

# Towards Speech Dialogue Translation Mediating Speakers of Different Languages

Shuichiro Shimizu<sup>1</sup> Chenhui Chu<sup>1</sup> Sheng Li<sup>2</sup> Sadao Kurohashi<sup>1,3</sup>

<sup>1</sup>Kyoto University, Japan

<sup>2</sup>National Institute of Information and Communications Technology, Japan

<sup>3</sup>National Institute of Informatics, Japan

{sshimizu, chu, kuro}@nlp.ist.i.kyoto-u.ac.jp sheng.li@nict.go.jp

## Abstract

We present a new task, speech dialogue translation mediating speakers of different languages. We construct the SpeechBSD dataset for the task and conduct baseline experiments. Furthermore, we consider context to be an important aspect that needs to be addressed in this task and propose two ways of utilizing context, namely monolingual context and bilingual context. We conduct cascaded speech translation experiments using Whisper and mBART, and show that bilingual context performs better in our settings.

## 1 Introduction

In this global era, it is becoming increasingly important for people from different countries/regions to interact with each other and have a mutual understanding. Recent advancements in machine translation (MT) technologies have enabled us to communicate with people worldwide, especially in text. Chat translation or dialogue machine translation (Liu et al., 2021) supports such communications, which enables people who use different languages to have cross-language chats. Speech translation (ST) has also recently shown success (e.g., Chen et al., 2022), especially in monologue translation (e.g., Di Gangi et al., 2019). However, to the best of our knowledge, no study has focused on ST of dialogues, which is an important aspect of language usage.

In this study, we propose a new task: speech dialogue translation (SDT) aiming to mediate speakers of different languages. We consider bilingual dialogues where several people who speak in different languages talk with each other mediated by an ST system.

It is important to consider context in SDT because we need to consider context in different languages, which cannot be readily handled by current ST systems that mainly focus on one translation

direction. Figure 1 shows an example of an ST-mediated dialogue between an English speaker and a Japanese speaker. They are discussing some ideas, and the English speaker says, “What do you think about it?” The Japanese speaker responds by saying the idea is naive, but without context it can be translated as “I think it’s a bit sweet” because “甘い” has two meanings, sweet and naive. By utilizing dialogue context, the meaning of “甘い” becomes clear so that the utterance can be translated properly.

For the proposed task, we construct the SpeechBSD dataset<sup>1</sup> based on an existing text dialogue corpus, BSD (Business Scene Dialogue) corpus (Rikters et al., 2019). We collect audio of the BSD corpus through crowdsourcing along with speaker attributes.

We conduct speech-to-text cascaded ST experiments on the dataset. There are two mainstream methods for ST, the cascade method (Stentford and Steer, 1988) where automatic speech recognition (ASR) and MT are chained together, and the end-to-end method (Duong et al., 2016; Bernard et al., 2016), where translations are directly predicted from speech. Recent study (Bentivogli et al., 2021; Tran et al., 2022) suggests that the two methods are on par. We conduct cascade ST experiments using Whisper (Radford et al., 2022) for ASR and mBART (Liu et al., 2020) for MT.

We consider three settings for translation: without context, with monolingual context, and with bilingual context. The monolingual context is composed in the language the utterance to be translated is spoken, whereas the bilingual context is composed in the original language of the spoken utterances (see examples in Figure 1). We show that translation with bilingual context performs better compared to the one without context by up to 1.9 BLEU points in MT and 1.7 BLEU points in

<sup>1</sup>The dataset is made public under the CC BY-NC-SA 4.0 license at <https://github.com/ku-nlp/speechbsd>.

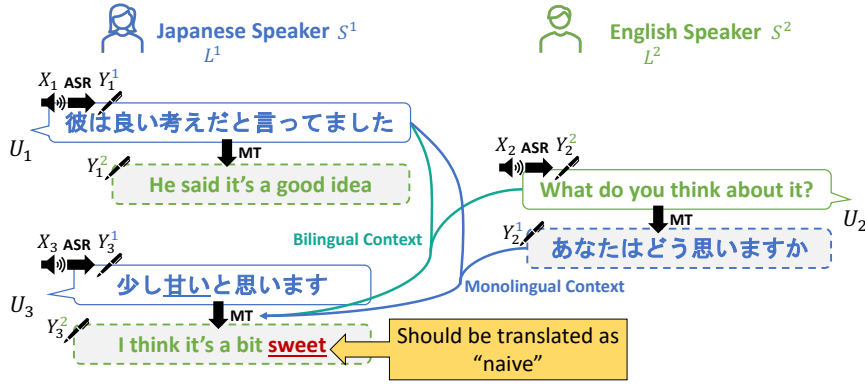


Figure 1: The importance of considering context in SDT. “甘い” can be translated into either “sweet” or “naive,” which can be disambiguated with the context. We consider two types of context for translation, monolingual context and bilingual context.

cascade ST with our settings. We also conduct a manual evaluation focusing on zero anaphora, a grammatical phenomenon where arguments of verbs are omitted when they are apparent from the context in Japanese. We show that with bilingual context, the MT models can often predict zero pronouns correctly.

## 2 Related Work

Although neural MT has greatly improved over the past few years, the translation of dialogues remains a challenging task because of its characteristics. Liu et al. (2021) summarizes the recent progress of dialogue MT and categorizes its issue into four categories, coherence, consistency, cohesion, and personality. The main approaches to address these problems include document MT (e.g., Liu et al., 2021), usage of pretrained models (e.g., Wang et al., 2020), and auxiliary task learning utilizing speaker information (e.g., Liang et al., 2021).

Considering context in ST is recently studied for the end-to-end approach (Zhang et al., 2021). We point out that although not addressed in this work, considering context for ASR is also an active research area (e.g., Inaguma and Kawahara, 2021).

In this work, we focus on the translation of speech dialogue. We use mBART, which performed best in a previous work of chat translation (Liu et al., 2021), and also consider utilizing context.

## 3 Speech Dialogue Translation (SDT)

In SDT, there are several speakers who speak different languages with the help of a translation system. In this work, we consider  $M$  speak-

ers  $\{S^m \mid m = 1, 2, \dots, M\}$  and 2 languages  $\{L^n \mid n = 1, 2\}$ . We consider a dialogue with  $T$  utterances  $D = (U_1, \dots, U_T)$ , where an utterance is  $U_t = (S_t^m, L_t^n, X_t)$ . Here,  $S_t^m$  is the speaker,  $L_t^n$  is the language spoken, and  $X_t$  is the speech signal of  $t$ -th utterance. Let  $Y_t^n$  ( $n = 1, 2$ ) be text that has the same meaning as  $X_t$  in language  $L^n$ . The task of SDT is to generate translation  $Y_t^2$  from speech signal  $X_t$  when the source language is  $L^1$  (or translation  $Y_t^1$  from  $X_t$  when the source language is  $L^2$ ) for every utterance  $U_t$ .

## 4 SpeechBSD Dataset

We construct the SpeechBSD dataset to study SDT. It is based on the existing dialogue dataset in text, BSD corpus (Rikters et al., 2019, 2021). We collect audio of all the sentences in the dataset along with speaker attributes (gender and homeplace) through crowdsourcing.

### 4.1 BSD Corpus

BSD corpus is a parallel corpus of English and Japanese composed of manually designed business scene dialogues. Each dialogue called scenario contains 30 sentences on average spoken by 2-5 speakers. The original language the scenarios were written in is half English and half Japanese so that the expressions are not biased toward one language.

### 4.2 Dataset Construction

First, we divided each scenario by speaker. For example in Figure 1, the original BSD corpus contains text of  $Y_1^1, Y_1^2, Y_2^1, Y_2^2, Y_3^1$ , and  $Y_3^2$ . In this case, we divide the dialogue into four parts: the Japanese speaker part ( $Y_1^1$  and  $Y_3^1$ ), the English speaker part ( $Y_2^2$ ), another Japanese speaker part

( $Y_2^1$ ), and another English speaker part ( $Y_1^2$  and  $Y_3^2$ ). In this way, we can compose two cross-language dialogues ( $Y_1^1 \rightarrow Y_2^2 \rightarrow Y_3^1$  and  $Y_1^2 \rightarrow Y_2^1 \rightarrow Y_3^2$ ) from one scenario of the BSD corpus. We collected audio through crowdsourcing so that each part is spoken by a different worker.<sup>2</sup> We designed a web application to record audio and collected English speech from the US using Amazon Mechanical Turk<sup>3</sup> and Japanese speech from Japan using Yahoo! crowdsourcing.<sup>4</sup> We also collected the gender and homeplace (the US states or Japanese prefecture) of the speakers as they may affect translation performance. The instructions given to the workers are shown in Appendix A.1.

### 4.3 Statistics of the SpeechBSD Dataset

The collected audio was 24.3 hours for English speech and 30.7 hours for Japanese speech in total. Details are provided in Appendix B Table 2. Regarding speaker gender, English speech was balanced, whereas there were more male speakers in Japanese. As for homeplace, in Japanese, the speakers were distributed roughly according to the population distribution. In English, it was less diverse (Appendix B Figure 3).

## 5 Considering Context for SDT

We propose two ways to consider context in SDT: monolingual context and bilingual context.

First, for every utterance  $U_t$ , an ASR system is used to obtain transcripts  $Y_t^n$ .

The monolingual context is composed in the source language of the utterance to be translated. For example, in Figure 1, when translating the third utterance  $U_3$  from Japanese to English, as the source language of the utterance is Japanese ( $L^1$ ), the context ( $Y_1^1$  and  $Y_2^1$ ) is also composed in Japanese. Let the context composed in this way be  $\mathbf{Y}_{<t}^n$ .

For monolingual context experiments, we use two translation models for each translation direction. The training objective of the MT model that translates from  $L^1$  to  $L^2$  is to maximize the follow-

ing log likelihood<sup>5</sup>:

$$\mathcal{L}^{1 \rightarrow 2} = \sum_t \log P(Y_t^2, \mathbf{Y}_{<t}^2 | Y_t^1, \mathbf{Y}_{<t}^1). \quad (1)$$

Similar objective  $\mathcal{L}^{2 \rightarrow 1}$  can be derived when  $L^2$  is the source language and  $L^1$  is the target language. Postprocessing is applied to extract  $Y_t^2$  from the output that contains both  $\mathbf{Y}_{<t}^2$  and  $Y_t^2$ .

The bilingual context is composed of the original language of the spoken utterances. For example, in Figure 1, when translating the third utterance  $U_3$  from Japanese to English, the bilingual context on the source side is  $Y_1^1$  and  $Y_2^2$ , which involves both languages. The bilingual context on the target side is  $Y_1^2$  and  $Y_2^1$ . Because there is no concept of source or target language in this case, let the source side utterance be  $Y_t$ , source side context be  $\mathbf{Y}_{<t}$ , target side utterance be  $\bar{Y}_t$ , and target side context be  $\bar{\mathbf{Y}}_{<t}$ . The MT model is trained with the following objective:

$$\mathcal{L} = \sum_t \log P(\bar{Y}_t, \bar{\mathbf{Y}}_{<t} | Y_t, \mathbf{Y}_{<t}). \quad (2)$$

Postprocessing is applied to extract  $\bar{Y}_t$  from the output.

We consider constrained context with context size  $c$  in practice, which shows the number of previous utterances used for translation in addition to the utterance to be translated. More formal definitions of monolingual, bilingual, and constrained context are provided in Appendix C.

## 6 Experiments

### 6.1 Automatic Speech Recognition

In SDT, ASR has to handle bilingual inputs. We used a multilingual ASR model Whisper (Radford et al., 2022). The medium model with 12 encoder and decoder layers was used without finetuning. Further details are provided in Appendix D.1. We evaluated the performance of the SpeechBSD test set. For English the word error rate was 8.3 %, and for Japanese the character error rate was 13.2 %.

### 6.2 Machine Translation

MT model also needs to handle bilingual inputs in SDT. We used mBART (Liu et al., 2020) and finetuned the model with SpeechBSD for MT. The large model with 12 encoder and decoder layers

<sup>2</sup>Because workers could participate in multiple tasks, it is possible that different parts are actually spoken by the same person.

<sup>3</sup><https://www.mturk.com/>

<sup>4</sup><https://crowdsourcing.yahoo.co.jp/>

<sup>5</sup>The utterances are generated token-wise. The notations in equations 1 and 2 are simplified for clarity.

was used. Although the dialogues are regarded as bilingual ones in this study, the predictions were re-composed to the monolingual dialogue form for evaluation because usually performance of MT models is evaluated on a single language pair. SacreBLEU (Post, 2018) was used for calculating BLEU scores. Further details are provided in Appendix D.2.

### 6.2.1 Context Settings

Three settings were considered: translation without context, with monolingual context, and with bilingual context.

**Without Context** Each utterance in a scenario was treated as a separate sentence in this setting. Finetuning was performed separately for each translation direction.

**Monolingual Context** For each utterance in a scenario, monolingual context with context width  $c = 5$  was composed in the way described in section 5. The context utterances and the utterance to translate were concatenated with the end of sentence token  $\langle /s \rangle$ . Finetuning was performed separately for each translation direction.

**Bilingual Context** For each utterance in a scenario, bilingual context with context width  $c = 5$  was composed in the way described in section 5. The context utterances and the utterance to translate were concatenated with the end of sentence token  $\langle /s \rangle$ . As there is no concept of source language or target language in this setting, a single model was finetuned in this setting.

### 6.2.2 Results

Table 1 (upper part) shows the results of the MT experiments. Comparing “Without” with “Monolingual,” more than 0.9 points of improvement were observed using monolingual context. Comparing “Monolingual” with “Bilingual,” the latter performed better, especially in Ja–En.

### 6.2.3 Manual Evaluation

To verify how context can help improve translations, we conducted a manual evaluation focusing on a grammatical phenomenon called zero anaphora, as discussed in Rikters et al. (2019). Similarly to Rikters et al. (2019), we counted the number of sentences with pronouns *I*, *you*, *he*, *she*, *it*, and *they* in English<sup>6</sup> and observed that 63 % of the

<sup>6</sup>We tokenized sentences with the NLTK toolkit (Bird et al., 2009).

	Context	En–Ja	Ja–En
MT	Without	15.9	18.2
	Monolingual	16.8 <sup>†</sup>	19.5 <sup>†</sup>
	Bilingual	<b>17.0<sup>†</sup></b>	<b>20.1<sup>†‡</sup></b>
Cascade ST	Without	15.2	15.4
	Monolingual	15.9 <sup>†</sup>	16.5 <sup>†</sup>
	Bilingual	<b>16.4<sup>†</sup></b>	<b>17.1<sup>†‡</sup></b>

Table 1: BLEU scores of the SpeechBSD test set for MT and Cascade ST experiments. “<sup>†</sup>” and “<sup>†‡</sup>” indicate that the results are significantly better than “without context” and “monolingual context” at  $p < 0.05$ , respectively.

test sentences included them. We sampled 50 of those sentences from the test set. First, we checked if the subjects of the Japanese sentences were zero pronouns by comparing Japanese and English gold references. Then we checked if the zero pronouns were translated into English correctly for the predictions of each Ja–En system.

Out of the 50 sentences, 29 were sentences with zero pronoun subjects. The number of sentences that the missing pronoun was translated correctly was 19, 20, and 24 for without context, monolingual context, and bilingual context settings, respectively. This shows that context can help disambiguate zero pronouns, and using bilingual context can help generate correct pronouns. Examples of the sentences are shown in Appendix E.

### 6.3 Cascade Speech Translation

Cascade ST experiments were performed by using Whisper recognition results as input to the MT models described in section 6.2.

Table 1 (lower part) shows the results. Similarly to MT, BLEU score improved by more than 0.7 points by using monolingual context. Further improvements by more than 0.5 points were observed using bilingual context.

We also performed manual evaluation as in Section 6.2.3. The number of sentences that the missing pronoun was translated correctly was 16, 18, and 22 for without context, monolingual context, and bilingual context settings, respectively. It showed a similar trend to the results of section 6.2.3 with lower translation accuracy. Examples of the sentences are shown in Appendix E.

## 7 Conclusion

We presented a new task, SDT aiming to mediate speakers of different languages. We constructed the SpeechBSD dataset via crowdsourcing. We



performed MT experiments utilizing context and showed its effectiveness. In the future, we plan to perform experiments in end-to-end ST settings and SDT utilizing speaker attributes.

## Limitations

The experiments were performed only on Japanese and English bilingual dialogue collected from a limited number of native speakers. Although the methods proposed in this work can work on any language pair, drawing conclusions for other language pairs should be avoided. The experiments were performed using existing pretrained models, Whisper and mBART, and the method used to pre-train those models would have affected the translation performances in this work. The dialogues in the SpeechBSD dataset are the read speech of pre-composed text dialogues, and further research is required for more realistic settings such as spontaneous dialogues.

## Ethics Statement

Consent was obtained from the crowdsourcing workers when collecting audio, gender, and homeplace. The SpeechBSD dataset is made public under the Creative Commons Attribution-NonCommercial-ShareAlike (CC BY-NC-SA) 4.0 license, which is the same as the license of the BSD corpus, and shall be used only for research purposes. Caution should be exercised when using gender or homeplace information included in the dataset so that the identities of the speakers are not revealed.

## Acknowledgements

This work was supported by JSPS KAKENHI Grant Numbers JP23H03454 and JP23KJ1356.

## References

Luisa Bentivogli, Mauro Cettolo, Marco Gaido, Alina Karakanta, Alberto Martinelli, Matteo Negri, and Marco Turchi. 2021. [Cascade versus direct speech translation: Do the differences still make a difference?](#) In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pages 2873–2887, Online. Association for Computational Linguistics.

Alexandre Berard, Olivier Pietquin, Christophe Servan, and Laurent Besacier. 2016. Listen and Translate:

A Proof of Concept for End-to-End Speech-to-Text Translation. In *NIPS Workshop on End-to-end Learning for Speech and Audio Processing*.

Steven Bird, Edward Loper, and Ewan Klein. 2009. *Natural Language Processing with Python*. O’Reilly Media Inc.

Zhehuai Chen, Yu Zhang, Andrew Rosenberg, Bhuvana Ramabhadran, Pedro J. Moreno, Ankur Bapna, and Heiga Zen. 2022. [MAESTRO: Matched Speech Text Representations through Modality Matching](#). In *Proc. Interspeech 2022*, pages 4093–4097.

Mattia A. Di Gangi, Roldano Cattoni, Luisa Bentivogli, Matteo Negri, and Marco Turchi. 2019. [MuST-C: a Multilingual Speech Translation Corpus](#). In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 2012–2017, Minneapolis, Minnesota. Association for Computational Linguistics.

Long Duong, Antonios Anastasopoulos, David Chiang, Steven Bird, and Trevor Cohn. 2016. An Attentional Model for Speech Translation Without Transcription. In *Proceedings of the 2016 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*.

Hirofumi Inaguma and Tatsuya Kawahara. 2021. [VAD-Free Streaming Hybrid CTC/Attention ASR for Unsegmented Recording](#). In *Proc. Interspeech 2021*, pages 4049–4053.

Taku Kudo and John Richardson. 2018. [SentencePiece: A simple and language independent subword tokenizer and detokenizer for neural text processing](#). In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing: System Demonstrations*, pages 66–71, Brussels, Belgium. Association for Computational Linguistics.

Yunlong Liang, Chulun Zhou, Fandong Meng, Jinan Xu, Yufeng Chen, Jinsong Su, and Jie Zhou. 2021. [Towards making the most of dialogue characteristics for neural chat translation](#). In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pages 67–79, Online and Punta Cana, Dominican Republic. Association for Computational Linguistics.

Siyu Liu, Yuqi Sun, and Longyue Wang. 2021. [Recent advances in dialogue machine translation](#). *Information*, 12(11).

Yinhan Liu, Jiatao Gu, Naman Goyal, Xian Li, Sergey Edunov, Marjan Ghazvininejad, Mike Lewis, and Luke Zettlemoyer. 2020. [Multilingual denoising pre-training for neural machine translation](#). *Transactions of the Association for Computational Linguistics*, 8:726–742.

- Myle Ott, Sergey Edunov, Alexei Baevski, Angela Fan, Sam Gross, Nathan Ng, David Grangier, and Michael Auli. 2019. [fairseq: A fast, extensible toolkit for sequence modeling](#). In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics (Demonstrations)*, pages 48–53, Minneapolis, Minnesota. Association for Computational Linguistics.
- Matt Post. 2018. [A call for clarity in reporting BLEU scores](#). In *Proceedings of the Third Conference on Machine Translation: Research Papers*, pages 186–191, Brussels, Belgium. Association for Computational Linguistics.
- Alec Radford, Jong Wook Kim, Tao Xu, Greg Brockman, Christine McLeavey, and Ilya Sutskever. 2022. [Robust Speech Recognition via Large-Scale Weak Supervision](#).
- Stefan Riezler and John T. Maxwell. 2005. [On some pitfalls in automatic evaluation and significance testing for MT](#). In *Proceedings of the ACL Workshop on Intrinsic and Extrinsic Evaluation Measures for Machine Translation and/or Summarization*, pages 57–64, Ann Arbor, Michigan. Association for Computational Linguistics.
- Matiss Rikters, Ryokan Ri, Tong Li, and Toshiaki Nakazawa. 2019. [Designing the business conversation corpus](#). In *Proceedings of the 6th Workshop on Asian Translation*, pages 54–61, Hong Kong, China. Association for Computational Linguistics.
- Matiss Rikters, Ryokan Ri, Tong Li, and Toshiaki Nakazawa. 2021. [Japanese–english conversation parallel corpus for promoting context-aware machine translation research](#). *Journal of Natural Language Processing*, 28(2):380–403.
- F. W. M. Stentiford and M. G. Steer. 1988. Machine translation of speech. *British Telecom technology journal*.
- Viet Anh Khoa Tran, David Thulke, Yingbo Gao, Christian Herold, and Hermann Ney. 2022. [Does Joint Training Really Help Cascaded Speech Translation?](#)
- Longyue Wang, Zhaopeng Tu, Xing Wang, Li Ding, Liang Ding, and Shuming Shi. 2020. [Tencent AI lab machine translation systems for WMT20 chat translation task](#). In *Proceedings of the Fifth Conference on Machine Translation*, pages 483–491, Online. Association for Computational Linguistics.
- Biao Zhang, Ivan Titov, Barry Haddow, and Rico Sennrich. 2021. [Beyond sentence-level end-to-end speech translation: Context helps](#). In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pages 2566–2578, Online. Association for Computational Linguistics.

## A Crowdsourcing Details

### A.1 Crowdsourcing Instructions Given to the Workers

Figure 2 shows the instructions given to the crowdsourcing workers and the interface used to record audio. We asked the workers to speak clearly and formally and to check that the audio was properly recorded. With the interface, we made sure that the workers agreed that their voices would be released and that the utterances were properly recorded.

### A.2 Crowdsourcing Payment

The crowdsourcing tasks were divided according to the number of utterances to record. The authors performed preliminary crowdsourcing tasks and estimated how long the tasks would take for each case. We paid the workers according to the estimated time and predefined wage per hour determined for each country.

## B Statistics of the SpeechBSD Dataset

Table 2 shows the statistics of the SpeechBSD dataset. Figure 3 shows the homeplace distribution of the speakers of the SpeechBSD dataset. The Japanese one (3(b)) roughly reflects Japan’s demographics (concentrated around Tokyo, Osaka, and Nagoya), whereas the English one (3(a)) is more biased (concentrated too much on California and Virginia). We believe these biases are caused by the differences in the crowdsourcing platforms used.

## C Formal Definition of Context

Here, we formally formulate monolingual, bilingual, and constrained contexts introduced in Section 5.

For simplicity, we consider the case where  $M = 2$  and  $m = n$  (i.e., speaker  $S^i$  speaks in language  $L^i$  ( $i = 1, 2$ )). In addition, we suppose the speakers speak interchangeably, and speaker  $S^1$  starts the conversation.<sup>7</sup> In other words, defining a map  $L : U_t \mapsto L^i$ ,

$$\forall U \in \{U_t \mid t \equiv i \pmod{2}\}, \quad L(U) = L^i.$$

The monolingual context is composed of previous utterances in a single language. In other words,

<sup>7</sup>In the experiments, consecutive utterances by the same speaker are treated as separate utterances. If there are more than three speakers, we number speakers in the order of appearance and regard speakers with the same parity speak in the same language.

monolingual context text of utterance  $U_t$  in language  $L^i$  is

$$\mathbf{Y}_{<t}^i = \{Y_\tau^i \mid \tau < t\}.$$

For example in Figure 1, when translating the third utterance  $U_3$  from Japanese to English, the monolingual context of the source side is “彼は良い考えだと言っていました。あなたはどう思いますか?”, and that of the target side is “He said it’s a good idea. What do you think?” Using this formulation, we can formally define the training objective of Equation 1. During inference, for the source language of the current utterance, ASR transcripts are used, and for the target language of the current utterance, the translations of ASR transcripts are used to compose context. During training, the corresponding gold text is used.

The Bilingual context is composed of transcripts of the two languages. ASR transcripts are used during inference, and gold transcripts are used for training. The bilingual context of utterance  $U_t$  is  $\mathbf{Y}_{<t} = \tilde{\mathbf{Y}}_{<t}^1 \cup \tilde{\mathbf{Y}}_{<t}^2$ , where

$$\tilde{\mathbf{Y}}_{<t}^i = \{Y_\tau^i \mid \tau < t \wedge \tau \equiv i \pmod{2}\}.$$

For example in Figure 1, when translating the third utterance  $U_3$  from Japanese to English, the bilingual context of the source side is “彼は良い考えだと言っていました。What do you think about it?”, and that of the target side is “He said it’s a good idea. あなたはどう思いますか?”

For bilingual context experiments, the MT system has to be able to handle two translation directions. Let the translation of  $\mathbf{Y}_{<t}$  be  $\overline{\mathbf{Y}}_{<t} = \overline{\tilde{\mathbf{Y}}_{<t}^1} \cup \overline{\tilde{\mathbf{Y}}_{<t}^2}$ , where

$$\overline{\tilde{\mathbf{Y}}_{<t}^i} = \{Y_\tau^j \mid \tau < t \wedge \tau \equiv i \pmod{2}\}, \\ (i, j) = (1, 2), (2, 1).$$

$\overline{\mathbf{Y}}_t$  is  $Y_t^2$  when  $L(U_t) = L^1$  and  $Y_t^1$  when  $L(U_t) = L^2$ . By setting  $\mathbf{Y}_{<t}$  as source side context and target side context as  $\overline{\mathbf{Y}}_{<t}$ , we can formally define the training objective of Equation 2.

In practice, we consider context width  $c$  for context  $U_{<t} = \{U_\tau \mid \tau < t\}$  because the maximum length the MT models can handle is limited. The constrained context of utterance  $U_t$  with context width  $c$  is

$$U_{<t} = \{U_\tau \mid \tau = t - 1, \dots, t - c \wedge \tau > 0\}.$$

### Audio Data Collection

Read aloud the displayed sentences, record and check them.

Notes

- Record in a **quiet** environment.
- Speak **clearly and formally**. If you make a mistake, you can record again.
- Use a **microphone** if possible.
- Make sure that the audio is **properly recorded** before submission.
- Please note that the recorded voice will be released as a dataset for academic purposes in the future.

I agree that my voice will be made public as a part of a dataset.

Sentences to record	Record / Stop	Check the audio
Hi I'm Ben Sherman.		
I'm the manager and Sam is the sales representative in our department.		
Hey Sam, how about we start with Andy telling us a bit more about himself.		
So why did you decide to apply for this position at our company?		
Alright, let's go to the next question.		

I checked all the utterances are properly recorded.

[Submit all recordings](#)

(a) English

### クラウドソーシング 音声録音ページ

表示された文を音読及び録音し、確認してください。

注意点

- 録音は静かな環境で行ってください。
- はっきりと丁寧に発音するようにしてください。読み間違えた場合は録り直しが可能です。
- 可能であればマイクなどの録音機器を使用してください。
- 録音した音声はきちんと録音されているかどうか確認してから提出してください。
- 録音された音声は学術研究のためにデータセットとして公開される予定ですので、ご了承ください。

録音した音声データセットの一部として公開されることに同意します。

録音する文	録音/停止	録音した音声の確認
こんにちは、ベンシャーマンです。		
私はマネージャーで、サムはうちの部署の営業担当です。		
なあ、サム、アンディさんにもっと御自身のことについて話してもらうとこから始めたらどうか。		
では、なぜ弊社のこのポジションに応募を決められたんでしょうか？		
そうですね、では、次の質問へ進みましょう。		

全ての音声正常に録音されていることを確認しました。

[全ての録音を提出する](#)

(b) Japanese

Figure 2: Crowdsourcing interface used to record audio. The upper part shows the instructions given to the workers.

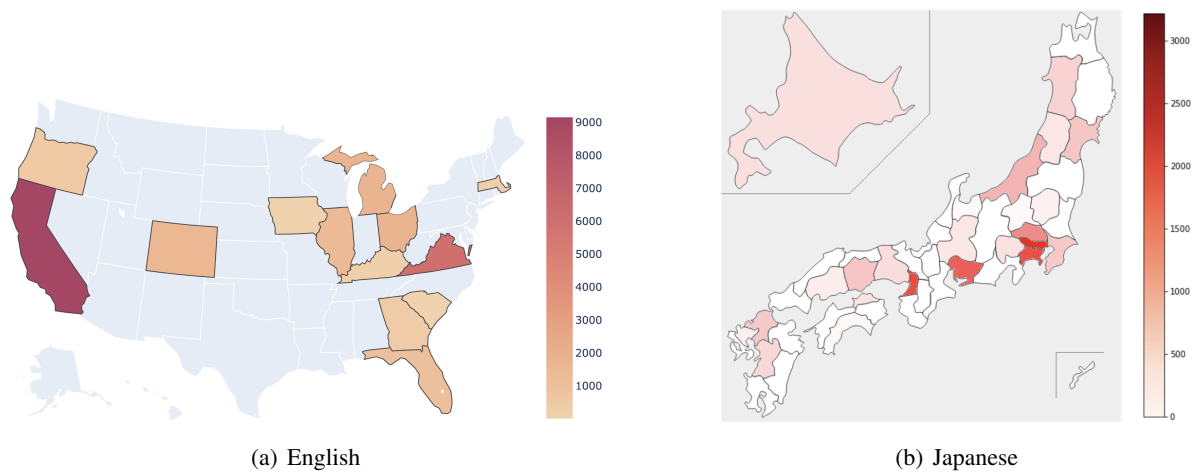


Figure 3: Homeplace distribution of the speakers of the SpeechBSD dataset by the number of utterances.



	Train	Dev.	Test
# of scenarios	670	69	69
# of sentences	20,000	2,051	2,120
English speech (h)	20.1	2.1	2.1
Japanese speech (h)	25.3	2.7	2.7
English gender (M / F %)	47.2 / 52.8	50.1 / 49.9	44.4 / 55.6
Japanese gender (M / F %)	68.0 / 32.0	62.3 / 37.7	69.0 / 31.0

Table 2: Statistics of the SpeechBSD dataset. The number of sentences is the same as the number of utterances in this dataset as it originally was in the BSD corpus.

## D Experimental Settings

### D.1 ASR

Whisper is a Transformer-based model that uses 80-channel log-Mel spectrograms converted from audio sampled with 16,000 Hz as input. As it is trained with 680,000 hours of data in various domains the model is robust enough to be able to work without any finetuning. We used the byte-level BPE vocabulary (size 50,257) of the pretrained model. We assumed the language of the utterances was given beforehand and fed the language tag to the model as a prefix token. We evaluated the development set of the SpeechBSD dataset using the base, small, medium, and large models with either greedy decoding or beam search decoding with beam size 5. We observed that the medium model with greedy decoding performed the best for both English and Japanese, which are the settings used for further experiments.

### D.2 MT

We used mBART trained with 25 languages for the experiments. BPE vocabulary of size 25,001 was used. As a preprocessing step, BPE was applied to all utterances with the sentencepiece (Kudo and Richardson, 2018) toolkit. Fairseq (Ott et al., 2019) was used for training and inference. The same hyperparameters as in Liu et al. (2020) were used, except that the training epochs were determined according to early stopping with patience 10 on validation loss. We did not use different random seeds for the experiments because Liu et al. (2020) reported that the finetuning process was stable with different seeds. When evaluating the model, the averaged weights of the last 10 checkpoints were used. The SacreBLEU signatures were `nrefs:1|case:mixed|eff:no|tok:ja-mecab-0.996-IPA|smooth:exp|version:2.0.0` for En-Ja and `nrefs:1|case:mixed|eff:no|tok:13a|`

`smooth:exp|version:2.0.0` for Ja-En. We conducted significance tests with paired approximate randomization (Riezler and Maxwell, 2005) with 10,000 approximate randomization trials and a  $p$ -value threshold of 0.05 to compare the BLEU scores of “without context” with the others, and “monolingual context” with “bilingual context.”

For bilingual context MT experiments, in order to match the finetuning style of mBART, language tags like `ja_XX` or `en_XX` have to be appended at the last of each translation unit. However, in bilingual context settings, both the source and the target side contain both languages, which does not comply with the finetuning style described in the original mBART paper (Liu et al., 2020). We conducted two kinds of experiments, appending `ja_XX` to the input and `en_XX` to the output and the other way around. The statistical significance test showed that they were not significantly different. We report the results of the systems where the language pair of the utterance to be translated matches the language pair specified by the appended language tags.

As to the context size  $c$ , we changed it from 1 to 8 in the bilingual context setting and evaluated the models with BLEU score on the validation set. The results are shown in Figure 4. In the bilingual context setting, 5 was the best for both En-Ja and Ja-En. For the monolingual context setting, 5 and 6 were the best for En-Ja and 3 for Ja-En. The difference between setting 3 and 5 as context width did not show a statistically significant difference in the BLEU scores for Ja-En. Therefore, for a consistent comparison, we reported the results on the test set with  $c = 5$  in Table 1.

We used 4 Tesla V100 or Titan RTX GPUs for the experiments. The total computation hours, including hyperparameter searching, were 278 hours.

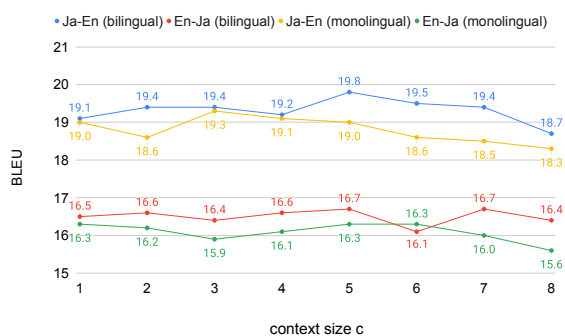


Figure 4: BLEU score on the development set when changing the context size  $c$ .

## E Example Sentences from Manual Evaluation

Table 3 shows examples from manual evaluation described in Section 6.2.3. In the first example, it is observed that the zero pronoun (*She*) is predicted correctly when monolingual or bilingual context is used in both MT and cascade ST experiments. In the second example, the zero pronoun (*They*) could not be correctly predicted by any system.

Context		
Ja reference	-	もう諦めて、仕事なら仕方ないわねって。
En reference	-	She's given up and just says it can't be helped if it's work.
MT	Without	I just gave up and said I can't do it if it's my job.
	Monolingual	She just gave up and said it's okay if it's the job.
	Bilingual	She's giving up on it and says if it's work then it's all right.
Cascade ST	Without	If I just gave up and gave up my job, I can't do anything about it.
	Monolingual	She said if I just gave up and gave up on the job then it should be fine.
	Bilingual	She said if I just give up and give up on the job then it's all right.

(a) An example where the “monolingual” and “bilingual” context predictions were better than the “without context” one. In this scenario, Patrick complains to Gary that he does not want to go to his company’s drinking party. Gary asks what Patrick’s wife thinks about it, and this is Patrick’s response. The pronoun *She* is omitted in the Japanese utterance. Word-by-word translation of the Japanese utterance with omitted words is: “(彼女は)-she / もう-already / 諦めて-give up / それが-it's / 仕事-work / なら-if / (それは)-it / 仕方ない-can't be helped / わね-I think / って(言ってる)-says .”

Context		
Ja reference	-	いつ在庫が入るか、でしょう?
En reference	-	They all want to know when it will be restocked, don't they?
MT	Without	When will the inventory start?
	Monolingual	So when will the inventory be available?
	Bilingual	I wonder when it will be in stock?
Cascade ST	Without	When will the inventory arrive?
	Monolingual	I wonder when it will be in stock.
	Bilingual	I wonder when it will be in stock.

(b) An example where all systems failed to predict the correct pronoun. In this scenario, Mr. Ogawa and Ms. Pace are talking about their company’s stock of a product. The previous utterances by Mr. Ogawa are, “We have 28 backorders for this product. I have been receiving many inquiries from the customers lately.”

This is the subsequent Ms. Pace’s response. The pronoun *They* is omitted in the Japanese utterance. Word-by-word translation of the Japanese utterance with omitted words is: “(彼らは)-they / いつ-when / 在庫-stock / が入るか-becomes available / (を聞くの)-ask / でしょう-don't they.” The translation is difficult because the word corresponding to “ask” is also omitted.

Table 3: Examples from manual evaluation of Ja-En translations focusing on zero pronouns.

## ACL 2023 Responsible NLP Checklist

---

### A For every submission:

- A1. Did you describe the limitations of your work?  
*The Limitations section is after Section 7.*
- A2. Did you discuss any potential risks of your work?  
*The Ethics Statement before References.*
- A3. Do the abstract and introduction summarize the paper’s main claims?  
*Abstract and Section 1.*
- A4. Have you used AI writing assistants when working on this paper?  
*Left blank.*

### B Did you use or create scientific artifacts?

*Section 4.*

- B1. Did you cite the creators of artifacts you used?  
*Section 4.*
- B2. Did you discuss the license or terms for use and / or distribution of any artifacts?  
*The Ethics Statement before References.*
- B3. Did you discuss if your use of existing artifact(s) was consistent with their intended use, provided that it was specified? For the artifacts you create, do you specify intended use and whether that is compatible with the original access conditions (in particular, derivatives of data accessed for research purposes should not be used outside of research contexts)?  
*Section 4.*
- B4. Did you discuss the steps taken to check whether the data that was collected / used contains any information that names or uniquely identifies individual people or offensive content, and the steps taken to protect / anonymize it?  
*The Ethics Statement before References.*
- B5. Did you provide documentation of the artifacts, e.g., coverage of domains, languages, and linguistic phenomena, demographic groups represented, etc.?  
*Section 4.*
- B6. Did you report relevant statistics like the number of examples, details of train / test / dev splits, etc. for the data that you used / created? Even for commonly-used benchmark datasets, include the number of examples in train / validation / test splits, as these provide necessary context for a reader to understand experimental results. For example, small differences in accuracy on large test sets may be significant, while on small test sets they may not be.  
*Section 4 and Appendix B.*

### C Did you run computational experiments?

*Section 6.*

- C1. Did you report the number of parameters in the models used, the total computational budget (e.g., GPU hours), and computing infrastructure used?  
*Appendix D.2.*

---

*The Responsible NLP Checklist used at ACL 2023 is adopted from NAACL 2022, with the addition of a question on AI writing assistance.*

- C2. Did you discuss the experimental setup, including hyperparameter search and best-found hyperparameter values?  
*Section 6 and Appendix D.*
- C3. Did you report descriptive statistics about your results (e.g., error bars around results, summary statistics from sets of experiments), and is it transparent whether you are reporting the max, mean, etc. or just a single run?  
*Appendix D.*
- C4. If you used existing packages (e.g., for preprocessing, for normalization, or for evaluation), did you report the implementation, model, and parameter settings used (e.g., NLTK, Spacy, ROUGE, etc.)?  
*Section 6 and Appendix D.*
- D**  **Did you use human annotators (e.g., crowdworkers) or research with human participants?**  
*Section 4.*
- D1. Did you report the full text of instructions given to participants, including e.g., screenshots, disclaimers of any risks to participants or annotators, etc.?  
*Appendix A.*
- D2. Did you report information about how you recruited (e.g., crowdsourcing platform, students) and paid participants, and discuss if such payment is adequate given the participants' demographic (e.g., country of residence)?  
*Section 4 and Appendix A.*
- D3. Did you discuss whether and how consent was obtained from people whose data you're using/curating? For example, if you collected data via crowdsourcing, did your instructions to crowdworkers explain how the data would be used?  
*Appendix A.*
- D4. Was the data collection protocol approved (or determined exempt) by an ethics review board?  
*Left blank.*
- D5. Did you report the basic demographic and geographic characteristics of the annotator population that is the source of the data?  
*Appendix B.*