowledge is used to extract/identify potential words/commands that indicate a specific role of speaker. For example, the words such as "identified", "approved", "wind" would most probably only be spoken by an ATCO and the words "wilco", "maintaining", "we", "our" would probably be spoken only by a pilot. Currently we have made a list of 25 words for ATCO and 9 words for pilot that indicate each role. This list was generated by manual curation and expert feedback. A list of callsigns⁶ is also prepared from available airline codes.

Since this method operates at word level, manual (if available) or automatically generated transcripts are required for the corresponding speech recordings. In order to identify if an utterance is spoken by an ATCO or a pilot, we check the corresponding transcript for the conditions below: if the callsign appears at the beginning of an utterance, this utterance is classified as ATCO, else it is classified as a pilot. As there is greeting at the beginning quite often, we check if the callsign appears within the first four words. If one of the words in the utterance is in the list of ATCO words or in the list of pilot words, then the respective role is assigned.

Once each utterance in the training data is classified as ATCO or pilot, we propose to train two versions of ASR. In the first system there are two acoustic models: one for ATCO and one for pilot. In the second system we train a multitask network with one task as ATCO ASR and other as pilot ASR. The procedure is illustrated in Figure 2.

4.1. Speaker Role Classification Performance

This method has been tested on manually speaker segmented and transcribed data from two different Air Navigation Service Providers (ANSPs) as a part of the HAAWAI project: i) NATS for London Approach and ii) ISAVIA for Icelandic en-route. In the first set, there are 1060 ATCO utterances and 1280 pilot utterances. From the confusion matrix shown in Figure 3, we can observe that this method provides a true positive rate (TPR) of 81% (correctly classified ATCO) and true negative rate (TNR) of 85% (correctly classified pilot). The second set used consists of 775 ATCO utterances and 887 pilot utterances. From the confusion matrix shown in Figure 4, we see that this method provides a TPR of 85% and TNR of 80%.

⁶https://en.wikipedia.org/wiki/List_of_airline_codes

Predicted Class	Actual	
	ATCO	Pilot
ATCO	856 81%	188 15%
Pilot	204 19%	1092 85%

Figure 3: Confusion matrix for speaker role identification based on text for manually speaker segmented data for London Approach. Total number of ATCO utterances are 1060 and the total number of pilot utterances are 1280.

Predicted Class	Actual	
	ATCO	Pilot
ATCO	660 85%	179 20%
Pilot	115 15%	708 80%

Figure 4: Confusion matrix for speaker role identification based on text for manually speaker segmented data for Icelandic en-route. Total number of ATCO utterances are 775 and the total number of pilot utterances are 887.

4.2. Error Analysis

As there exists many variants for any given callsign, checking only for the airline code (e.g. lufthansa) is a major factor contributing to the misclassification of ATCO as pilot. A reason for the misclassification of pilot as ATCO is the occurrence of callsigns at the beginning of the utterance. Analysis of misclassification errors show that the accuracy can be improved by i) matching the callsign spoken with its allowed variants (e.g. LUF189AF → lufthansa one eight nine alfa foxtrot, one eight nine alfa foxtrot, etc) and ii) using the context prior to the callsigns (e.g. the pilot may mention the place of the control they want to communicate followed by the callsign). We will consider applying the aforementioned improvements as a part of our future work.

5. Experiments

For all our experiments, conventional biphone Convolutional Neural Network (CNN) [25] + TDNN-F [18] based acoustic models trained with Kaldi [16] toolkit (i.e. nnet3 model architecture) is used. AMs are trained with the LF-MMI [17] training framework considered to produce state-of-the-art performance for hybrid ASR systems. In all the experiments, 3-fold speed perturbation [26] and i-vectors are used. The multi-task training script used can be found in Kaldi [16]⁷. The value of the task dependent weight α_t used in our experiments is 0.5. Language model (LM) is trained with all the manual transcripts available from datasets described in Section 3.2 and used for all

⁷[egs/babel/multilang/s5d/local/chain2/run_tdnm.sh](https://github.com/kaldi-asr/kaldi/blob/master/egs/babel/multilang/s5d/local/chain2/run_tdnm.sh)

Table 1: WER comparison for AMs trained with data from other ATC datasets and tested on LiveATC ATCO and pilot test sets.

Model	WER %	
	ATCO test	Pilot test
Clean	36.9	47.7
Noise	31.3	41.1
Combined	36.1	45.8
Multitask	31.6	41.1

Table 2: WER comparison for models trained with only the data collected from VHF receivers and tested on LiveATC ATCO and pilot test sets.

Model	WER %	
	ATCO test	Pilot test
VHF ATCO	43.2	51.6
VHF Pilot	40.3	45
Combined	46	50
Multitask	38.2	44

the experiments.

The performance of different models is evaluated on LiveATC test set with the Word Error Rate (WER) metric which is based on the Levenshtein distance at the word level. The total duration of the test set is 1h 50 mins. The set is split into two subsets: ATCO set (52 mins) and Pilot set (58 mins).

In each group of experiments, results are given for i) AM trained for each task separately, ii) AM trained by combining all data and iii) AM trained with multitask learning.

5.1. Experiments on ATC databases

In this setup, we use data from the ATC databases mentioned in Section 3.2 as Clean data and its noise augmented part as Noise data. As shown in Table 1, both ATCO and pilot test sets provide better performance when the model is trained with Noise data compared to the model trained with only Clean data. This shows that the noise augmented version of the clean data matches with the test sets much better than the clean version. Moreover, the Combined system performs significantly worse than the Noise system. This shows that using the Clean dataset in fact hurts ASR performance. This is one of the reasons why the multitask system performs only on par with the Noise system. Therefore only the noise augmented data is used for training in the next experiments.

5.2. Experiments on VHF data

Results in Table 2 are presented for AMs trained with only the VHF data. Applying speaker role identification for the pre-processed data (66 h) yields 43 h for ATCO and 23 h for Pilot. Similar to Table 1, the results in Table 2 show that using multitask learning instead of training AM by combining all the data provides better ASR performance. Furthermore, the results reveal that due the low amount of data, multitask learning outperforms its single task counterparts.

5.3. Experiments on VHF+other ATC datasets

In this subsection we report results with models trained from both VHF and ATC datasets used in the previous two

Table 3: WER comparison for models trained with all ATCO data from all databases and all pilot data with noise augmented data

Model	WER %	
	ATCO test	Pilot test
ATCO	30.3	43.2
Pilot	32.8	40.3
Combined	31.2	41.3
Multitask	31.9	41.3

experiments. By investigating the ATC databases used in Section 5.1, we discovered that some of the datasets also contain pilot speech. Since no speaker role labels are available for these sets, we applied the proposed method to split the noise augmented speech as ATCO or pilot and combined them with their respective classes of the VHF data. This provided 123h of data for ATCO and 80h for pilot. The results in Table 3 show that training AMs for each task separately performs relatively better by 2.9% for ATCO and 2.4% for pilot than using the Combined system. This suggests that when more data is available, using our grammar-based approach to obtain speaker role information to train separate ATCO and pilot ASR is better than the Combined approach. The Multitask system does not perform better than the Combined; suggesting a negative transfer when considering ATCO and pilot tasks. This is expected as the ATC data dominates in size during training.

6. Conclusions

In this work, we compared different types of training AMs with state-of-the-art LF-MMI framework for ATCO and pilot speech recordings. The developed ASR systems were evaluated separately on ATCO and pilot test sets built from LiveATC. Due to the noisy nature of both ATCO and pilot test sets, AM trained with only noise augmented speech data boosts the ASR performance. We proposed a simple grammar based approach to identify speaker roles automatically and train acoustic models either by speaker role or in a multitask fashion. The results show that multitask training approach outperforms other training methods when limited training data is available. When sufficient data is available, we show that training AMs separately provides better ASR performance for both ATCO and pilot compared to the model trained by combining all data. Relative improvements of 3.2% for the ATCO set and 1.9% for the pilot set were obtained.

7. Acknowledgements

The work was supported by the European Union’s Horizon 2020 project No. 864702 - ATCO2 (Automatic collection and processing of voice data from air-traffic communications), which is a part of Clean Sky Joint Undertaking. The work was also partially supported by SESAR EC project No. 884287 - HAAWAI (Highly automated air-traffic controller workstations with artificial intelligence integration).

8. References

- [1] A. Srinivasamurthy, P. Motlicek, I. Himawan, G. Szaszak, Y. Oualil, and H. Helmke, “Semi-supervised learning with semantic knowledge extraction for improved speech recognition in air traffic control,” in *Proc. of the 18th Annual Conference of the International Speech Communication Association*, 2017.

- [2] M. Kleinert, H. Helmke, G. Siol, H. Ehr, A. Cerna, C. Kern, D. Klakow, P. Motlicek, Y. Oualil, M. Singh *et al.*, "Semi-supervised adaptation of assistant based speech recognition models for different approach areas," in *37th Digital Avionics Systems Conference (DASC)*. IEEE, 2018, pp. 1–10.
- [3] H. Helmke, O. Ohneiser, T. Mühlhausen, and M. Wies, "Reducing controller workload with automatic speech recognition," in *2016 IEEE/AIAA 35th Digital Avionics Systems Conference (DASC)*. IEEE, 2016, pp. 1–10.
- [4] H. Helmke, O. Ohneiser, J. Buxbaum, and C. Kern, "Increasing atm efficiency with assistant based speech recognition," in *Proc. of the 13th USA/Europe Air Traffic Management Research and Development Seminar, Seattle, USA*, 2017.
- [5] J. Zuluaga-Gomez, P. Motlicek, Q. Zhan, K. Veselý, and R. Braun, "Automatic speech recognition benchmark for air-traffic communications," in *Interspeech 2020, 21st Annual Conference of the International Speech Communication Association, Virtual Event, Shanghai, China, 25-29 October 2020*, H. Meng, B. Xu, and T. F. Zheng, Eds. ISCA, 2020, pp. 2297–2301. [Online]. Available: <https://doi.org/10.21437/Interspeech.2020-2173>
- [6] J. Zuluaga-Gomez, K. Veselý, A. Blatt, P. Motlicek, D. Klakow, A. Tart, I. Szöke, A. Prasad, S. Sarfjoo, P. Kolčárek *et al.*, "Automatic call sign detection: Matching air surveillance data with air traffic spoken communications," in *Multidisciplinary Digital Publishing Institute Proceedings*, vol. 59, no. 1, 2020, p. 14.
- [7] S. Ruder, "An overview of multi-task learning in deep neural networks," *arXiv preprint arXiv:1706.05098*, 2017.
- [8] X. Anguera, S. Bozonnet, N. Evans, C. Fredouille, G. Friedland, and O. Vinyals, "Speaker diarization: A review of recent research," *IEEE Transactions on Audio, Speech, and Language Processing*, vol. 20, no. 2, pp. 356–370, 2012.
- [9] T. J. Park, N. Kanda, D. Dimitriadis, K. J. Han, S. Watanabe, and S. Narayanan, "A review of speaker diarization: Recent advances with deep learning," *arXiv preprint arXiv:2101.09624*, 2021.
- [10] ALLCLEAR, "Icao phraseology reference guide," 2020. [Online]. Available: <https://www.skybrary.aero/bookshelf/books/115.pdf>
- [11] S. Madikeri, B. K. Khonglah, S. Tong, P. Motlicek, H. Bourlard, and D. Povey, "Lattice-free maximum mutual information training of multilingual speech recognition systems," in *Proc. of Interspeech*, vol. 2020, 2020.
- [12] L. Burget, P. Schwarz, M. Agarwal, P. Akyazi, K. Feng, A. Ghoshal, O. Glembek, N. Goel, M. Karafiát, D. Povey *et al.*, "Multilingual acoustic modeling for speech recognition based on subspace gaussian mixture models," in *2010 IEEE International Conference on Acoustics, Speech and Signal Processing*. IEEE, 2010, pp. 4334–4337.
- [13] D. Imseng, P. Motlicek, H. Bourlard, and P. N. Garner, "Using out-of-language data to improve an under-resourced speech recognizer," *Speech communication*, vol. 56, pp. 142–151, 2014.
- [14] N. T. Vu, D. Imseng, D. Povey, P. Motlicek, T. Schultz, and H. Bourlard, "Multilingual deep neural network based acoustic modeling for rapid language adaptation," in *2014 IEEE international conference on acoustics, speech and signal processing (ICASSP)*. IEEE, 2014, pp. 7639–7643.
- [15] M. Karafiát, M. K. Baskar, P. Matějka, K. Veselý, F. Grézl, and J. Černocký, "Multilingual blstm and speaker-specific vector adaptation in 2016 but babel system," in *2016 IEEE Spoken Language Technology Workshop (SLT)*. IEEE, 2016, pp. 637–643.
- [16] D. Povey, A. Ghoshal, G. Boulianne, L. Burget, O. Glembek, N. Goel, M. Hannemann, P. Motlicek, Y. Qian, P. Schwarz *et al.*, "The kaldi speech recognition toolkit," in *IEEE workshop on automatic speech recognition and understanding*, no. CONF. IEEE Signal Processing Society, 2011.
- [17] D. Povey, V. Peddinti, D. Galvez, P. Ghahremani, V. Manohar, X. Na, Y. Wang, and S. Khudanpur, "Purely sequence-trained neural networks for asr based on lattice-free mmi," in *Interspeech*, 2016, pp. 2751–2755.
- [18] D. Povey, G. Cheng, Y. Wang, K. Li, H. Xu, M. Yarmohammadi, and S. Khudanpur, "Semi-orthogonal low-rank matrix factorization for deep neural networks," in *Interspeech*, 2018, pp. 3743–3747.
- [19] S. Pigeon, W. Shen, A. Lawson, and D. A. v. Leeuwen, "Design and characterization of the non-native military air traffic communications database (nmmatc)," in *Eighth Annual Conference of the International Speech Communication Association*, 2007.
- [20] J. Segura, T. Ehrette, A. Potamianos, D. Fohr, I. Illina, P. Breton, V. Clot, R. Gemello, M. Matassoni, and P. Maragos, "The hiwire database, a noisy and non-native english speech corpus for cockpit communication," *Online*. <http://www.hiwire.org>, 2007.
- [21] K. Hofbauer, S. Petrik, and H. Hering, "The atcosim corpus of non-prompted clean air traffic control speech," in *LREC*, 2008.
- [22] E. Delpech, M. Laignelet, C. Pimm, C. Raynal, M. Trzos, A. Arnold, and D. Pronto, "A Real-life, French-accented Corpus of Air Traffic Control Communications," in *Proceedings of the Eleventh International Conference on Language Resources and Evaluation (LREC 2018)*, 2018.
- [23] J. Godfrey, "The Air Traffic Control Corpus (ATCO) - LDC94S14A," 1994. [Online]. Available: <https://catalog.ldc.upenn.edu/LDC94S14A>
- [24] H. Helmke, M. Slotty, M. Poiger, D. F. Herrler, O. Ohneiser, N. Vink, A. Cerna, P. Hartikainen, B. Josefsson, D. Langr *et al.*, "Ontology for transcription of atc speech commands of sesar 2020 solution pj. 16-04," in *IEEE/AIAA 37th Digital Avionics Systems Conference (DASC)*. IEEE, 2018, pp. 1–10.
- [25] Y. LeCun, Y. Bengio *et al.*, "Convolutional networks for images, speech, and time series," *The handbook of brain theory and neural networks*, vol. 3361, no. 10, p. 1995, 1995.
- [26] T. Ko, V. Peddinti, D. Povey, and S. Khudanpur, "Audio augmentation for speech recognition," in *Sixteenth Annual Conference of the International Speech Communication Association*, 2015.