

Automatic Meeting Transcription System for the Japanese Parliament (Diet)

Tatsuya Kawahara

Kyoto University, School of Informatics, Sakyo-ku, Kyoto 606-8501, Japan

Abstract—Applications of automatic speech recognition (ASR) have been extended to a variety of tasks and domains, including spontaneous human-human speech. We have developed an ASR system for the Japanese Parliament (Diet), which has been operated every day since 2011. By exploiting official records made by human stenographers, we have realized an efficient training scheme of acoustic and language models, which does not require faithful transcripts and thus is scalable to enormous data. This article gives an overview of the system performance and how it changed the meeting record generation process in the Parliament.

I. INTRODUCTION

Since the Japanese Parliament (Diet) was founded in 1890, verbatim records had been made by manual shorthand over a hundred years. However, early in this century, the government terminated recruiting stenographers, and investigated alternative methods¹. The House of Representatives has chosen our ASR technology for the new system. The system was deployed and tested in 2010, and it has been in official operation from April 2011[1]. This is the first automatic transcription system deployed in national Parliaments, and it is still only one running system in the world at this moment.

Requirements for the ASR system are as follows. The first is high accuracy; over 90% is preferred. This can be easily achieved for plenary sessions, but is difficult for committee meetings, which are interactive, spontaneous, and often heated. The second requirement is fast turn-around. In the House, reporters are assigned speech for transcription by 5-minute segments. ASR should be performed almost in real-time, so Parliamentary reporters can start working promptly even during the meeting. The third issue is compliance to the standard transcript guidelines of the House. This can be guaranteed by using only the past Parliamentary meeting records for building the lexicon and language model.

In order to achieve high recognition performance, acoustic and language models must be customized to Parliamentary speech; that is, they need to be trained with a large amount of the matched data. Fortunately, there is a large amount of data of Parliamentary meetings. There is a huge archive of official meeting records in text, amounting to 15M words per year, which is comparable to newspapers. There is also a huge archive of meeting speech, which amounts to 1200 hours per year. However, official meeting records are different from actual utterances due to the editing process by Parliamentary reporters. There are several reasons for this:

¹Similar changes have taken place in many countries over last decades.

differences between spoken-style and written-style, disfluency phenomena such as fillers and repairs, redundancy such as discourse markers, and grammatical corrections.

From these reasons, we need to build a corpus of Parliamentary meetings, which consists of faithful transcripts of utterances including fillers. We prepared this kind of corpus in the size of 200 hours in speech or 2.4M words in text. The corpus is vital for satisfactory performance, but very costly. The methodology has been adopted in the conventional ASR research and development [2],[3],[4],[5],[6],[7]. In fact, the amount of our corpus is large, even compared with corpora prepared in these projects, however, uses only a fraction of the available data as mentioned above. Moreover, the corpus needs to be updated; otherwise, the performance would degrade in time.

In order to exploit the huge archives of Parliamentary meetings in a more efficient manner, we investigated a novel training scheme by focusing on the differences between the official meeting record and the faithful transcript. We found that the majority of the differences can be computationally modeled by a scheme of simple statistical machine translation (SMT). With the statistical model of the difference, we can predict what is uttered from the official records. By applying the SMT model to a huge scale of the past Parliamentary meeting records, a precise language model is generated. Moreover, by matching the audio data with the model predicted for every speaker turn, we can reconstruct what was actually uttered. This results in an effective lightly-supervised training of the acoustic model, by exploiting a huge speech archive that is not faithfully transcribed. As a result, we could build precise models of spontaneous speech in Parliament, and these models will evolve in time, reflecting the change of Members of Parliament (MPs) and topics discussed.

In this article, following the analysis of the differences between faithful transcripts and official meeting records, the scheme of language model transformation is reviewed. Then, the specification of the transcription system deployed in the Japanese Parliament is described. Finally, its evaluations through long-term operation in the real setting and its impact on the Parliamentary reporting are reported.

II. DIFFERENCES BETWEEN “ACTUALLY” SPOKEN LANGUAGE AND “NORMALLY” TRANSCRIBED TEXT

In any languages universally, there is a difference between spoken style and written style. Moreover, spontaneous speech essentially includes disfluency phenomena such as fillers and

TABLE I
ANALYSIS OF EDITS MADE FOR OFFICIAL RECORDS

edit type	category	ratio
Deletion	Fillers	46.7%
	*Repeats/Repairs	9.4%
	Discourse markers	18.4%
	Extraneous expressions	3.0%
Insertion	Function words	7.5%
Substitution	Colloquial expressions	8.4%
*Reordering		5.9%

* means not-simple edits, which are not dealt in the proposed scheme.

repairs as well as redundancy such as repeats and redundant words. Therefore, faithful transcripts of spontaneous speech are not good in terms of readability and documentation. As a result, there is a significant mis-match between faithful transcripts and official meeting records of Parliament because of the editing process by Parliamentary reporters. Similar phenomena are observed in TV programs (between spoken utterances and their captions) and public speaking (between actual speeches and scripts).

Unlike European Parliament Plenary Sessions (EPPS), which were targeted by the TC-Star project [7][8], the great majority of meetings in the Japanese Parliament are in committee meetings. They are more interactive and thus spontaneous, compared to plenary sessions. The characteristics pose a big challenge to ASR; no existing systems could meet the requirements which were demanded by Parliament and stated in Introduction. The characteristics of the meetings also result in many differences between the actual utterances and the official meeting records. It was observed that the difference in the words (edit distance or word error rate) between the two transcripts ranges 10-20% (16% on average). Table I lists the breakdown of the edits made by Parliamentary reporters. The analysis was made for transcripts of 45 meetings held in 2005-2007, which consist of 380K words in total. Among them, approximately a half were by deletion of fillers, and rests were not so trivial. However, most of the edits (around 93%) can be modeled by simple substitution/deletion/insertion of a word (or two words).

When we compare the analysis with the European Parliament Plenary Sessions, we observe Japanese has more disfluency and redundancy, but less grammatical corrections, because the Japanese language has a relatively free grammatical structure. Still, the above conclusion generally holds.

III. SCHEME OF LANGUAGE MODEL TRANSFORMATION

We proposed a statistical scheme to cope with the differences between spontaneous utterances (verbatim text: V) and human-made transcripts (written-style text: W) [9][10]. In this scheme, the two are regarded as different languages and statistical machine translation (SMT) is applied (Figure 1). It can be applied in both directions: to clean a faithful transcript of the spoken utterances to a document-style text, and to recover the faithful transcript from a human-made text.

The decoding process is formulated in the same manner as SMT, which is based on the following Bayes' rule.

$$p(W|V) = \frac{p(W) \cdot p(V|W)}{p(V)} \quad (1)$$

$$p(V|W) = \frac{p(V) \cdot p(W|V)}{p(W)} \quad (2)$$

Here the denominator is usually ignored in the decoding.

We have extended the simple noisy channel model to a log-linear model which can incorporate joint probability $p(W, V)$ and contextual information [11], for the task of cleaning transcripts (eq. 1).

A. Estimation of Verbatim-style Language Model

On the other hand, the process to uniquely determine V (eq. 2) is much more difficult than the cleaning process (eq. 1) because there are more arbitrary choices in this direction; for example, fillers can be randomly inserted in eq. 2 while all fillers are removed in eq. 1. Therefore, we are more interested in estimating the statistical language model of V , rather than recovering the text of V . Thus, we derive the following estimation formula.

$$p(V) = p(W) \cdot \frac{p(V|W)}{p(W|V)} \quad (3)$$

The key point of this scheme is that the available text size of the document-style texts W is much larger than that of the verbatim texts V needed for training ASR systems. For the Parliamentary meetings, we have a huge archive of official meeting records. Therefore, we fully exploit their statistics $p(W)$ to estimate the verbatim-style language model $p(V)$ for ASR (Figure 1 right-hand side).

The transformation is actually performed on occurrence counts of N-grams as below. Here we count up to trigrams.

$$N_{gram}(v_1^n) = N_{gram}(w_1^n) \cdot \frac{p(v|w)}{p(w|v)} \quad (4)$$

Here v and w are individual transformation patterns. We model substitution $w \rightarrow v$, deletion of w , and insertion of v , by considering their contextual words.² $N_{gram}(w_1^n)$ is an N-gram entry (count) including them, to be revised to $N_{gram}(v_1^n)$. Estimation of the conditional probabilities $p(v|w)$ and $p(w|v)$ requires an aligned corpus of verbatim transcripts and their corresponding document-style texts. We constructed the "parallel" corpus by using a part of the official records of the Parliamentary meetings, which was mentioned in Introduction. The conditional probabilities are estimated by counting the corresponding patterns observed in the corpus. Their neighboring words are taken into account in defining the transformation patterns for precise modeling. For example, an insertion of a filler "ah" is modeled by $\{w = (w_{-1}, w_{+1}) \rightarrow v = (w_{-1}, ah, w_{+1})\}$, and the N-gram entries affected by this insertion are revised. A smoothing technique based on POS

²Unlike ordinary SMT, permutation of words is not considered in this transformation.

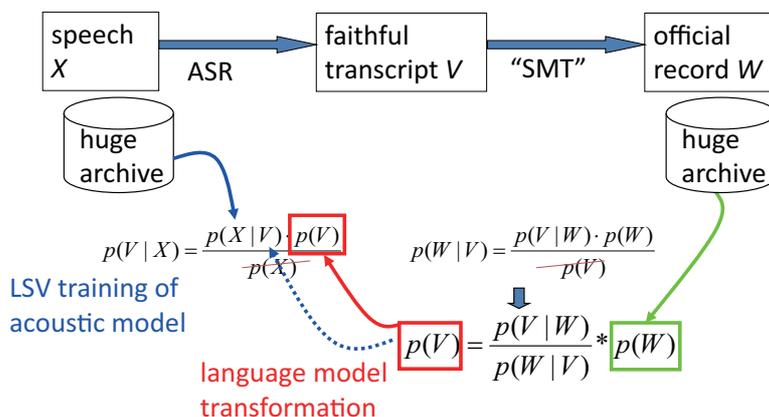


Fig. 1. Overview of the proposed scheme of ASR model training

(Part-Of-Speech) information is introduced to mitigate the data sparseness problem. Please refer to [10] for implementation details and evaluation.

B. Lightly Supervised Training of Acoustic Model

The language model transformation scheme is also applied to lightly supervised training (LSV) of an acoustic model [12]. For the Parliamentary meetings, we have a huge archive of speech which are not faithfully transcribed but have edited texts in the official records. Since it is not possible to uniquely recover the faithful verbatim transcripts from texts of the official records, as mentioned in the previous sub-section, we generate a dedicated language model for decoding the speech using the corresponding text segment of the official record. As a result of ASR, we expect to obtain a faithful transcript with high accuracy. (Figure 1 left-hand side).

For each turn (typically ten seconds to three minutes, and on the average one minute) of the meetings, we compute N-gram counts from the corresponding text segment of the official record. Here, we adopt a turn as a processing unit, because the whole meeting (typically two to five hours) is too long, containing a variety of topics and speakers. The N-gram entries and counts are then converted to the verbatim style using the transformation model (eq. 4). Insertion of fillers and omission of particles in the N-gram chain are also modeled by considering their context in this process. Then, ASR is conducted using the dedicated model to produce a faithful transcript. The model is very constrained and still expected to accurately predict spontaneous phenomena such as filler insertion. It is also compact compared with the former methods of lightly supervised training [13][14][15][16], which interpolate the cleaned text with the baseline language model, resulting in a very large model.

With the proposed method applied to turn-based segments, we could get character accuracy of 94.3% (baseline ASR 83.1%) for the additional training data set, which is used for re-training of the acoustic model. The best phone hypothesis is used as the label for the standard HMM training based on ML (Maximum Likelihood) criterion. For discriminative training such as MPE (Minimum Phone Error) criterion, we also

generate competing hypotheses using the baseline language model.

IV. TRANSCRIPTION SYSTEM FOR THE JAPANESE PARLIAMENT AND ITS INITIAL EVALUATIONS

The ASR system was deployed as a core part of the new transcription system in the House of Representatives. By applying the scheme of language model transformation to a huge scale of the past Parliamentary meeting records (200M words in text over 12 years), a faithfully spoken-style language model is generated. Moreover, the lightly-supervised training of an acoustic model was applied to 1000 hours of speech that are not faithfully transcribed. As a result, we could build precise models of spontaneous speech in Parliament.

Figure 2 shows the relationship between the size of acoustic model training data and the ASR accuracy in character correctness for a test set of 12 meetings. It is clearly seen that the increase of the data size continuously results in the improvement of accuracy, though the improvement becomes smaller when the size gets larger. This suggests that the semi-automated training scheme is critical.

These acoustic and language models, developed by Kyoto University, have been integrated into the recognition engine or decoder of NTT Corporation, which is based on the fast on-the-fly composition of WFST (Weighted Finite State Transducers) [17]. The overview of the entire system is depicted in Figure 3. The module integration was enabled effectively and smoothly since we have developed an open-source ASR platform Julius³, which decouples acoustic and language models from the recognition engine in an open format.⁴

Speech is captured by stand microphones in meeting rooms for plenary sessions and committee meetings. Separate channels are used for interpellators and ministers. Channel selection and speaker segmentation modules, which were also developed in NTT, were also incorporated. The development of the entire system took almost one year, after the research presented in this paper was conducted over a couple of years.

³<http://julius.sourceforge.jp/>

⁴Julius itself was not adopted in the running system because it is no-warranty, and the NTT decoder performs much faster.

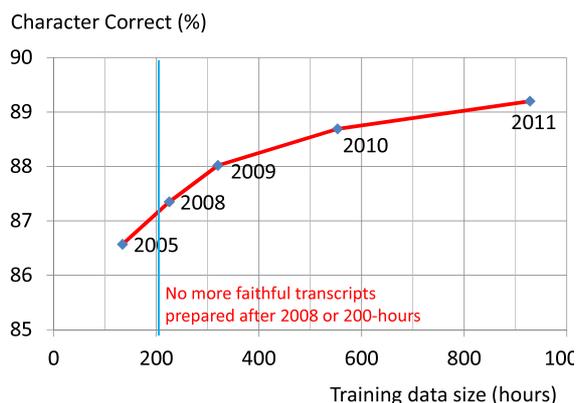


Fig. 2. Relationship of acoustic model training data size and character correctness

Trials and evaluations of the system have been conducted since the system was deployed in March 2010. The accuracy defined by the character correctness compared against the official record⁵ is 89.3% for 60 meetings done in 2010. When limited to plenary sessions, it is over 95%. No meetings got accuracy of less than 85%. The processing speed is 0.5 in real-time factor, which means it takes about 2.5 minutes for a 5-minute segment assigned to each Parliamentary reporter. The system can also automatically annotate and remove fillers, but automation of other edits is still under ongoing research [11]. Furthermore, the semi-automated update of the acoustic and language models using the data throughout the trial operations has brought an additional gain in accuracy by 0.7% absolute.

After the trials in FY 2010, the system has been in official operation from April 2011. The new system now handles all plenary sessions and committee meetings. The speaker-independent ASR system generates an initial draft, which is corrected by Parliamentary reporters. The averaged character correctness measured for 118 meetings held throughout 2011 was 89.8%, and the additional insertion errors, excluding fillers, account for 8.2%. It translates that, roughly speaking, the system's recognition error rate was around 10%, and disfluencies and colloquial expressions to be deleted or corrected also accounted for around 10%. Thus, Parliamentary reporters still play an important role although the entire transcription process is streamlined by the new system compared with the conventional short-hand scheme.

The post-editor used by Parliamentary reporters is vital for efficient correction of ASR errors and cleaning transcripts. Designed by reporters, it is a screen editor, which is similar to the word-processor interface. The editor provides easy

⁵This is slightly different from accuracy normally used in ASR evaluations, which is computed against faithful transcripts. Preparing faithful transcripts costs very much and cannot be done for a long period. On the other hand, computing character correctness against official meeting records is consistent with the goal of the system.

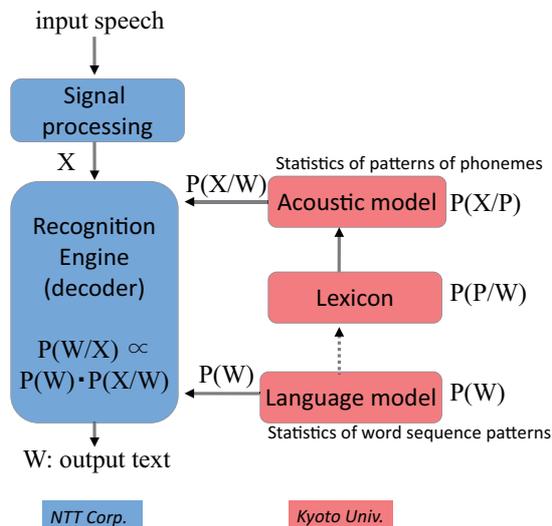


Fig. 3. Overview of the ASR system

reference to original speech and video, by time, by utterance, and by character. It can speed up and down the replay of speech. A side effect of the ASR-based system is all of text, speech, and video are aligned and hyperlinked by speakers and by utterance. It will allow for efficient search and retrieval of the multi-media archive.

V. EVALUATION AND IMPACT THROUGH LONG-TERM OPERATION

A. System Update and Long-term Evaluation

For system maintenance, we continuously monitor the ASR accuracy, and update ASR models. Specifically, the lexicon and language model are updated once a year to incorporate new words and topics. New words can be added by Parliamentary reporters at any time. The acoustic model is updated after the general election.⁶ Note that these updates can be semi-automated without manual transcription in our lightly-supervised training scheme. With this mechanism, the system has improved with more data accumulated.

Figure 4 shows the ASR accuracy plotted against the years of operation. Here, “S” stands for a spring session of usually six months, and “F” for a fall session of typically three months.⁷ The acoustic model was updated in 2013 and 2015, and the training data size became almost 2000 hours and 2700 hours in the respective update. The ASR performance was improved to over 90% after the initial operation and it has been stable close to 91% over these years. Note that since 2013 there has been no government change and only one Prime Minister in the office.

Recently, we have also incorporated the acoustic model based on Deep Neural Network (DNN). It shows a promising improvement of accuracy by 3-4% absolute, and is to be deployed in the running system.

⁶To be accurate, it is updated after collecting meeting data of one or two sessions for a couple of months after the general election

⁷There was not a fall session in 2015.

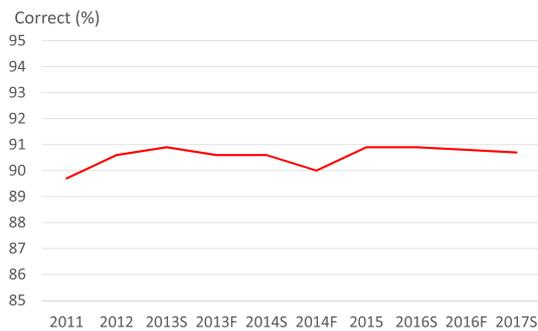


Fig. 4. ASR performance by years (sessions) since system deployment

This year, a survey was conducted among reporters in the House. A majority of the reporters expressed that it took less time and labor to make a draft with the ASR system, and more than 80% replied they are satisfied with the performance of the ASR system.

B. Impact on Parliamentary Reporting

The system was introduced to minimize the change on the reporting process of Parliament. Stenographers are engaged with the 5-minute segments of a meeting to make drafts, which are combined and proof-read by editors.

The new system, however, made it possible for those who have not been trained on stenography to produce a draft upon proper training. This is a big impact on recruiting and training of reporters in Parliament.

Another interesting impact is on how the meeting transcript is edited. In Section 2, we reported that a number of edits (16% of words on average) were made in the transcripts in ten years ago (2005-2007). When we made a similar analysis for the meetings in 2016, we found the edits are significantly reduced (10% of characters on average).

There are several causes of this reduction. Phrase reordering is not done any more. Some discourse markers are now kept, and some repeats are allowed such as those expressing emphasis. Moreover, many colloquial expressions are getting accepted. These suggest that the transcripts become more verbatim than before.

There are some possible reasons for this trend such as Internet broadcasting and references in social media. But one of the causes is deployment of the ASR-based system. Now reporters edit a faithful transcript generated by the ASR system, compared to the old scheme based on stenography when they typed in text with editing in their mind.

VI. CONCLUSIONS

This article addresses the transcription system based on our “sustainable” approach to ASR systems, which can evolve

without requiring manual transcription. It is not totally unsupervised, but assumes the final output which is generated during the system operation. With this scheme, we believe we have realized one of the highest-standard ASR system dedicated to Parliamentary meetings. Specifically, the accuracy reaches over 90% through the long-term operation. The system also made impacts on Parliamentary reporting by shifting it from manual labor to fully ICT-based work. The system has also attracted attentions from many other national Parliaments.

ACKNOWLEDGMENT

The work presented in this article is indebted to Dr. Yuya Akita and Mr. Masato Mimura of Kyoto University as well as a number of researchers and engineers of NTT Corporation.

REFERENCES

- [1] T.Kawahara. Transcription system using automatic speech recognition for the Japanese Parliament (Diet). In *Proc. AAAI/IAAI*, pages 2224–2228, 2012.
- [2] Sadaoki Furui and Tatsuya Kawahara. Transcription and distillation of spontaneous speech. In J.Benesty, M.M.Sondhi, and Y.Huang, editors, *Springer Handbook on Speech Processing and Speech Communication*, chapter 32, pages 627–651. Springer, 2008.
- [3] J.Glass, T.J. Hazen, S.Cyphers, I.Malioutov, D.Huynh, and R.Barzilay. Recent progress in the MIT spoken lecture processing project. In *Proc. InterSpeech*, pages 2553–2556, 2007.
- [4] S.Matsoukas, J.-L.Gauvain, G.Adda, T.Colthurst, Chia-Lin Kao, O.Kimball, L.Lamel, F.Lefevre, J.Z.Ma, J.Makhoul, L.Nguyen, R.Prasad, R.Schwartz, H.Schwenk, and B.Xiang. Advances in transcription of broadcast news and conversational telephone speech within the combined EARS BBN/LMSI system. *IEEE Trans. Audio, Speech & Language Process.*, 14(5):1541–1556, 2006.
- [5] S.F.Chen, B.Kingsbury, L.Mangu, D.Povey, G.Saon, H.Soltau, and G.Zweig. Advances in speech transcription at IBM under the DARPA EARS program. *IEEE Trans. Audio, Speech & Language Process.*, 14(5):1596–1608, 2006.
- [6] S.Renals, T.Hain, and H.Boulevard. Recognition and understanding of meetings: The AMI and AMIDA projects. In *Proc. IEEE Workshop Automatic Speech Recognition & Understanding*, 2007.
- [7] C.Gollan, M.Bisani, S.Kanthak, R.Schluter, and H.Ney. Cross domain automatic transcription on the TC-STAR EPPS corpus. In *Proc. IEEE-ICASSP*, volume 1, pages 825–828, 2005.
- [8] B.Ramabhadran, O.Siohan, L.Mangu, G.Zweig, M.Westphal, H.Schulz, and A.Soneiro. The IBM 2006 speech transcription system for European parliamentary speeches. In *Proc. InterSpeech*, pages 1225–1228, 2006.
- [9] Y.Akita and T.Kawahara. Efficient estimation of language model statistics of spontaneous speech via statistical transformation model. In *Proc. IEEE-ICASSP*, volume 1, pages 1049–1052, 2006.
- [10] Y.Akita and T.Kawahara. Statistical transformation of language and pronunciation models for spontaneous speech recognition. *IEEE Trans. Audio, Speech & Language Process.*, 18(6):1539–1549, 2010.
- [11] G.Neubig, Y.Akita, S.Mori, and T.Kawahara. Improved statistical models for SMT-based speaking style transformation. In *Proc. IEEE-ICASSP*, pages 5206–5209, 2010.
- [12] T.Kawahara, M.Mimura, and Y.Akita. Language model transformation applied to lightly supervised training of acoustic model for congress meetings. In *Proc. IEEE-ICASSP*, pages 3853–3856, 2009.
- [13] L.Lamel, J.Gauvain, and G.Adda. Investigating lightly supervised acoustic model training. In *Proc. IEEE-ICASSP*, volume 1, pages 477–480, 2001.
- [14] H.Y.Chan and P.Woodland. Improving broadcast news transcription by lightly supervised discriminative training. In *Proc. IEEE-ICASSP*, volume 1, pages 737–740, 2004.
- [15] L.Nguyen and B.Xiang. Light supervision in acoustic model training. In *Proc. IEEE-ICASSP*, volume 1, pages 185–188, 2004.
- [16] M.Paulik and A.Waibel. Lightly supervised acoustic model training EPPS recordings. In *Proc. InterSpeech*, pages 224–227, 2008.
- [17] T.Hori and A.Nakamura. Generalized fast on-the-fly composition algorithm for wfst-based speech recognition. In *Proc. InterSpeech*, pages 557–560, 2005.