

Virtual Machines and Containers as a Platform for Experimentation

Florian Metze¹, Eric Riebling¹, Anne S. Warlaumont², and Erika Bergelson³

¹Language Technologies Institute, Carnegie Mellon University; U.S.A.

²Cognitive and Information Sciences, University of California, Merced; U.S.A.

³Center for Language Sciences, University of Rochester; U.S.A., and
Psychology & Neuroscience Department, Duke University; U.S.A.

fmetze@cs.cmu.edu

Abstract

Research on computational speech processing has traditionally relied on the availability of a relatively large and complex infrastructure, which encompasses data (text and audio), tools (feature extraction, model training, scoring, possibly on-line and off-line, etc.), glue code, and computing. Traditionally, it has been very hard to move experiments from one site to another, and to replicate experiments. With the increasing availability of shared platforms such as commercial cloud computing platforms or publicly funded super-computing centers, there is a need and an opportunity to abstract the experimental environment from the hardware, and distribute complete setups as a virtual container or some other shareable resource, that can be deployed and worked with anywhere.

In this paper, we discuss our experience with this concept and present some tools that the community might find useful. We outline, as a case study, how such tools can be applied to a naturalistic language acquisition audio corpus.

Index Terms: speech processing, reproducible research, citizen science, shared platforms, cloud computing

1. Introduction

Building and maintaining a state-of-the-art Automatic Speech Recognition (ASR) or Natural Language Processing (NLP) system has become a very complex task: it requires installing large amounts of text and audio data, normalizing them, training acoustic and language models with different tools, combining them to decode test data, perform error analysis, etc. It requires more than one expert to maintain an end-to-end system, and adapt it to new languages, tasks, or conditions, and even rebuilding an existing baseline and adapting it to new data is a daunting task, prone to errors.

This effort attempts to extend the model of lab-internal knowledge transfer to a community-wide effort through the use of Virtual Machines and/ or Containers. The Speech Recognition Virtual Kitchen (SRVK) [1]¹ presents a way to share not just “recipes”, but ready-to run experimental environments – hence the analogy to a kitchen: our setup contains tools, data, log-files and results from baseline runs, so that classes, evaluations, workshops, or collaborating researchers can easily share setups and compare results meaningfully.

In this paper, we present our recently released “Eesen Transcriber” Virtual Machine (VM), and discuss its applicability to the HomeBank repository. At the special session, we will invite researchers to “taste” the latest version of the transcriber

¹Visit us at <https://github.com/srvk> or <http://speechkitchen.org/>

VM, and give visitors a guided tour of its possibilities using the HomeBank database as an example application.

2. Related Work

The need for shared, reproducible tools in research has been recognized by several other groups and communities, and has already led to many tools being made available on public, shared repositories such as GitHub [2], and under an Open Source license (e.g. Apache 2.0 [3]). In the speech community, Kaldi [4] is certainly the most successful such project, given the large number of “recipes” it contains, some of which rely on Open Source data, e.g. the LIUM corpus [5].

Box [6] is a similar project from the NLP community, which attempts to provide a disk image with several tools already pre-installed. Code and data repositories such as CodaLab [7] provide an environment for reproducible research, while repositories such as Covarep [8] strive to make sharing of implementations easier. In the speech community, WebASR [9] and CloudCAST [10] aim to provide a shared infrastructure, and some access to components for researchers, rather than treating ASR as a black box.

The SRVK combines aspects from all of these, in that we aim to provide a ready-to-use tool, that can be run in the cloud or at the client’s site. Distributing it as a VM provides complete control over all aspects of processing to the user, allowing for customizations easily, and access to source code and data. As such, the concept is also suitable for education; this concept is discussed in a companion paper [11].

Previously, an SRVK Virtual Machine has been used to distribute a baseline to MediaEval’s “QUESST” task [12], which participants found very helpful. An earlier, related effort created tools to facilitate the building of speech recognizers by non-experts [13, 14, 15]. Many of the tools and techniques developed within this thesis [16] have been integrated into the “Eesen Transcriber” VM described next.

3. The “Kitchen” Concept

The main idea behind the kitchen is to distribute Virtual Machines (VMs) or containers² to the community, which contain Open Source software, so that users can freely download, copy, and/ or modify them. The kitchen’s resources are typically

²Going forward, we will simply refer to “VMs”, although the concepts apply equally to Virtual Machines and (Docker) containers. Both of these are supported by the kitchen, and community feedback and contributions are encouraged to provide the best mix of technologies. To those readers unfamiliar with these concepts of “virtualization”, we recommend they read e.g. [17].

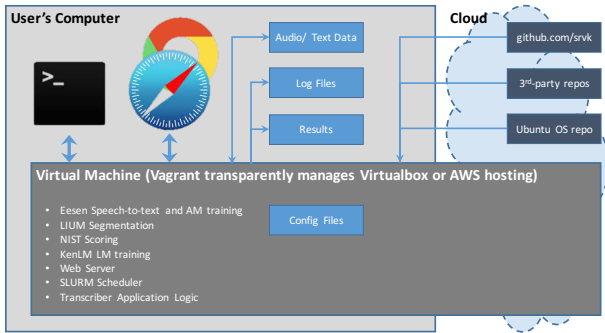


Figure 1: “Eesen Transcriber” architecture: the user controls the VM with a terminal window and/ or a browser on his or her own computer; it is transparent if the VM has been instantiated on the user’s computer, or on the cloud. Data, log files, and results reside on the user’s computer, while configuration files mostly reside inside the VM, to make it portable. During instantiation, the VM pulls resources from online repositories as required.

made available as a git repository containing mostly scripts, in particular `Dockerfiles` [18] or `Vagrantfiles` [19]. The user can install these either on his or her own computer (using e.g. Virtualbox [20] or cluster (irrespective of operating system) or on the cloud (currently, we support Amazon Web Services (AWS) [21]). With this mechanism, it is not necessary to distribute pre-compiled, monolithic VMs, which can easily reach Gigabytes in size, but rather small text files, that describe how a VM can be built from scratch, usually in a matter of minutes.

As of today, the kitchen provides a number of “dishes”, yet we **encourage and require community feedback and contributions** in order to further develop all the ideas presented here.

A typical “work-flow” follows, also compare with the architecture shown in Figure 1:

Check out a VM using e.g. `git clone https://github.com/srvk/eesen-transcriber.git`, and `cd` into the newly created directory structure (unfortunately, some use of the command line is required), the so-called “shared” folder

Configure the VM to better match the amount of memory or number of processors used by the VM to the resources available on the host (often, this step may not be needed)

Bring up the VM using `vagrant up` or `sudo docker build`. Often, a `build.sh` script is provided to facilitate the download of additional data; at the end of this “instantiation” step, the VM will print further instructions on how to continue

Work with the VM: This can involve copying data into the shared folder, or accessing a web browser, executing commands in the VM using `vagrant exec`, or logging into the VM to execute commands, debug the application, or further develop the VM

Shut down the VM, leaving it available for re-starting, continuing work from where one left off, or wiping it from the computer entirely

In all cases, the repository contains documentation and pointers to shared documentation and community resources that are available online.

4. The “Eesen Transcriber” VM

After discussion with several potential users, easy-to-use speech-to-text that can be treated as a black box, but does permit access to the internals if required, emerged as a central need of the community.

We thus combined our recently released “Eesen” ASR toolkit [22] with an Open Source “transcriber” [23] and segmentation [24] framework, and added several additional functionality, releasing the system with English acoustic models trained on TEDLIUM data [25], and a large general purpose language model [26].

Once the system has been started as described in the previous section, the user is given a URL, which points to a control and help page for the current VM. This page is being hosted on the VM, either locally or in the cloud, and allows the user to check the status of the VM, and perform certain tasks without having to use the command line interface (which gives complete control over the VM, including `sudo rm -rf /`).

4.1. Acoustic Model

The acoustic model is based on a bi-directional Long Short-Term Memory (LSTM) network, which has been trained using the Connectionist Temporal Classification (CTC) criterion [27], which directly optimizes the sequence of output symbols, rather than the frame likelihood. The VM comes with a pre-trained Acoustic Model, so that the user is able to immediately start decoding, and does not need to provide a GPU-based instance in order to train it. All the scripts are provided however, so that a user can start the VM on a GPU-enabled machine (for example an AWS `g2.8xlarge` instance), in order to re-train the model, adapt it, or use the provided recipe as a blueprint to train a model for a different language or task.

CTC is a good choice as an optimization criterion, because the model does not depend on the initialization (the alignment is being marginalized out during the optimization of the network), and there is only a single, context independent (CI) model that has to be built. As an example, the “Eesen” training recipe to build a Wall Street Journal (WSJ) system has less than 100 lines of code including data preparation and testing, while the corresponding “Kaldi” recipe (from which the “Eesen” one is derived) has more than 450 lines of code, and both reach around 7% Word Error Rate (WER).

CI acoustic models also result in smaller Weighted Finite State Transducer decoding graphs, and decode faster.

4.2. Language Model

Similarly, a pre-compiled n-gram Language Model is provided, but the VM comes with instructions on how to rebuild the language model, or how to adapt it to different domains. As we expect that more users will want to adapt the language model (and training does not require a GPU node), these instructions are quite detailed, and all required software (currently, we use KenLM [28]) is already installed and set up.

Future developments will also allow the integration of neural network language models, resulting in an all-neural, “end-to-end” experimental work-bench in a VM.

4.3. Running on Own Data

“Eesen transcriber” implements a complete end-to-end pipeline for processing and indexing audio. Processing can be triggered by either “dropping” an audio file in a “watched folder” on the user’s host, in which case a daemon on the VM will pick up

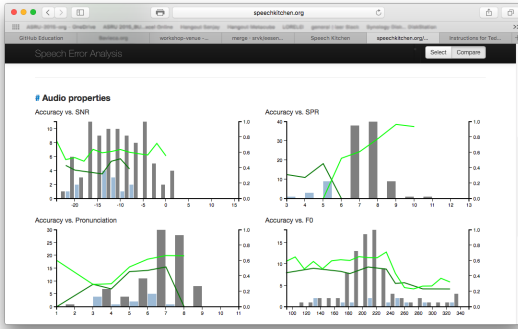


Figure 2: The data analysis tool originally developed in [16]. In this example, the user selected two parts of the data: young children (blue histogram bars), and older children (grey bars). The accuracy for young children (dark green line) is lower than for older children (light green line). The main difference (see the “Accuracy vs. SPR” plot in the top right) is that young children speak slower (lower “speaking rate” value).

the file and trigger processing, or by manually triggering processing from the host using a command like `vagrant ssh -c "vids2web.sh /videos/demo.mp4"` which will directly execute a script on the VM with a file that can be found on the host. The VM uses shared folders or `sshfs`s to make file systems accessible in a transparent way.

The following steps are currently supported out of the box:

Normalization of arbitrary audio files to 16kHz sampling frequency and PCM file format

Diarization of the audio file: mono and stereo files can be segmented automatically into speakers, or a provided STM file can be read in order to use a given segmentation. The VM contains documentation on how to alter some of the parameters inherent to the segmentation, if the segmentation is found to be performing badly.

Feature extraction of IMEL and pitch features (by default); other approaches can be implemented easily if required

Speech-to-text to produce lattices, word confidences, and n-est lists with time alignments; the pipeline can be configured to produce a number of output file formats, e.g. CTM, VTT, SBV, etc.

Indexing of the ASR output in a simple database, implemented within a standard web server

Browsing and Searching in a web browser on the user’s computer, permitting play-back of audio and video with subtitles, similar to a video library [29] or a lecture browser [30]

Processing is transparently controlled by the SLURM workload manager, so that any amount of data can be “dumped” on the VM with a single action. Command shortcuts are being provided for the most important tasks such as rebuilding the index, cleaning cached or intermediate files, processing files, etc. While most of the configuration files are kept within the VM, most of the data files and indices are kept outside of the VM, so that the user has easy access to all of them, and can process them further or store them as required. Figure 1 shows the most important components.

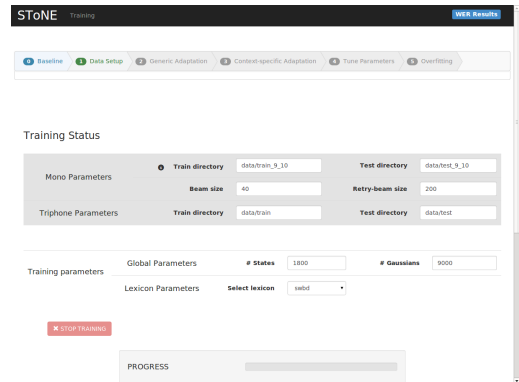


Figure 3: The “SToNE” training module (from [16]), which is being adapted to the “Eesen” framework, and its unique set of parameters. Such a framework allows a user to train a system by selecting fundamental parameters in a GUI, and follow along a typical development process.

4.4. Analyzing Results

Error analysis is an important part of system development, which is often tedious for experts and sometimes impossible for non-experts, because they do not know what to look for. Figure 2 (generated on <http://speechkitchen.org/home/error-analysis-tools/>) shows a selection of online tools, which we are currently porting into the VM. It will compute a number of features such as signal-to-noise ratio, speaking rate, pronunciation score, fundamental frequency, and power for training and test data, and allow the user to correlate these with known demographics of the data (age, gender, dialect, etc.). The VM further integrates the NIST sclite scoring tools [31], and keeps track of out-of-vocabulary words.

The VM thus provides tools to identify mis-matched data and the reasons behind poor performance, and enables the user to further develop the system accordingly.

4.5. Re-training the Models

Leveraging earlier work [16], we are currently porting the Kaldi-based “SToNE” (Speech Toolkit for Non-Experts) “guided” training module to Eesen and the Transcriber VM, which will allow a non-expert user to develop a usable speech recognition system. It uses an iterative process, in which a rule-based expert system analyzes the system’s errors, and gives recommendations as to which parts of the system should be modified, in order to improve performance.

The focus is less on the choice of different Deep Neural Network (DNN) architectures, but on getting the fundamentals in place, e.g. what is the out-of-vocabulary rate, is the test data too conversational, are errors correlated with noise, etc.

4.6. Performance and Characteristics

The “Eesen Transcriber” is currently distributed with a fully open TEDLIUM version 2 [25] system, which delivers a WER of about 15% on the test set, using the pruned Cantab LM [26]. It decodes in less than real-time on a typical single-processor VM, and can be retrained from scratch within 6 days on a single AWS GPU (we are currently working to parallelize the training to multi-GPU cloud architectures).

Any “Eesen” or Kaldi recipe can be integrated into the “Transcriber” VM, if the user has access to the required data.

4.7. Prospects

The “Eesen Transcriber” VM thus presents a complete speech processing toolbox which can be used to process and analyze arbitrary audio. It will ingest any audio or video file in a format known to SoX and/ or FFmpeg, and allows the user to index, browse, and search his collection. If a reference is provided, it will compute various statistics and visualizations that can be helpful in tuning or developing a system.

The “Transcriber” VM separates data (outside of the VM) from tools (inside of the VM), and keeps the “non-user-serviceable” parts hidden from the novice user. It gives a semi-expert easy access to all the components of the system and allows him to make modifications or improvements as required. Because all the code is assembled using a provisioning solution such as Vagrant or Docker, a system can always be re-created from scratch, and variants can be created and shared easily using a version control system (git in our case).

5. Case Study: the HomeBank Repository

HomeBank [32] is a repository for long-form child audio recordings, their meta-data, and code used to process and analyze this data (see [33] for details). HomeBank is part of the Talkbank/ CHILDES family of child language corpora [34]. Presently, the modal use of HomeBank is as a repository for data generated by LENA recorders, which store up to 16h of continuous audio on small recording devices worn by children [35]. The audio files can then be processed using LENA’s segmentation and diarization, or be fed further into other analysis pipelines, like those implemented by the Speech Kitchen. HomeBank is also part of DARCLE (Daylong Audio Recordings of Childrens Language Environments, [36]), an international group focused on connecting child language researchers with speech technology experts and computational modelers.

One of the largest challenges for HomeBank has been figuring out the ethics guidelines for sharing daylong audio recordings with other researchers. Such naturalistic recordings often have personal information in them (e.g. names, addresses), or take place in public or semi-public settings where some of the individuals recorded may not have given consent for their voices to be recorded (e.g. at pre-schools). Moreover, families often permit sharing in certain research settings without realizing what content may be on the daylong recordings. Naturally, cleaning such long-form recordings for personal content is an extremely laborious process. HomeBank has developed an approved membership protocol, where recordings that are in the “password” section of the repository are only available to vetted members of the research community, while “public” recordings are available more broadly.

From a speech technology standpoint, one challenge has been in developing automatically derived measures that are linguistically meaningful from daylong audio files that contain noisy recordings from a single microphone source. This makes speaker diarization and segmentation even more challenging for even state-of-the-art algorithms. Nevertheless, recordings of this kind have been successfully used in semi- or fully-automated applications to, e.g. help diagnose children with language delays and deficits in a reliable way [37], to help understand the role of parental responses in speech development [38], and to help understand the differences in language input to children in low-income households [39]. Such recordings have also been used to examine how the acoustic properties of real-world mother-child exchanges evolve over time [40].

While the LENA system’s segmentation and diarization methods have shown reasonable reliability compared to human listeners’ labelings, and have been demonstrated to be sufficient for various scientific and applied purposes, the algorithm has not been substantially updated since the software was made publicly available about a decade ago. Users of the system could benefit from application of alternative segmentation and diarization methods, which have the potential to be more accurate as well as to provide more detailed information (for example, one limitation of the LENA software is that segments labeled as containing overlapping speech do not indicate which speaker types are the ones that overlap).

Speech Kitchen VMs enable users with limited technical skills to test and adapt alternative segmentation and diarization techniques from state-of-the-art providers. The challenges mentioned earlier in the paper, of configuring one’s system (so that all the necessary packages are installed, file system matches what’s assumed by the code, and so on), inhibit sharing of tools even among engineering labs; this problem is expected to be that much greater when the users are behavioral scientists with little or no computer science training.

Additionally, the LENA system does not attempt any transcription of the speech. It seems worthwhile to apply the latest ASR methods to the HomeBank dataset to see if it is possible to train/ adapt them. This would be a hugely beneficial addition to the automatic segmentation and diarization. For example, an ability to detect a specified set of keywords could be very useful for scientists interested in how the natural audio environment supports word learning. Prior work suggests this is a reasonable goal [16, 41].

6. Outlook and Future Work

SRVK is currently preparing several dishes, to be served over the summer. For example, a Kaldi Tutorial at SLTU 2016 [42] will be using VMs created with techniques developed in the SRVK project. In another effort, the “Eesen Transcriber” VM will be expanded to include not just speech retrieval, but a complete multi-media indexing and retrieval system, which can also query visual or paralinguistic information about input material.

All components of the “Eesen Transcriber” VM are currently available to users of the kitchen, and are currently being integrated into a simple and re-usable `Vagrantfile`. A demo of this VM will be shown at INTERSPEECH 2016, and copies will be made available to attendants.

HomeBank and DARCLE are especially potent diners for the dishes served by the Speech Kitchen because one main goal of the group is to develop a fully open-source pipeline for data analysis. The “Eesen Transcriber” VM is perfectly suited for this task. Relatedly, a subset of HomeBank data is currently being prepared for a Computational Paralinguistics Challenge.

7. Acknowledgments

The authors would like to thank Anuj Kumar for his work on the error analysis and system building web application. We would also like to thank Brian MacWhinney and Melanie Soderstrom.

This work used the Extreme Science and Engineering Discovery Environment (XSEDE), which is supported by U.S. National Science Foundation grant number OCI1053575. The Speech Recognition Virtual Kitchen is supported by the NSF under grant numbers IIS-1247368, CNS-1305365, 1305319, and 1305215. HomeBank is supported by NSF SBE-1539129, 1539133, and 1539010.

8. References

- [1] F. Metze, E. Fosler-Lussier, and R. Bates, "The speech recognition virtual kitchen," in *Proc. INTERSPEECH*. Lyon; France: ISCA, Aug. 2013, <https://github.org/srvk>.
- [2] <https://github.com/>.
- [3] <http://www.apache.org/licenses/>.
- [4] D. Povey, A. Ghoshal, G. Boulianne, L. Burget, O. Glembek, N. Goel, M. Hannemann, P. Motlíček, Y. Qian, P. Schwarz, J. Silovsky, G. Stemmer, and K. Veselý, "The Kaldi speech recognition toolkit," in *Proc. Automatic Speech Recognition and Understanding (ASRU), 2011 IEEE Workshop on*. Waikoloa, HI; U.S.A.: IEEE, Dec. 2011.
- [5] A. Rousseau, P. Deléglise, and Y. Estève, "TED-LIUM: an automatic speech recognition dedicated corpus," in *Proc. LREC*. Istanbul; Turkey: ELRA, May 2012.
- [6] A. Axelrod, "Box: Natural language processing research using amazon web services," *The Prague Bulletin of Mathematical Linguistics*, vol. 104, no. 1, pp. 27–38, 2015.
- [7] <http://codalab.org>.
- [8] G. Degottex, J. Kane, T. Drugman, T. Raitio, and S. Scherer, "Covarepa collaborative voice analysis repository for speech technologies," in *Acoustics, Speech and Signal Processing (ICASSP), 2014 IEEE International Conference on*. Florence; Italy: IEEE, May 2014, pp. 960–964.
- [9] T. Hain, A. El Hannani, S. N. Wrigley, and V. Wan, "Automatic speech recognition for scientific purposes-webasr," in *Proc. INTERSPEECH*, Brisbane; Australia, Sep. 2008.
- [10] P. Green, R. Marxer, S. Cunningham, H. Christensen, F. Rudzicz, M. Yancheva, A. Coy, M. Malavasi, and L. Desideri, "Remote speech technology for speech professionals - the CloudCAST initiative," in *6th Workshop on Speech and Language Processing for Assistive Technologies (SLPAT)*, 2015.
- [11] R. Bates, E. Fosler-Lussier, F. Metze, M. Larson, G.-A. Levow, and E. M. Provost, "Experiences with shared resources for research and education in speech and language processing," in *Proc. INTERSPEECH*. San Francisco, CA; U.S.A.: ISCA, Sep. 2016, submitted.
- [12] X. Anguera, L. J. Rodríguez-Fuentes, I. Szöke, A. Buzo, and F. Metze, "Query by example search on speech at Mediaeval 2014," in *Proc. MediaEval Workshop*, Barcelona; Spain, Oct. 2014, <http://ceur-ws.org/Vol-1263>.
- [13] A. Kumar, F. Metze, W. Wang, and M. Kam, "Formalizing expert knowledge for developing accurate speech recognizers," in *Proc. INTERSPEECH*. Lyon; France: ISCA, Aug. 2013.
- [14] A. Kumar, F. Metze, and M. Kam, "Enabling the rapid development and adoption of speech-user interfaces," *IEEE Computer Magazine*, vol. 46, no. 1, Jan. 2014, <http://dx.doi.org/10.1109/mc.2014.11>.
- [15] A. Kumar, F. Metze, E. Riebling, and M. Kam, "Demystifying development of speech recognizers for novices," in *Proc. Designing Speech and Language Interactions Workshop at CHI*. Toronto; Canada: ACM, Apr. 2014.
- [16] A. Kumar, "Enabling non-speech experts to develop usable speech-user interfaces," Ph.D. dissertation, School of Computer Science, Carnegie Mellon University, Aug. 2014, cMU-HCII-14-105.
- [17] VirtualBox, "Virtual machines," <https://www.virtualbox.org/wiki/Virtualization>, Mar. 2016.
- [18] <https://docker.com/>.
- [19] <https://vagrantup.com/>.
- [20] <https://virtualbox.org/>.
- [21] <https://aws.amazon.com/>.
- [22] Y. Miao, M. Gowayed, and F. Metze, "EESSEN: End-to-End Speech Recognition using Deep RNN Models and WFST-based Decoding," in *Proc. Automatic Speech Recognition and Understanding Workshop (ASRU)*. Scottsdale, AZ; U.S.A.: IEEE, Dec. 2015, <https://github.com/srvk/eessen>.
- [23] <https://github.com/alumae/kaldi-offline-transcriber>.
- [24] M. Rouvier, G. Dupuy, P. Gay, E. Khoury, T. Merlin, and S. Meignier, "An open-source state-of-the-art toolbox for broadcast news diarization," in *Proc. INTERSPEECH*. Lyon; France: ISCA, Sep. 2013.
- [25] A. Rousseau, P. Deléglise, and Y. Estève, "Enhancing the ted-lium corpus with selected data for language modeling and more ted talks," in *Proc. LREC*. Reykjavik; Iceland: ELRA, May 2014.
- [26] W. Williams, N. Prasad, D. Mrva, T. Ash, and T. Robinson, "Scaling recurrent neural network language models," in *Acoustics, Speech and Signal Processing (ICASSP), 2015 IEEE International Conference on*. Brisbane; Australia: IEEE, May 2015.
- [27] A. Graves, S. Fernández, F. Gomez, and J. Schmidhuber, "Connectionist temporal classification: labelling unsegmented sequence data with recurrent neural networks," in *Proceedings of the 23rd international conference on Machine Learning*. ACM, 2006, pp. 369–376.
- [28] K. Heafield, "KenLM: faster and smaller language model queries," in *Proceedings of the EMNLP 2011 Sixth Workshop on Statistical Machine Translation*, Edinburgh, Scotland, United Kingdom, July 2011, pp. 187–197. [Online]. Available: <http://kheafield.com/professional/avenue/kenlm.pdf>
- [29] A. G. Hauptmann and M. J. Witbrock, "Informedia: News-on-demand multimedia information acquisition and retrieval," in *Intelligent Multimedia Information Retrieval*, M. T. Maybury, Ed. Cambridge, MA, USA: MIT Press, 1997, pp. 215–239. [Online]. Available: <http://dl.acm.org/citation.cfm?id=266260.266279>
- [30] J. R. Glass, T. J. Hazen, D. S. Cyphers, K. Schutte, and A. Park, "The mit spoken lecture processing project," in *Proc. NAACL*, 2005.
- [31] <http://www.nist.gov/itl/iad/mig/tools.cfm>.
- [32] <http://homebank.talkbank.org/>.
- [33] M. VanDam, A. S. Warlaumont, E. Bergelson, A. Cristia, M. Soderstrom, P. DePalma, and B. MacWhinney, "Homebank, an online repository of daylong child-centered audio recordings," *Seminars in Speech and Language*, 2016, in press.
- [34] B. MacWhinney, *The CHILDES project: The database*, 2nd ed. Psychology Press, 2000.
- [35] <https://www.lenafoundation.org>.
- [36] <http://www.darcl.org/>.
- [37] D. Oller, P. Niyogi, S. Gray, J. Richards, J. Gilkerson, D. Xu, U. Yapanel, and S. Warren, "Automated vocal analysis of naturalistic recordings from children with autism, language delay, and typical development," *Proceedings of the National Academy of Sciences*, vol. 107, no. 30, pp. 13 354–13 359, 2010.
- [38] A. S. Warlaumont, J. A. Richards, J. Gilkerson, and D. K. Oller, "A social feedback loop for speech development and its reduction in autism," *Psychological science*, p. 0956797614531023, 2014.
- [39] A. Weisleder and A. Fernald, "Talking to children matters early language experience strengthens processing and builds vocabulary," *Psychological science*, vol. 24, no. 11, pp. 2143–2152, 2013.
- [40] E.-S. Ko, A. Seidl, A. Cristia, M. Reimchen, and M. Soderstrom, "Entrainment of prosody in the interaction of mothers with their young children," *Journal of child language*, pp. 1–26, 2015.
- [41] J. Trmal, G. Chen, D. Povey, S. Khudanpur, P. Ghahremani, X. Zhang, V. Manohar, C. Liu, A. Jansen, D. Klakow, D. Yarowsky, and F. Metze, "A keyword search system using open source software," in *Proc. IEEE Workshop on Spoken Language Technology*. South Lake Tahoe, NV; USA: IEEE, Dec. 2014.
- [42] *5th International Workshop on Spoken Language Technologies for Under-resourced Languages*. Yogyakarta; Indonesia: ISCA, May 2016, <http://mica.edu.vn/sltu2016/>.