



MOVER: Combining Multiple Meeting Recognition Systems

Naoyuki Kamo^{*,1}, Tsubasa Ochiai^{*,1}, Marc Delcroix¹, Tomohiro Nakatani¹

¹NTT Corporation, Japan
naoyuki.kamo@ntt.com

Abstract

In this paper, we propose Meeting recognizer Output Voting Error Reduction (MOVER), a novel system combination method for meeting recognition tasks. Although there are methods to combine the output of diarization (e.g., DOVER) or automatic speech recognition (ASR) systems (e.g., ROVER), MOVER is the first approach that can combine the outputs of meeting recognition systems that differ in terms of both diarization and ASR. MOVER combines hypotheses with different time intervals and speaker labels through a five-stage process that includes speaker alignment, segment grouping, word and timing combination, etc. Experimental results on the CHiME-8 DASR task and the multi-channel track of the NOTSOFAR-1 task demonstrate that MOVER can successfully combine multiple meeting recognition systems with diverse diarization and recognition outputs, achieving relative tcpWER improvements of 9.55 % and 8.51 % over the state-of-the-art systems for both tasks.

Index Terms: Meeting recognition, system combination, multi-speaker ASR, speaker diarization, CHiME-8

1. Introduction

With the advancements in deep learning technologies, the performance of automatic speech recognition (ASR) has remarkably improved [1–3]. Thus, the research interests have moved to more challenging tasks such as meeting recognition [4–9], which requires estimating *who* speaks *what* and *when*, i.e., generating transcription with speaker labels and timing information for long multi-talker recordings including overlapped speech.

System combination is often used to boost meeting recognition performance [10–18]. For example, many of the CHiME-7 and 8 submissions used some kind of system combination, e.g., [12–18]. In these challenges, many participants used system combination with diarization-first pipelines [5,6,19], which perform diarization, then separate each segment, and finally perform ASR on the separated signals. With such systems, we can combine the outputs of multiple diarization systems into a single diarization hypothesis using e.g., diarization output voting error reduction with overlap (DOVER-LAP) [20]. We can then generate multiple recognition hypotheses for each segment of the diarization hypothesis using multiple ASR back-ends. Finally, we can combine the hypotheses for each segment using recognizer output voting error reduction (ROVER) [21], minimum Bayes risk decoding [22], or N-best rescoring [23].

This approach can successfully reduce diarization and ASR errors. However, we need to combine the diarization results into a single hypothesis to make sure that the different ASR back-ends generate hypotheses for the same speech content. This

limitation prevents the correction of ASR errors due to the diarization (e.g., speaker confusion or missed detection) through the latter ASR system combination.

Other pipelines also exist for meeting recognition, such as the continuous speech separation (CSS) pipeline [7, 19, 24, 25], which performs separation, then transcribes the separated signals, and finally assigns speaker labels by clustering speaker embeddings for each word. Unlike diarization-first pipelines, it is not straightforward to combine the output of different CSS-based systems using, e.g., DOVER-LAP and ROVER, as there is no clear distinction between diarization and ASR steps.

Although it is well known that system combination is particularly effective when combining hypotheses with diverse error patterns [26], there are no approaches that can combine very distinct meeting recognition systems, such as diarization-first and CSS-based ones. To fully exploit such diversity, it is important to develop a system combination approach that can accommodate systems with different diarization and ASR outputs.

In this paper, we propose meeting recognizer output voting error reduction (MOVER) to combine the hypotheses of the multiple meeting recognition systems with different diarization and ASR outputs (i.e., transcriptions, timing information, and speaker labels). Therefore, MOVER allows the combination of meeting recognition systems that could not be combined by conventional approaches. Moreover, even for cases where we could use, e.g., DOVER-LAP and ROVER, we could improve the effect of the system combination because we can preserve the diversity of the hypotheses in terms of diarization and ASR. The proposed MOVER algorithm can combine the hypotheses relying only on the meeting recognition outputs regardless of the system architecture, which allows us to combine even black box systems such as meeting recognition APIs.

We tested the effectiveness of the proposed MOVER on tasks 1 (DASR) and 2 (NOTSOFAR-1) of the CHiME-8 challenge. We combined the meeting recognition systems of the different organizations with MOVER, and achieved relative tcpWER improvements of around 9% compared to the best system of each task. These results brought new state-of-the-art (SOTA) for these tasks, and demonstrate the effectiveness of the proposed MOVER, showing that it can combine the meeting recognition systems with diverse diarization and ASR characteristics. The scripts of the proposed MOVER are publicly released.¹

2. Related work

The NIST Scoring Toolkit (SCTK)² provides the implementation of ROVER. Although it basically assumes segmented transcriptions as input, it also has an optional mechanism that lo-

* Contributed equally

¹<https://github.com/nttclub-sp/mover>

²<https://github.com/usnistgov/SCTK>

cates common silence regions among the hypotheses to segment long transcriptions [27]. This function requires the hypotheses with the start and end timing information for each word, but many meeting recognition systems only provide segment-level timing information. Moreover, the tool supports only single-speaker recordings and does not handle the speaker ambiguities among hypotheses. Thus, it is not possible to directly apply the tool to the meeting recognition tasks, which include long multi-speaker recordings with overlapped speech.

The authors in [15] have mentioned they applied ROVER to merge meeting recognition systems with different diarization systems by matching the speaker labels of different systems and then performing ROVER on the whole text of each speaker. However, the paper is mainly about the overall meeting recognition system submitted to the CHiME-8 challenge [15]. They did not describe the details of the system combination algorithm, e.g., how to merge the timing information of the segments or words. To the best of our knowledge, this study is the first work describing a detailed algorithm extending ROVER for meeting recognition hypotheses and demonstrating its effectiveness through comprehensive experiments.

3. Proposed method

3.1. System formulation

The aim of meeting recognition is transcribing a long recording with timing information and speaker labels for all segments in the recording (i.e., who speaks when and what?). We assume that we have H meeting recognition systems producing transcriptions with different errors. Let $\mathbf{u}_{s,j}^h = (\mathbf{w}_{s,j}^h, \mathbf{v}_{s,j}^h)$ be a pair of transcription $\mathbf{w}_{s,j}^h$ and segment timing $\mathbf{v}_{s,j}^h = [b_{s,j}^h, e_{s,j}^h]$ for the j -th segment of the speaker s estimated by the meeting recognition system h , where $b_{s,j}^h$ and $e_{s,j}^h$ are the start and end times of the segment. Given a long recording, the meeting recognition system h outputs the set of all segments of all speakers $\mathbf{O}^h = \{\mathbf{o}_s^h\}_{1 \leq s \leq S^h}$, where $\mathbf{o}_s^h = \{\mathbf{u}_{s,j}^h\}_{1 \leq j \leq J_s^h}$ denotes the set of all segments associated with the speaker s . Here, S^h is the total number of speakers, and J_s^h is the number of segments of speaker s .

System combination aims at merging the output of the H meeting recognition systems as: $\mathbf{O} = \{\mathbf{o}_s\}_{1 \leq s \leq S} = \text{Combine}(\{\mathbf{O}^h\}_{1 \leq h \leq H})$, where $\text{Combine}(\cdot)$ denotes the functional representation of the system combination, and S is the number of speakers determined as a result of the combination.

3.2. Conventional: ROVER

ROVER is a classic system combination approach that merges the transcriptions of a segment estimated by multiple ASR systems. The idea of ROVER is that each system has different error patterns. Consequently, we may obtain more accurate transcriptions than any single system by combining multiple systems. Let $\mathbf{w}^h = [w_1^h, w_2^h, \dots, w_{N^h}^h]$ be the transcription (i.e., word sequence) estimated by the system h , where N^h denotes its corresponding total number of words in the segment. Since we consider only a single segment by a single speaker, we removed the s and j indices in this section.

ROVER mainly consists of two stages; 1) word alignment and 2) word voting. First in the word alignment stage, by repeatedly applying a variant of dynamic programming (DP) matching [28], we can obtain a confusion network (CN) $\{(\hat{w}_1^{h=1}, \dots, \hat{w}_1^{h=H}), \dots, (\hat{w}_{N^{\text{CN}}}^{h=1}, \dots, \hat{w}_{N^{\text{CN}}}^{h=H})\}$, which represents the word alignment among the H hypotheses. Here, N^{CN} is the number of aligned words. \hat{w}_n^h is the n -th word in




Figure 1: Processing flow of MOVER

the CN associated with hypothesis h , which corresponds to a word from the word sequence \mathbf{w}^h or the *null symbol*. Then in the word voting stage, we can select the word according to some criterion, e.g., the word with the highest count, among all words in the same position i and obtain the merged hypothesis $\mathbf{w}^{\text{ROV}} = \{w_1^{\text{ROV}}, w_2^{\text{ROV}}, \dots, w_{N^{\text{ROV}}}^{\text{ROV}}\}$, where N^{ROV} denotes the number of words determined as a result of voting.

3.3. Proposed: MOVER

As described in Section 3.2, the ROVER algorithm assumes the transcriptions estimated by the multiple systems, which correspond to the same speech content with different errors. When the systems share a diarization system (and use different ASR systems), the estimated segment boundaries are the same, and the transcriptions of each segment are aligned without any pre-processing. However, the hypotheses of meeting recognition systems with different diarization systems can have a different number of segments and also different segment boundaries. In addition, different meeting recognition systems may use different speaker label spaces; the order and number of speaker labels do not necessarily match across systems. Consequently, to apply ROVER to the meeting recognition hypotheses, we need pre-processing to align the set of segments.

Figure 1 shows the processing flow of the proposed MOVER, which comprises five stages; 1) speaker label mapping, 2) segment grouping, 3) word-level time annotation, 4) ROVER and timing combination, and 5) order consistency resolution (OCR).

1) Speaker label mapping: First, we align the transcriptions uttered by the same speaker across all systems. To obtain such speaker-aligned hypotheses, we re-map the system-dependent speaker label space s to the common label space s' using $s' = \text{Map}(s, h)$, where $\text{Map}(\cdot)$ denotes the functional representation of the speaker label mapping. If system h has no output for speaker s' , we set $\mathbf{o}_{s'}^h = \{(\phi, \varphi)\}$, where ϕ and φ denote the *empty text and timing*, which is not judged to overlap with any segment.

In this study, we adopted the global speaker mapping algorithm of the DOVER [29], which is based on the similarities between speaker time activities for the whole recording.

2) Segment grouping: In the second step, we group together transcriptions considered to be transcribing the same speech content. We considered two approaches; 1) re-grouping the multiple segments and applying ROVER to each group (subset grouping) and 2) applying ROVER to the hypotheses concatenating all the segments over the entire recording (full set grouping).

The subset grouping divides the segments into groups as $\{(\mathbf{g}_{s,k=1}^{h=1}, \dots, \mathbf{g}_{s,k=1}^{h=H}), \dots, (\mathbf{g}_{s,k=K}^{h=1}, \dots, \mathbf{g}_{s,k=K}^{h=H})\}$, where K denotes the number of the re-grouped segments. Here, $\mathbf{g}_{s,k}^h = \{\mathbf{u}_{s,j=l_k^h}^h, \dots, \mathbf{u}_{s,j=m_k^h}^h\}$ represents the re-grouped segment. l_k^h and m_k^h denotes the start and end segment indices for the k -th re-grouped segment of the hypothesis h .

Algorithm 1 Segment grouping

Input: $\{(\mathbf{w}_{s,j}^h, [b_{s,j}^h, e_{s,j}^h])\}_{(h,j) \in \Omega_s}$ \triangleright Set of segments
Output: $\tilde{\mathbf{G}}$ \triangleright Set of grouped segments
 $\tilde{\mathbf{G}} \leftarrow \phi, (\tilde{\mathbf{g}}^1, \dots, \tilde{\mathbf{g}}^H) \leftarrow (\phi, \dots, \phi), \tilde{e} \leftarrow 0$
for h, j **in** $\text{ARGSORT}_{(h,j)}(\{b_{s,j}^h\}_{(h,j) \in \Omega_s})$ **do** \triangleright Sort by start time
if $(\tilde{e} = 0)$ **or** $(b_{s,j}^h < \tilde{e})$ **then** \triangleright If empty-sets or overlapped
 $\tilde{\mathbf{g}}^h \leftarrow \tilde{\mathbf{g}}^h \cup \{(\mathbf{w}_{s,j}^h, [b_{s,j}^h, e_{s,j}^h])\}$
 $\tilde{e} \leftarrow \max(\tilde{e}, e_{s,j}^h)$ \triangleright Get the maximum end time of segments
else \triangleright If not overlapped
 $\tilde{\mathbf{G}} \leftarrow \tilde{\mathbf{G}} \cup \{(\tilde{\mathbf{g}}^1, \dots, \tilde{\mathbf{g}}^H)\}$ \triangleright Append the group
 $(\tilde{\mathbf{g}}^1, \dots, \tilde{\mathbf{g}}^H) \leftarrow (\phi, \dots, \phi), \tilde{e} \leftarrow 0$ \triangleright Initialize
end if
end for

In our implementation, we group the segments over the hypotheses so that a segment in one group overlaps any segments in the group but does not overlap segments in other groups. Algorithm 1 describes the procedure of the subset grouping, where Ω_s is the set of all indices (h, j) of the hypotheses and segments for the speaker s , except empty hypotheses. $\text{ARGSORT}(\cdot)$ returns the indices (h, j) in ascending order of the start time of the segment.

On the other hand, the full set grouping simply makes a single group containing all the segments as $\{(\hat{\mathbf{g}}_s^{h=1}, \dots, \hat{\mathbf{g}}_s^{h=H})\}$, where $\hat{\mathbf{g}}_s^h = \{\mathbf{u}_{s,j=1}^h, \dots, \mathbf{u}_{s,j=N_{s,j}^h}^h\}$.

3) Word-level time annotation: MOVER utilizes the word-level timing information for the ROVER-based merging procedure of the transcription and timing information. Meeting recognition systems do not always provide word-level timing (only segment-level timing). Thus, when the hypotheses do not have word-level time information, we require assigning it by, e.g., applying forced alignment [30], pseudo-word-level annotation [31], etc. The segment-level timing information $[b_{s,j}^h, e_{s,j}^h]$ is converted to word-level one $[b_{s,j,n=1}^h, e_{s,j,n=1}^h], \dots, [b_{s,j,n=N_{s,j}^h}^h, e_{s,j,n=N_{s,j}^h}^h]$, where $b_{s,j,n}^h$ and $e_{s,j,n}^h$ are the start and end times of the n -th word of the j -th segment, and $N_{s,j}^h$ is the number of words in it.

In this study, we adopted the pseudo-word-level annotations [31] with *character-based* option for each segment, which assigns the time interval for each word in proportion to its number of characters.

4) ROVER and timing combination: Finally, given the set of grouped segments $(\mathbf{g}_{s,k}^{h=1}, \dots, \mathbf{g}_{s,k}^{h=H})$, we apply ROVER to each group of segments and obtain the merged meeting recognition hypothesis $\mathbf{g}_{s,k}^{\text{ROV}} = (\mathbf{w}_{s,k}^{\text{ROV}}, \mathbf{v}_{s,k}^{\text{ROV}})$. By applying DP matching to the hypotheses, we can obtain the aligned words $(\hat{w}_{s,j,n}^{h=1}, \dots, \hat{w}_{s,j,n}^{h=H})$ and their timing information $([b_{s,j,n}^{h=1}, e_{s,j,n}^{h=1}], \dots, [b_{s,j,n}^{h=H}, e_{s,j,n}^{h=H}])$. Then, we obtain the merged transcription $\mathbf{w}_{s,k}^{\text{ROV}}$ by the word voting as for ROVER (see Section 3.2). In our implementation, we simply select the most frequent word for voting. When the null symbol is selected by the voting, it is removed from the merged hypothesis.

We also merge the timing information for each word as $b_{s,j,n}^{\text{ROV}} = \frac{1}{N_{s,j,n}} \sum_{h \in \mathcal{H}_{s,j,n}} b_{s,j,n}^h$ and $e_{s,j,n}^{\text{ROV}} = \frac{1}{N_{s,j,n}} \sum_{h \in \mathcal{H}_{s,j,n}} e_{s,j,n}^h$, where $\mathcal{H}_{s,j,n}$ is the set of hypotheses that have the same word as that selected by the voting process and $N_{s,j,n}$ is the number of hypotheses in $\mathcal{H}_{s,j,n}$.

As a result of the segment grouping, the duration of the re-grouped segment may become very long (especially, for the full set grouping), which could cause the word alignments whose word timings are temporally distant and may cause a negative impact on the system combination. To prevent such implausible alignments, we proposed to extend the DP matching pro-

Algorithm 2 Order consistency resolution

Input: $\{(\mathbf{w}_{s,j}, [b_{s,j}, e_{s,j}])\}_{1 \leq j \leq J}$ \triangleright Set of segments of merged hypothesis
Output: $\{(\tilde{\mathbf{w}}_j, [\tilde{b}_j, \tilde{e}_j])\}_{1 \leq j \leq \tilde{J}}$ \triangleright Set of revised segments
 $\{(\tilde{\mathbf{w}}_j, [\tilde{b}_j, \tilde{e}_j])\}_{1 \leq j \leq J} \leftarrow \{(\mathbf{w}_{s,j}, [b_{s,j}, e_{s,j}])\}_{1 \leq j \leq J}$ \triangleright Copy input
 $j \leftarrow 1, \tilde{J} \leftarrow J$
while $j < \tilde{J}$ **do**
if $\tilde{b}_{j+1} < \tilde{e}_j$ **then** \triangleright If order is inconsistent or overlapped
 $[\tilde{b}, \tilde{e}] \leftarrow [\min(\tilde{b}_j, \tilde{b}_{j+1}), \max(\tilde{e}_j, \tilde{e}_{j+1})]$ \triangleright Merge two segments
 $\tilde{\mathbf{w}} \leftarrow [\tilde{\mathbf{w}}_j, \tilde{\mathbf{w}}_{j+1}]$ \triangleright Merge two segments
 $(\tilde{\mathbf{w}}_j, [\tilde{b}_j, \tilde{e}_j]) \leftarrow (\tilde{\mathbf{w}}, [\tilde{b}, \tilde{e}])$ \triangleright Replace old to new element
 $\text{delete } (\tilde{\mathbf{w}}_{j+1}, [\tilde{b}_{j+1}, \tilde{e}_{j+1}])$ \triangleright Remove old element from set
 $j \leftarrow 1, \tilde{J} \leftarrow \tilde{J} - 1$ \triangleright Go to start of the loop
else
 $j \leftarrow j + 1$
end if
end while

cedure of ROVER’s word alignment stage by incorporating the time-constraint mechanism (TC-DP), like tcpWER [31]. Unlike the DP for a pair of hypothesis and reference transcriptions in the tcpWER, ROVER requires performing the DP between the CN (aligned hypotheses) and a hypothesis. When we compute the score (i.e., *correct*, *substitution*, *insertion*, and *deletion*) between a word in the hypothesis and one of the words in the CN, we apply a time constraint, allowing *correct* and *substitution* only if the word timings overlap. To allow for a certain amount of time gaps between hypotheses, we adopted a collar c like [31], which we set to 5 seconds based on our preliminary experiments.

5) Order consistency resolution: The simple merging procedure of the timing information can cause a mismatch between the order of the words and their timing information, which may change the order of the words from that obtained with ROVER. Here, we hypothesize that the word order of ROVER is more accurate and propose a post-processing procedure called OCR to correct the timing information.

Algorithm 2 describes the OCR procedure, where **delete** stands for removing an element from a list. The idea of OCR is to find order inconsistent words that have an end time after the start time of the next word, concatenate these words, and re-assign the timing information of the concatenated words. In this study, we applied the OCR to the merged hypothesis $\{(\mathbf{w}_{s,j=1}^{\text{ROV}}, \mathbf{v}_{s,j=1}^{\text{ROV}}), \dots, (\mathbf{w}_{s,j=J}^{\text{ROV}}, \mathbf{v}_{s,j=J}^{\text{ROV}})\}$ for each speaker s , where J is the total number of the segments in the entire hypothesis. Each of the element $(\mathbf{w}_{s,j}, \mathbf{v}_{s,j})$ contains only one word $\mathbf{w}_{s,j} = [w_{s,j}]$ and its word-level timing information $\mathbf{v}_{s,j} = [b_{s,j}, e_{s,j}]$. Here, we assume that the words are arranged in the order of ROVER’s output, and $w_{s,j}$ is the j -th word in the entire transcription.

4. Experiments

4.1. Experimental setup

We evaluated the effectiveness of the proposed MOVER on task-1 (DASR) and the multi-channel track of task-2 (NOTSOFAR/mc) of the recent CHiME-8 Challenge. The evaluation set of task 1 comprises four datasets, i.e., CHiME-6 [32], DiPCO [33], Mixer 6 [34], and NOTSOFAR-1 [7], while that of the Task 2 focuses on the NOTSOFAR-1 dataset. Following CHiME-8’s regulation, we evaluated the performance using the tcpWER with a collar of 5 seconds. All texts in the hypotheses were normalized using the CHiME-8 text normalization tool, which includes unifying numerical representations, removing fillers, etc [6].

For our experiments, we used the system outputs provided

Table 1: Comparison of the tcpWER [%] on the eval set of NOT-SOFAR/mc for the 1-best system and different configurations of MOVER for the combination of nine systems.

System	Grouping	OCR	TC-DP	Eval
1-Best	-	-	-	10.81
				11.29
MOVER	Full set	✓		10.26
		✓	✓	9.89
	Subset			10.81
		✓	✓	9.95
		✓	9.96	




Figure 2: tcpWER [%] when combining different systems with MOVER. Each bar represents the performance of the combined system (the pink one is the best). The horizontal lines show the performance of the individual systems. The numbers at the bottom indicate the number of systems from each team (i.e., from left to right: USTC-NERC SLIP, STCON, NTT, and NPU-TEA).

by the participants. For task 1, we combined systems from the two participating teams (STCON [13] and NTT [14]), and two additional post-challenge systems [35]. For task 2, we used the systems from four teams, i.e., USTC-NERC SLIP [12], STCON [13], NTT [14], and NPU-TEA [15].

4.2. Experimental result of Task 2 (NOTSOFAR/mc)

Table 1 shows the results of the NOTSOFAR/mc task. We selected nine systems among all CHiME-8 submissions for combination, based on the performance on the development set; four systems from USTC-NERC SLIP [12], two systems from STCON [13], two systems from NTT [14], and one system from NPU-TEA [15]. We evaluated the three options of MOVER, i.e., Full set or Subset “Grouping”, without or with “TC-DP”, and without or with “OCR”. As a baseline, the table includes the USTC-NERC SLIP/sys1, which achieved the best performance on the evaluation set using a single system (1-Best).

From Table 1, we observe that MOVER can improve performance over the 1-best system for both full set and subset grouping. The ablation results show that resolving the consistency of the word orders with OCR is essential for MOVER to outperform the 1-Best system. TC-DP does not improve the performance for the subset grouping, while it is essential for the full set grouping. This result is reasonable, because the DP matching for the long segment may cause implausible word alignments with a large time discrepancy. Overall, MOVER+TC-DP+OCR with Full set grouping achieved the best performance and a 0.9-point improvement compared to the baseline.³ The difference between 1-Best and the best MOVER approaches statistical significance according to a one-sided Student’s t-test (p

³We also evaluated system combination with SCKT ROVER using speaker label mapping and pseudo-word-level annotation. It failed to improve performance over the 1-best, achieving performance comparable to that of MOVER with full set grouping without OCR and TC-DP (11.32 vs. 11.29 %).

Table 2: tcpWER [%] results on the eval set of the DASR task.

	CHiME-6	DiPCo	Mixer 6	NOTSOFAR-1	Macro-Ave
NTT/best	31.99	21.68	13.74	13.83	20.31
STCON/sys2	33.13	19.88	10.87	14.61	19.62
MOVER	30.32	18.61	12.28	12.97	18.54

= 0.057).⁴ In the following, we report only the results of this configuration.

Note that the performance after applying MOVER (i.e., 9.89 %) is the new SOTA performance of the NOTSOFAR-1/mc task, outperforming the top challenge system (i.e., 10.81 %, NERC SLIP/sys1). These results demonstrate that the proposed MOVER can successfully merge diverse meeting recognition systems, designed by different organizations.

Figure 2 shows the results obtained by combining different systems with MOVER. The results show that MOVER stably improved performance for various combinations of the systems. Depending on the combined systems, the performance improvement varies. For example, combining the 4 best systems did not improve performance, probably because of the lack of variety, as they were all built by the same team (USTC-NERC SLIP). The best performance was achieved by combining nine systems from four teams, but interestingly, a very close performance could be achieved by combining only each team’s best system.

These results suggest that it is more effective to merge systems with diverse diarization and ASR characteristics than to merge systems with good performance but similar characteristics. These results re-confirm the potential of MOVER as it allows merging systems with different diarization results, resulting in more diverse ASR results.

4.3. Experimental result of Task 1 (DASR)

We evaluated the effectiveness of MOVER on four different datasets of task 1 (DASR) of CHiME-8. Only two teams (STCON and NTT) submitted a total of five systems. In addition, we included the output of two more systems built after the challenge [35], bringing the total to seven systems.

Table 2 compares the best-performing systems of STCON and NTT with the results obtained with MOVER. MOVER improves the macro-averaged tcpWER by 1.1 points over the SOTA (i.e., 19.6 %, STCON/sys2). Except for Mixer-6, MOVER stably achieved the performance gains over the best-performing single systems. These results confirm the effectiveness of the proposed method not only for NOTSOFAR-1 but also for other meeting recognition datasets. The STCON and NTT systems utilized system combination methods for both diarization and ASR outputs. Even for such systems, the proposed MOVER obtained further performance gains.

5. Conclusion

In this paper, we proposed a novel system combination method for meeting recognition, called MOVER. MOVER is a voting-based method that combines multiple hypotheses with different transcriptions, timing information, and speaker labels. With the proposed MOVER, we could combine systems from different CHiME-8 participants, achieving new SOTA performance on both the DASR and NOTSOFAR/mc tasks, demonstrating its effectiveness. The proposed implementation of MOVER uses global speaker mapping. In future work, we would like to explore using local speaker mapping to allow further speaker confusion error correction based on textual information.

⁴The tcpWER is evaluated for each session, and thus the number of samples is small (i.e., 160 samples), which may explain why the test did not reach a smaller p-value.

6. References

- [1] R. Prabhavalkar, T. Hori, T. N. Sainath, R. Schlüter, and S. Watanabe, “End-to-end speech recognition: A survey,” *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, 2023.
- [2] A. Radford, J. W. Kim, T. Xu, G. Brockman, C. Mcleavey, and I. Sutskever, “Robust speech recognition via large-scale weak supervision,” in *Proc. ICML*, vol. 202, 2023, pp. 28 492–28 518.
- [3] R. Haeb-Umbach, J. Heymann, L. Drude, S. Watanabe, M. Delcroix, and T. Nakatani, “Far-field automatic speech recognition,” *Proc. IEEE*, vol. 109, no. 2, pp. 124–148, 2020.
- [4] W. Kraaij, T. Hain, M. Lincoln, and W. Post, “The AMI meeting corpus,” in *Proc. International Conference on Methods and Techniques in Behavioral Research*, 2005, pp. 1–4.
- [5] S. Cornell, M. S. Wiesner, S. Watanabe, D. Raj, X. Chang, P. Garcia *et al.*, “The CHiME-7 DADR challenge: Distant meeting transcription with multiple devices in diverse scenarios,” in *Proc. International Workshop on Speech Processing in Everyday Environments (CHiME 2023)*, 2023, pp. 1–6.
- [6] S. Cornell, T. J. Park, H. Huang, C. Boeddeker, X. Chang, M. Maciejewski *et al.*, “The CHiME-8 DADR challenge for generalizable and array agnostic distant automatic speech recognition and diarization,” in *Proc. International Workshop on Speech Processing in Everyday Environments (CHiME 2024)*, 2024, pp. 1–6.
- [7] A. Vinnikov, A. Ivry, A. Hurvitz, I. Abramovski, S. Koubi, I. Gurvich *et al.*, “NOTSOFAR-1 challenge: New datasets, baseline, and tasks for distant meeting transcription,” in *Proc. Interspeech*, 2024, pp. 5003–5007.
- [8] Y. Fu, L. Cheng, S. Lv, Y. Jv, Y. Kong, Z. Chen, Y. Hu, L. Xie, J. Wu, H. Bu, X. Xu, J. Du, and J. Chen, “Aishell-4: An open source dataset for speech enhancement, separation, recognition and speaker diarization in conference scenario,” in *Interspeech 2021*, 2021, pp. 3665–3669.
- [9] F. Yu, S. Zhang, Y. Fu, L. Xie, S. Zheng, Z. Du *et al.*, “M2MeT: The ICASSP 2022 multi-channel multi-party meeting transcription challenge,” in *Proc. ICASSP*, 2022, pp. 6167–6171.
- [10] F. Yu, S. Zhang, P. Guo, Y. Fu, Z. Du, S. Zheng *et al.*, “Summary on the ICASSP 2022 multi-channel multi-party meeting transcription grand challenge,” in *Proc. ICASSP*, 2022, pp. 9156–9160.
- [11] Y. Liang, M. Shi, F. Yu, Y. Li, S. Zhang, Z. Du *et al.*, “The second multi-channel multi-party meeting transcription challenge (M2MeT 2.0): A benchmark for speaker-attributed ASR,” in *Proc. ASRU*, 2023, pp. 1–8.
- [12] S. Niu, R. Wang, J. Du, G. Yang, Y. Tu, S. Wu *et al.*, “The USTC-NERCSLIP systems for the CHiME-8 NOTSOFAR-1 challenge,” in *Proc. International Workshop on Speech Processing in Everyday Environments (CHiME 2024)*, 2024, pp. 31–36.
- [13] A. Mitrofanov, T. Prisyach, T. Timofeeva, S. Novoselov, M. Korenevsky, Y. Khokhlov *et al.*, “STCON system for the CHiME-8 challenge,” in *Proc. International Workshop on Speech Processing in Everyday Environments (CHiME 2024)*, 2024, pp. 13–17.
- [14] N. Kamo, N. Tawara, A. Ando, T. Kano, H. Sato, R. Ikeshita *et al.*, “NTT Multi-Speaker ASR System for the DADR Task of CHiME-8 Challenge,” in *Proc. International Workshop on Speech Processing in Everyday Environments (CHiME 2024)*, 2024, pp. 69–74.
- [15] K. Huang, Y. Li, Z. Wang, H. Wang, W. Rao, Z. Sun *et al.*, “The NPU-TEA system for the CHiME-8 NOTSOFAR-1 challenge,” in *Proc. International Workshop on Speech Processing in Everyday Environments (CHiME 2024)*, 2024, pp. 45–48.
- [16] C. B. Boeddeker, T. Cord-Landwehr, T. von Neumann, and R. Haeb-Umbach, “Multi-stage diarization refinement for the CHiME-7 DADR scenario,” in *Proc. International Workshop on Speech Processing in Everyday Environments (CHiME 2023)*, 2023, pp. 51–56.
- [17] M. Karafiat, K. Veselý, I. Szoke, L. Mosner, K. Benes, M. Witkowski *et al.*, “BUT CHiME-7 system description,” in *Proc. International Workshop on Speech Processing in Everyday Environments (CHiME 2023)*, 2023, pp. 67–72.
- [18] K. Deng, X. Zheng, and P. Woodland, “The university of Cambridge system for the CHiME-7 DADR task,” in *Proc. International Workshop on Speech Processing in Everyday Environments (CHiME 2023)*, 2023, pp. 73–76.
- [19] D. Raj, P. Denisov, Z. Chen, H. Erdogan, Z. Huang, M. He *et al.*, “Integration of speech separation, diarization, and recognition for multi-speaker meetings: System description, comparison, and analysis,” in *Proc. SLT*, 2021, pp. 897–904.
- [20] D. Raj, L. P. Garcia-Perera, Z. Huang, S. Watanabe, D. Povey, A. Stolcke *et al.*, “DOVER-Lap: A method for combining overlap-aware diarization outputs,” *Proc. SLT*, 2021.
- [21] J. G. Fiscus, “A post-processing system to yield reduced word error rates: Recognizer Output Voting Error Reduction (ROVER),” in *Proc. ASRU*, 1997, pp. 347–354.
- [22] H. Xu, D. Povey, L. Mangu, and J. Zhu, “Minimum Bayes risk decoding and system combination based on a recursion for edit distance,” *Computer Speech & Language*, vol. 25, no. 4, pp. 802–828, 2011.
- [23] B. Roark, M. Saraclar, and M. Collins, “Discriminative n-gram language modeling,” *Computer Speech & Language*, vol. 21, no. 2, pp. 373–392, 2007.
- [24] T. Yoshioka, I. Abramovski, C. Aksoylar, Z. Chen, M. David, D. Dimitriadis *et al.*, “Advances in online audio-visual meeting transcription,” in *Proc. ASRU*, 2019, pp. 276–283.
- [25] T. Von Neumann, C. Boeddeker, T. Cord-Landwehr, M. Delcroix, and R. Haeb-Umbach, “Meeting recognition with continuous speech separation and transcription-supported diarization,” in *Proc. ICASSPW*, 2024, pp. 775–779.
- [26] Z.-H. Zhou, *Ensemble Methods: Foundations and Algorithms*. Chapman and Hall/CRC, 2012.
- [27] J. S. Garofolo, J. G. Fiscus, and C. D. Laprun, *The rich transcription 2004 spring meeting recognition evaluation*. US Department of Commerce, National Institute of Standards and Technology, 2004.
- [28] D. Sankoff and J. B. Kruskal, *Time Warps, String Edits, and Macromolecules: The Theory and Practice of Sequence Comparison*. Addison-Wesley, 1983.
- [29] A. Stolcke and T. Yoshioka, “DOVER: A method for combining diarization outputs,” in *Proc. ASRU*, 2019, pp. 757–763.
- [30] M. McAuliffe, M. Socolof, S. Mihuc, M. Wagner, and M. Sonderegger, “Montreal forced aligner: Trainable text-speech alignment using kaldii,” in *Proc. Interspeech*, 2017, pp. 498–502.
- [31] T. von Neumann, C. Boeddeker, M. Delcroix, and R. Haeb-Umbach, “MeetEval: A toolkit for computation of word error rates for meeting transcription systems,” in *Proc. International Workshop on Speech Processing in Everyday Environments (CHiME 2023)*, 2023, pp. 27–32.
- [32] S. Watanabe, M. Mandel, J. Barker, E. Vincent, A. Arora, X. Chang *et al.*, “CHiME-6 challenge: Tackling multispeaker speech recognition for unsegmented recordings,” in *Proc. International Workshop on Speech Processing in Everyday Environments (CHiME 2020)*, 2020, pp. 1–7.
- [33] M. V. Segbroeck, Z. Ahmed *et al.*, “DiPCo—Dinner Party Corpus,” in *Proc. Interspeech*, 2020, pp. 434–436.
- [34] L. Brandschain, D. Graff, C. Cieri, K. Walker, C. Caruso, and A. Neely, “The Mixer 6 corpus: Resources for cross-channel and text independent speaker recognition,” in *Proc. LREC*, 2010, pp. 2441–2444.
- [35] N. Kamo, N. Tawara, A. Ando, T. Kano, H. Sato, R. Ikeshita *et al.*, “Microphone array geometry independent multi-talker distant ASR: NTT system for the DADR task of the CHiME-8 challenge,” *arXiv preprint arXiv:2502.09859*, 2025.