

Filtering Noisy Dialogue Corpora by Connectivity and Content Relatedness

Reina Akama^{1,2} Sho Yokoi^{1,2} Jun Suzuki^{1,2} Kentaro Inui^{1,2}

¹Tohoku University ²RIKEN

{reina.a, yokoi, jun.suzuki, inui}@ecei.tohoku.ac.jp

Abstract

Large-scale dialogue datasets have recently become available for training neural dialogue agents. However, these datasets have been reported to contain a non-negligible number of unacceptable utterance pairs. In this paper, we propose a method for scoring the quality of utterance pairs in terms of their connectivity and relatedness. The proposed scoring method is designed based on findings widely shared in the dialogue and linguistics research communities. We demonstrate that it has a relatively good correlation with the human judgment of dialogue quality. Furthermore, the method is applied to filter out potentially unacceptable utterance pairs from a large-scale *noisy* dialogue corpus to ensure its quality. We experimentally confirm that training data filtered by the proposed method improves the quality of neural dialogue agents in response generation.¹

1 Introduction

Some million-scale datasets such as movie scripts and social media posts have become available in recent years for building neural dialogue agents (Lison and Tiedemann, 2016; Henderson et al., 2019). Such large-scale datasets can be expected to improve the performance of dialogue response generation models based on deep neural networks (DNNs) since the combination of DNNs and large-scale training datasets has led to considerable performance improvement in many sentence generation tasks (Koehn and Knowles, 2017; Sennrich and Zhang, 2019; Adiwardana et al., 2020).

In contrast to the quantity of the data, the quality of the data has often been problematic. For example, OpenSubtitles (Lison and Tiedemann, 2016; Lison et al., 2018), the most widely used large-scale English dialogue corpus, was constructed by

¹The code is publicly available at <https://github.com/jqk09a/CoRe-dialogue-filtering>.




Figure 1: *Is the sequence of the two utterances acceptable as a dialogue?* Response acceptability scores are given by humans on the English OpenSubtitles corpus.

collecting two consecutive lines of movie subtitles under the simplified assumption that one line of a movie subtitle is one utterance and the next line is the next utterance follow it. Inevitably, this corpus includes unacceptable utterance pairs from the viewpoint of a conversational sequence, e.g., caused by scene switching or flashback. Several previous studies have identified such flaws and reported that the corpus is *noisy* (Vinyals and Le, 2015; Li et al., 2016; Baheti et al., 2018), where *noisy* refers to unacceptable utterance pairs in this context. Figure 1 shows the result of our experimental investigation regarding the acceptability rate of the utterance pairs in the OpenSubtitles corpus.² It can be noticed from the figure that only half of the utterance pairs can be considered *acceptable* (i.e., were rated with score 5: Strongly agree or 4: Agree), and over 25% of utterance pairs are clearly *unacceptable* (i.e., were rated with score 1: Strongly disagree or 2: Disagree) from the human perspective.³

With this situation, a straightforward research question arises, namely, *Can we further improve the performance of neural response generation models by ablating unacceptable utterance pairs*

²See Appendix A for detailed experimental settings.

³See Table 1 for samples of unacceptable/acceptable utterance pairs annotated by humans.

from training data? To the best of our knowledge, no previous study has explicitly focused on this question. Thus, the goal of this paper is to provide an answer to this question. Furthermore, it is not clear whether and how one can effectively discover unacceptable utterance pairs within large-scale training datasets. This study explores a way of constructing a scoring method for filtering *noisy* data filtering to improve the performance of response generation models.

To achieve the set goals, we started with a review of previous arguments about the criteria for identifying appropriate utterances in dialogues and designed our scoring function that is consistent with reflects as much of the community’s consensus as possible. In particular, the proposed scoring method estimates the quality of utterance pairs based on the following two aspects: (i) the **connectivity** between source and target utterances and (ii) their **content relatedness** (Section 4).

The contributions of this study are the following:

- We propose a scoring method for estimating the quality of utterance pairs in an unsupervised manner (Section 5);
- We reveal that our scoring method effectively detects unacceptable utterance pairs, and thus, be appropriate for noisy data filtering (Section 6);
- We empirically prove that our proposed data filtering method improves the performance of neural response generation models (Section 7); and
- We confirm that our noisy data filtering approach is effective across different languages and dataset sizes (Section 8).

2 Task Definition: Noisy Data Filtering

Let x be an **utterance** and y be a **response** to x . Then, an **utterance pair** can be denoted as we refer to (x, y) . Let \mathcal{D} be a dataset that comprising a set of utterance pairs, $\mathcal{D} = \{(x, y)\}$. Then, the task can be formulated as ablating unacceptable utterance pairs from \mathcal{D} to obtain a less noisy subset $\mathcal{D}' \subseteq \mathcal{D}$, hereinafter referred to as filtered dataset. \mathcal{D}' can then be used to train response generation models. This paper refers to this process as **noisy data filtering**, where *noisy* means unacceptable utterance pairs in this context. Furthermore, we establish a function $S: \mathcal{D} \rightarrow \mathbb{R}$ is used to score the degree of *acceptability* of each utterance pair $(x, y) \in \mathcal{D}$.

3 Background

Response generation using noisy data. The following two approaches are widely used to address the problem of dialogue response generation noisy dialogue corpora. According to the *model approach*, models are trained while handling noise at the same time. For example, Shang et al. (2018) proposed a method with a calibration framework and demonstrated its effectiveness on a Chinese corpus. According to the *data approach*, training data are pre-processed with the aim of improving their quality before training models. In this study, we take the data approach in light of the success of noisy parallel corpus filtering in machine translation (MT). Additionally, it has become a reasonable strategy to reduce the size of training data since enormous dialogue data has been available. Csáky et al. (2019)’s method is most relevant to our study in that it cleanses dialogue corpora. However, the main goal of their method is to eliminate generic, or boring, responses, whereas the goal of the method proposed here is to eliminate unacceptable utterance pairs. This difference in goals leads to the essential difference in filtering strategies.

Effectiveness of filtering noisy data in neural machine translation. Researchers in the field of neural machine translation (NMT) have recognized that collecting high-quality training data to be equally or even more important than exploring sophisticated model architectures (Koehn et al., 2018; Junczys-Dowmunt, 2018; Morishita et al., 2018). Techniques used in neural response generation and NMT are nearly identical; e.g., sequence-to-sequence models (Sutskever et al., 2014) and Transformers (Vaswani et al., 2017) are often used as base model architectures. We hypothesize that high-quality filtered dialogue data can also improve the performance of dialogue response generators. However, the straightforward application of methods proposed for filtering noisy data in NMT may not work well due to the different nature of NMT and neural response generation tasks. In particular, MT data have one-to-one (ignoring paraphrases) correspondence in source and target sentences, whereas dialogues have many-to-many mappings (Zhao and Eskenazi, 2017). The experiments presented in this paper provide an answer to whether NMT filtering methods can perform well in dialogue response generation.

4 Requirements to Utterance Pairs

In this section, we investigate the requirements that should be satisfied by an acceptable utterance pair.

4.1 Criteria for Manual Evaluation

The instructions for manual evaluation provided by the dialogue community explain the key factors for distinguishing acceptable and unacceptable utterance pairs.

In many previous studies, human raters were asked to evaluate the **connectivity** of utterance pairs. For instance, Shang et al. (2015) asked whether a response could be considered as *an appropriate and natural response to the post*. Xing et al. (2017) asked whether *the response can be used as a reply*. Pei and Li (2018) asked whether *the answer is natural* for the question. Other studies have also evaluated the same or similar aspects by using keywords related to the connectivity, such as *semantically appropriate for* (Akama et al., 2017) or *coherent with* (Shen et al., 2017), and *coherence* (Lowe et al., 2017b).

Another frequently used metric is **content relatedness**. For instance, Galley et al. (2015) asked human evaluators to evaluate *each response in terms of their relevance to a given utterance*. Li et al. (2016) asked for the preference of responses *that were more specific to certain utterances*. Ritter et al. (2011) suggested that *an appropriate response should be on the same topic as the utterances*. Several other studies have also focused on evaluating the *relevance* between an utterance and its response (Xu et al., 2018b; Pei and Li, 2018; Lowe et al., 2017b).

In summary, the most widely used criteria can be categorized into connectivity and content relatedness of utterance pairs. In fact, these two aspects are considered in the field of sociolinguistics as crucial features of conversation (Sacks, 1989; Sidnell, 2010).

4.2 Observation

Furthermore, we investigated how the two aforementioned aspects can be observed in actual utterance pairs. For this investigation, we use the utterance pairs scored by human raters that were used in our preliminary experiments shown in Figure 1. Some examples are shown in Table 1.

We observe that typical phrase pair patterns can often be found in utterance pairs with high scores. For example, the pair (*where is*, *at*) in Table 1

is one of the typical phrase pair patterns that asks a place and provides an answer to it. Other typical examples include (*why*, *because*) and (*what do you want*, *I want*). In discourse linguistics, such phrase pair patterns are known as the concept of *cohesive devices*. Hereafter, we refer to such a typical phrase pair pattern as **key phrase pair**.

Moreover, in high scored utterance pairs, both an utterance and response are on the **same topic**. For example, in the third example listed in Table 1, both the utterance and response mention [money].

5 Proposed Method

As per the discussion in the previous section, each acceptable utterance pair should satisfy the following criteria:

- **connectivity** — existence of key phrase pairs
- **content relatedness** — topic commonality

This section presents the proposed scoring functions to assess the degree of satisfying the above two criteria in an unsupervised manner.⁴

5.1 Connectivity

Let f and e represent phrases obtained from x and y , respectively. Let $\phi(x, y)$ be a function that returns a set of all possible phrase (n -gram) pairs obtained from the utterance pair (x, y) . We can define a finite set of all possible phrase pairs obtained from the entire dialogue data \mathcal{D} as $\overline{\mathcal{P}}_{\mathcal{D}} = \bigcup_{(x,y) \in \mathcal{D}} \phi(x, y)$. Then, let \mathcal{P} represent a set of key phrase pairs (defined in Section 4.2). We assume that \mathcal{P} is a subset of $\overline{\mathcal{P}}_{\mathcal{D}}$, i.e., $\mathcal{P} \subseteq \overline{\mathcal{P}}_{\mathcal{D}}$.

To obtain \mathcal{P} , we take advantage of a phrase table extraction technique developed in statistical machine translation, e.g., Moses (Koehn et al., 2007). In this task, we require only some phrase pairs that can contribute to the connectivity of an utterance pair (as mentioned in Section 4.2), unlike the translation task where the whole sentence must correspond in mutual. Accordingly, in our experiments, we set the null alignment ratio (i.e., probability of no alignment) to 0.5 and extend the phrase extraction algorithm to include only the explicitly corresponding range as phrases in our table.

Then, we define the scoring function $S_{\mathcal{C}}$ to esti-

⁴The reason for focusing on an unsupervised approach the lack of data that can provide good supervision for utterance pair evaluation.

| Utterance | Response | Human |
|---|---|-------|
| 1: It'll be like you never left. [??] | I painted a white line on the street way over there. [painting] | 1.4 |
| 2: You're gonna get us assimilated. [??] | Switch to a garlic shampoo. [??] | 1.8 |
| 3: I probably asked for too much money. [money] | Money's always a problem, isn't it? [money] | 4.2 |
| 4: I wonder who I should call back. [phone] | They're saying they want to call one of you back. [phone] | 4.4 |
| 5: Okay, so where's the rest? [??] | Electronically scanned and archived at headquarters but you'll have to speak with them about that. [work] | 4.4 |

Table 1: Samples of pairs judged as unacceptable/acceptable in preliminary experiments. Human denotes the average score of five human evaluators on a scale of 1-5. Phrases considered to contribute to connectivity are highlighted. Estimated $_{\text{topic}}$ of utterance is written in the end of each utterance.

mate connectivity as:

$$S_{\mathbf{C}}(x, y) := \sum_{(f,e) \in \phi(x,y) \cap \mathcal{P}} \max(\text{nPMI}(f, e), 0) \cdot \frac{|f|}{|x|} \cdot \frac{|e|}{|y|}, \quad (1)$$

where $|\cdot|$ denotes the number of words in the phrase or utterance. To calculate the co-occurrence, we use the normalized pointwise mutual information (nPMI) (Bouma, 2009), which normalizes the value so that low-frequency phrases do not take an extremely large value. Note that we ignore the negative nPMI scores by the $\max(\cdot, 0)$ operation because we aim only to consider the positive effect of connectivity. The intuition behind Equation 1 is as follows:

- If a phrase pair (f, e) has a high co-occurrence, the association strength of (x, y) including (f, e) might also be high.
- If a phrase f or e occupies almost the entire sentence x or y , (f, e) is a strong indicator of the association of (x, y) .

5.2 Content Relatedness

Let $\mathbf{v}(x)$ and $\mathbf{v}(y)$ be sentence vector of x and y , respectively. We compute topic commonality of x and y , that is, content relatedness as follows:

$$S_{\mathbf{R}}(x, y) := \max(\cos(\mathbf{v}(x), \mathbf{v}(y)), 0). \quad (2)$$

Cosine similarity between certain kinds of sentence vectors is known to be a good proxy of the topical relatedness of two sentences (Conneau et al., 2017; Subramanian et al., 2018; Xu et al., 2018a). For the same reasons as Equation 1, we ignore the negative cos scores by the $\max(\cdot, 0)$ operation.

5.3 Summary

Eventually, combining the above two scoring measures, we propose the following function:

$$S_{\mathbf{C}+\mathbf{R}}(x, y) := \alpha S_{\mathbf{C}}(x, y) + \beta S_{\mathbf{R}}(x, y), \quad (3)$$

where $\alpha, \beta \in \mathbb{R}_{\geq 0}$ are hyperparameters that weigh the two viewpoints. For our experiments, we fix α and β as follows:

$$\alpha = \frac{1}{\frac{1}{|\mathcal{D}|} \sum_{(x,y) \in \mathcal{D}} S_{\mathbf{C}}(x, y)}, \quad \beta = \frac{1}{\frac{1}{|\mathcal{D}|} \sum_{(x,y) \in \mathcal{D}} S_{\mathbf{R}}(x, y)}. \quad (4)$$

6 Experiments: Data Scoring

In this section, we describe our experiments that validate the effectiveness of the proposed scoring method.

6.1 Experimental Setup

Dataset. We conducted our experiments on a noisy English dialogue corpus from OpenSubtitles (Lison et al., 2018) containing roughly 441M lines. As explained in Section 1, this corpus includes many unacceptable utterance pairs (Section 1). We first applied several rule-based filtering as rudimentary preprocesses, which are typically used in the related literature. Then, we obtained 79,445,453 utterance pairs as our training data, which excludes our test and validation data.⁵

Proposed method: detailed setup. To compute the connectivity $S_{\mathbf{C}}$, we obtained a phrase table on our training data by using Moses (Koehn et al., 2007) with fastAlign (Dyer et al., 2013). We then removed phrase pairs with a low co-occurrence frequency (here, less than 200 times) or composed

⁵See Appendix B for details on our data such as the preparation procedure and statistics.




Figure 2: Distributions between human scores and automatically computed scores by each method (English).

of the same phrases from the table. As a result, the phrase table included 68,891 phrase pairs, which were used as the key phrase set \mathcal{P} as described in Section 5.1.

To compute the content relatedness $S_{\mathbf{R}}$, we created a sentence vector from pre-trained fast-Text word embeddings (Bojanowski et al., 2017; Mikolov et al., 2018) following Arora et al. (2017)’s method, i.e., using SIF weighting and common component removal. Their method is reported to be useful for computing the relatedness of two given sentences and used in many studies (Marelli et al., 2014b,a; Conneau et al., 2017; Subramanian et al., 2018; Baheti et al., 2018). We learned common components using 30K sentences randomly selected from the training costs appropriately. We then removed the first common component for all sentence vectors.

Baselines. For comparison, we prepared the following:

- **Csáky et al. (2019):** Entropy-based filtering to remove generic utterances from the training data for promoting less-boring response generation. SRC/TRG indicates that using the entropy of source/target utterances.
- **Junczys-Dowmunt (2018):** Filtering for NMT based on the dual conditional cross-entropy computed by a neural encoder-decoder model. It achieved the best performance on the Parallel Corpus Filtering Task at WMT 2018.⁶

Human evaluation. To validate the ability of the proposed method to estimate the quality of utterance pairs, we measured the correlation between its scores and those assigned by humans through crowdsourcing. We used Amazon Mechanical Turk.⁷ We randomly extracted 200⁸ scored

| Scoring method | Spearman’s ρ | p -value |
|--|-------------------|--|
| Csáky et al. (2019) SRC | -0.1173 | 9.8×10^{-2} |
| Csáky et al. (2019) TRG | 0.0462 | 5.2×10^{-1} |
| Junczys-Dowmunt (2018) | 0.2973 | 1.9×10^{-5} |
| Ours $S_{\mathbf{C}+\mathbf{R}}$ | 0.3751 | 4.4×10^{-8} |
| Ours $S_{\mathbf{C}}$ (ablation study) | 0.2044 | 3.7×10^{-3} |
| Ours $S_{\mathbf{R}}$ (ablation study) | 0.3007 | 1.5×10^{-5} |

Table 2: Correlation coefficient between human scores and automatically computed scores (English).

utterance pairs and asked native English-speaking crowdworkers to answer the following question for each pair: *Is the sequence of the two utterances acceptable as a dialogue?* Workers were instructed to provide an answer on a five-point Likert scale (from 5: Strongly agree to 1: Strongly disagree) (Likert, 1932). Unqualified workers were filtered out using attention checks. Eventually, we used the average of the scores provided by five workers as the human score for each pair.

6.2 Results and Analysis

Table 2 shows the correlation between human scores and those automatically computed by each method. Among the methods, $S_{\mathbf{C}+\mathbf{R}}$ achieved the highest correlation with human scores. Additionally, we also evaluated $S_{\mathbf{C}}$ and $S_{\mathbf{R}}$ as an ablation study of $S_{\mathbf{C}+\mathbf{R}}$. We found that both scores were less correlated than $S_{\mathbf{C}+\mathbf{R}}$. This result supports the hypothesis that both aspects, namely, connectivity and content relatedness, should be considered when evaluating the quality of utterance pairs.

Figure 2 shows the distribution of automatically computed scores corresponding to human scores.⁹ As shown in (c), $S_{\mathbf{C}+\mathbf{R}}$ rarely overestimates utterance pairs with low human scores but underestimates those with high human scores. The baseline methods presented in (a) and (b) do not show such behavior. This behavior unique to $S_{\mathbf{C}+\mathbf{R}}$ is safe for

⁶<http://www.statmt.org/wmt18/>

⁷<https://www.mturk.com/>

⁸Same size as Sedoc et al. (2019) and Cho and May (2020).

⁹See Appendix E.1 for the distributions of other methods.

| Utterance | Response | S_C | S_R | S_{C+R} | Human |
|---|---------------------------------------|-------|-------|--------------|-------|
| 1: What is the anarchy facing the jail of the sick passion? | Gosh, it's really cold! | 0.32 | 0.00 | 0.32 | 1.4 |
| 2: Pushers won't let the junkie go free. | Across 110th Street. | 0.00 | 0.42 | 0.42 | 2.4 |
| 3: It started when I was 17. | They'd make a cash drop, | 0.63 | 0.00 | 0.63 | 2.0 |
| 4: A big nail should be put in your head | Who are they | 0.74 | 0.00 | 0.74 | 1.2 |
| 5: He told me so. | Oh, he did, huh? | 2.21 | 0.00 | 2.21 | 4.8 |
| 6: There's a laundry. | Have your clothes dry-cleaned, okay? | 0.81 | 2.89 | 3.70 | 4.4 |
| 7: Then if I win, what are you going to do? | When you win? | 1.04 | 7.01 | 8.05 | 4.2 |
| 8: But what do you want me to do? | We want you to kick her off the team. | 10.20 | 1.53 | 11.72 | 5.0 |

Table 3: Samples of utterance pairs scored with our method and human judgements (English). The scores of S_C and S_R were normalized by α, β .

the noisy data filtering task since it can successfully detect lower-quality pairs with high precision. On the other hand, improperly underestimating some acceptable pairs (i.e., low recall) is one downside of S_{C+R} , and we discuss its influences in Section 6.3. We emphasize that S_{C+R} has a desirable property for noisy data filtering in today's situation where a sufficiently large corpus is available; it allows us to obtain a sufficient amount of clean data even if discarding a certain portion of potentially clean data. Interesting future work is to investigate how to improve our methods not to underestimate acceptable pairs while maintaining high precision. It is nearly equivalent to develop an unsupervised approach of dialogue evaluation methods, and thus, this direction is a challenging and essential attempt.

Table 3 shows several examples of utterance pairs well-scored by S_C , S_R , and S_{C+R} . Note that the score ranges differ; e.g., human scores are in $[1, 5]$, while S_R is in the range $[0, 1]$.¹⁰ Thus, we discuss relative score values; the comparison of absolute score values across the different methods would be meaningless. These examples demonstrate that the complementary contributions of both S_C and S_R allow S_{C+R} to provide quality estimations close to human judgments.

6.3 Discussion on Low Recall Property

What types of pairs cause low recall? Since the proposed method prefers precision over recall, it tends to discard a certain number of acceptable utterance pairs during filtering. To investigate the characteristics of such discarded (yet acceptable) pairs, we analyzed 27 pairs.¹¹ These pairs were selected from those that obtained a human score

¹⁰See Appendix E.2 for score distributions on training data.

¹¹Some examples are listed in Appendix E.3

| Scored data | len | distinct-1 | distinct-2 |
|---------------------|------|------------|------------|
| Top 50% (remained) | 9.02 | 0.028 | 0.472 |
| Worst 50% (removed) | 9.00 | 0.030 | 0.470 |

Table 4: Comparison of the top and the worst utterance pairs' responses in the training data scored by our method (English).

of 4.0 or above (77 pairs) and were among the worst 50% as scored by S_{C+R} (100 pairs). Consequently, we found two potential issues. One is that human annotators may sometimes easily find the connectivity or the content relatedness for the utterance pairs with the low S_{C+R} scores. This observation indicates that S_C and S_R are still not perfect for scoring functions, and there remains room for improvement. The possible drawbacks we have already noticed in S_C and S_R are that S_C sometimes fails to capture the connectivity because of the limited coverage by a discrete phrase table-based approach, and S_R is not robust for out-of-vocabulary of word vector. The other case is that the human annotators gave high scores, but we found no connectivity and content relatedness in the utterance pairs. We found that some utterance pairs without any connectivity and content relatedness can be judged as acceptable by the human annotators since they can imagine the underlying context and situation of the utterance pairs using human world knowledge, such as commonsense. We think this is a challenging issue that exceeds our focus in this paper, and thus, remains as future work.

Does our filtering undermine diversity? One might think that our method succeeds in filtering by assigning high scores to generic responses such as dull responses. This concern makes sense since it is

| Training data | # of pairs | Automatic evaluation | | | | Human evaluation | | |
|-----------------------------|------------|----------------------|------------------|------------------|------------|------------------|----------------|----------------|
| | | len | distinct-1 | distinct-2 | BLEU-1 | Avg. | ✗ [↓] | ✓ [↑] |
| non-filtered | 79,445,453 | 8.44 | 127/0.030 | 238/0.064 | 8.8 | 3.37 | 38 % | 62 % |
| Csáky et al. (2019) SRC | 40,000,000 | 7.97 | 165/0.041 | 329/0.094 | 9.1 | 3.56 | 25 % | 75 % |
| Csáky et al. (2019) TRG | 40,000,000 | 18.25 | 213/0.023 | 591/0.069 | 5.4 | 2.85 | 65 % | 35 % |
| Junczys-Dowmunt (2018) | 40,000,000 | 8.63 | 206/0.048 | 478/0.125 | 9.4 | 3.43 | 32 % | 68 % |
| Ours S_{C+R} | 40,000,000 | 7.13 | 345/0.097 | 853/0.278 | 9.4 | 3.73 | 15 % | 85 % |
| Ours S_C (ablation study) | 40,000,000 | 7.31 | 201/0.055 | 466/0.148 | 9.2 | 3.69 | 19 % | 81 % |
| Ours S_R (ablation study) | 40,000,000 | 7.91 | 270/0.068 | 662/0.192 | 9.4 | 3.76 | 20 % | 80 % |
| reference | | 9.04 | 1301/0.288 | 3244/0.807 | - | - | - | - |

Table 5: Evaluation results for generated responses (English; filtered out 50%). **Bold** denotes the best results. The ✗/✓ shows the percentages of the low/high scored responses (i.e., human scores in [1, 3] or in [3, 5]).

known that dialogue systems learned from the training data, including many generic utterances, tend to generate bland responses (Csáky et al., 2019). To answer this interesting question, we confirmed the diversity of utterance pairs with a high score (i.e., remained as training data) and a low score (i.e., removed from training data) in our S_{C+R} (Table 4).¹² As a result, there was no significant difference between them. Therefore, we conclude that the proposed method does not prefer only generic responses and maintains the diversity of data. It is an essential future attempt to improve the quality of dialogue data further (e.g., more diversity) after using the proposed method to remove unacceptable pairs.

7 Experiments: Response Generation

This section reports on the effectiveness of the proposed method for filtering noisy data in neural response generation.

7.1 Experimental Setup

Training. We obtained the filtered training data \mathcal{D}' by removing utterance pairs with low scores from the original dataset \mathcal{D} (approximately 10% or 50% of total utterance pairs were removed). As a response generation model, we used a Transformer (Vaswani et al., 2017) based encoder-decoder model implemented in the fairseq toolkit (Ott et al., 2019).¹³ Transformer has demonstrated high performance in response generation (Dinan et al., 2019) and other NLP tasks.

Automatic evaluation. Here, we report the following metrics: the average response length in tokens (len), type-token ratio for {1, 2}-grams

¹²See Appendix E.3 for more extensive result.

¹³See Appendix D for training details.

1: **Well, I thought we'd wait for Mr. Oliver.**

A: I don't think so.

B: I'll be right back.

C: I'm sorry, but I have to get back to work.

D: No, I don't think so.

Ours: Oh, I'm sorry.

2: **Are you heading home?**

A: No, I don't think so.

B: No, I don't think so.

C: No, no, no, no, no, no, no, no, no, no, no, no.

D: No, I don't think so.

Ours: No, I'm not.

3: **Because we want to keep him safe.**

A: I don't know what you're talking about.

B: I don't know what you're talking about.

C: I don't care if he's dead or alive or dead.

D: I don't know what you're talking about.

Ours: We need to get him out of here.

4: **Quite soon, now, you'll be able to travel.**

A: I don't think so.

B: I'll be back in a minute.

C: I'll be back in a few days, then I'll be back.

D: I don't want to see you again.

Ours: No, I'm not going anywhere.

5: **Which is where exactly?**

A: I don't know.

B: I don't know.

C: I don't know what you're talking about, but I don't know where you are.

D: I don't know.

Ours: I don't know.

Table 6: Samples of input utterances (**bold**) and responses generated by models trained on filtered data by proposed method S_{C+R} (Ours). Other responses are generated by models trained on (A) non-filtered data and filtered data by baselines; (B) Csáky et al. (2019)'s SRC method, (C) their TRG method, and (D) Junczys-Dowmunt (2018)'s method, respectively.

(distinct- $\{1, 2\}$), and BLEU-1 (Papineni et al., 2002). The latter was used as a reference-based metric; while it is widely used in previous studies (Zhao and Eskenazi, 2017; Baheti et al., 2018; Csáky et al., 2019), some studies (e.g., (Liu et al., 2016)) have reported that BLEU-1 may not be highly correlated with the human evaluation of response generation.¹⁴

Human evaluation. We evaluated the quality of the generated responses manually. We asked human evaluators recruited via Amazon Mechanical Turk to evaluate responses that are generated for 100¹⁵ input utterances randomly sampled from the test data. We used the same task setting and protocol as described in Section 6.1 to obtain the human scores for each pair. Higher human scores indicate that the better results.

7.2 Results and Analysis

Table 5 shows the results of automatic and human evaluations of the generated responses. The model trained on the data filtered using the proposed method S_{C+R} produced more than three times as many distinct $\{1, 2\}$ -grams as the model trained on non-filtered data. Furthermore, it outperformed the model trained on non-filtered data in the human evaluation, achieving the highest percentage of acceptable responses of 85%. Additionally, these results of our S_{C+R} were better than other baselines. To conclude, these experimental results indicate that the proposed scoring method can help generate diverse responses that are judged as acceptable by humans.

This experiment provides empirical evidence for supporting our hypothesis that the performance of neural response generation models can be improved by just removing unacceptable utterance pairs from training data, which answers the research question formulated at the start of this paper.

8 Multilingual Availability

While the proposed method S_{C+R} was tested on an English corpus, it can potentially work for other languages as well. To demonstrate this, we selected Japanese dialogue data as another case study.¹⁶ The linguistic phenomena in Japanese are quite different from those in English, thus making this

¹⁴See Appendix E.5 for more extensive evaluation results.

¹⁵Same size as Shen et al. (2017) and Bao et al. (2020).

¹⁶See Appendix F for all experimental results on Japanese.

| Scoring method | Spearman’s ρ | p -value |
|-----------------------------|-------------------|--|
| Csáky et al. (2019) SRC | -0.0553 | 4.4×10^{-1} |
| Csáky et al. (2019) TRG | -0.0366 | 6.1×10^{-1} |
| Junczys-Dowmunt (2018) | 0.1074 | 1.3×10^{-1} |
| Ours S_{C+R} | 0.2491 | 3.8×10^{-4} |
| Ours S_C (ablation study) | 0.1395 | 4.9×10^{-2} |
| Ours S_R (ablation study) | 0.1504 | 3.3×10^{-2} |

Table 7: Correlation coefficient between human scores and automatically computed scores (Japanese).




Figure 3: Distribution between human scores and our S_{C+R} (Japanese).

experiment to be a good test of the applicability of the proposed method to non-English languages.

Japanese dataset. We prepare the Japanese dialogue data from Japanese OpenSubtitles (Lison et al., 2018) containing roughly 3M lines. We obtain 1,893,477 utterance pairs as our training data, which excludes our test and validation data.¹⁷

8.1 Data Scoring

Settings. To compute S_C , we defined a low co-occurrence frequency as less than 20, considering the size of the Japanese corpus, and consequently obtained the key phrase pairs $|\mathcal{P}| = 19,992$. To compute S_R , we used pre-trained fastText (Grave et al., 2018) and learned common components from all sentences in the training data.

For human evaluation, we used Yahoo! crowdsourcing¹⁸ to hire native Japanese-speaking workers. The task setting and protocol are the same as those for English (Section 6.1), regardless of the crowdsourcing platform.

Results and analysis. Table 7 shows the correlation between human scores and those automatically computed by each method. Our method S_{C+R} has the highest correlation with human scores, although the overall result is lower than that obtained for the English dataset. Figure 3 shows the distribution of

¹⁷See Appendix B for details on our data such as the preparation procedure and statistics.

¹⁸<https://crowdsourcing.yahoo.co.jp/>

| Training data | # of pairs | Automatic evaluation | | | | Human evaluation | | |
|-----------------------------|------------|----------------------|------------------|------------------|-------------|------------------|----------------|----------------|
| | | len | distinct-1 | distinct-2 | BLEU-1 | Avg. | ✗ [↓] | ✓ [↑] |
| non-filtered | 1,893,477 | 5.91 | 268/0.091 | 509/0.207 | 13.4 | 3.35 | 39 % | 61 % |
| Csáky et al. (2019) SRC | 1,700,000 | 5.75 | 295/0.102 | 550/0.231 | 13.2 | 3.47 | 37 % | 63 % |
| Csáky et al. (2019) TRG | 1,700,000 | 7.06 | 336/0.095 | 662/0.219 | 11.6 | 3.37 | 34 % | 66 % |
| Junczys-Dowmunt (2018) | 1,700,000 | 5.31 | 284/0.107 | 516/0.240 | 12.6 | 3.46 | 32 % | 68 % |
| Ours S_{C+R} | 1,700,000 | 5.68 | 319/0.112 | 582/0.249 | 13.9 | 3.61 | 27 % | 73 % |
| Ours S_C (ablation study) | 1,700,000 | 5.51 | 264/0.096 | 492/0.218 | 13.7 | 3.44 | 32 % | 68 % |
| Ours S_R (ablation study) | 1,700,000 | 5.73 | 296/0.103 | 555/0.234 | 12.5 | 3.56 | 30 % | 70 % |
| reference | | 7.29 | 750/0.206 | 1446/0.460 | - | - | - | - |

Table 8: Evaluation results for generated responses (Japanese; filtered out 10%). **Bold** denotes the best results. The **✗/✓** shows the percentages of the low/high scored responses (i.e., human scores in $[1, 3)$ or in $[3, 5)$).

our S_{C+R} corresponding to human scores. Similar to the result obtained for English as presented in Figure 2 (c), S_{C+R} rarely overestimates utterance pairs with low human scores but underestimates those with high human scores in Japanese.

8.2 Response Generation

Settings. We used the same experimental settings described in Section 7.1 for the preparation of filtered data \mathcal{D}' and model training.

Results and analysis. Table 8 shows the results of evaluations of the generated responses. The filtered data generated by S_{C+R} provided the best results in terms of almost all the metrics, including human evaluation. It supports our hypothesis that the proposed method is also suitable for non-English languages.

9 Relationship with Evaluation Metric

The proposed method S_{C+R} maps an utterance pair to a score (scalar value) in terms of the quality of dialogue. That is, formally, our method is similar to the reference-free automatic evaluation metrics for dialogue agents; both of them evaluate the response given an input utterance and also map into a score. Recently, the novel reference-free metrics for evaluating generated responses such as USR (Mehri and Eskenazi, 2020) or MAUDE (Sinha et al., 2020) were developed. While it is possible to use them as a scoring method for filtering noisy data, in theory, there are some concerns with applying them in practice. One is the difference of the data of interest; since evaluation metrics are intended for responses generated as dialogue, i.e., somewhat valid dialogue data, it is unclear whether they also work for apparently noisy data. Another one is the difference of desired

properties; evaluation metrics need to be sensitive to “how good is it?” while the filtering requires to detect “is it a dialogue?” with high accuracy. It would be interesting to investigate the effectiveness of reference-free metrics for noisy dialogue data filtering tasks, and vice versa. We leave these investigations for future work.

In contrast, reference-based metrics require a reference response (i.e., ground truth) when they calculate scores; such metrics include the traditional overlap-based BLEU, ROUGE, METEOR, embedding-based metrics (Liu et al., 2016), and neural network-based RUBER (Tao et al., 2018) and ADEM (Lowe et al., 2017a) Thus, these methods cannot straightforwardly be considered as alternatives to the proposed method, which aims at filtering.

10 Conclusion

In light of the success of noisy corpus filtering in neural machine translation, we attempted to filter out unacceptable utterance pairs from large dialogue corpora in an unsupervised manner. The proposed scoring method estimates the quality of utterance pairs by focusing on the two crucial aspects of dialogue, namely, the *connectivity* and *content relatedness* of utterance pairs. We demonstrated that our scoring method has a higher correlation with human judgment than recently proposed methods. Furthermore, we provided empirical evidence that our method improves the performance of a response generation model by removing unacceptable utterance pairs from its training data. We hope that this study will facilitate discussions in the dialogue response generation community regarding the issue of filtering noisy corpora.

Acknowledgments

This work was supported by JSPS KAKENHI Grant Numbers JP19H04162, JP19J21913.

References

- Daniel Adiwardana, Minh-Thang Luong, David R So, Jamie Hall, Noah Fiedel, Romal Thoppilan, Zi Yang, Apoorv Kulshreshtha, Gaurav Nemade, Yifeng Lu Quoc, and V Le. 2020. [Towards a Human-like Open-Domain Chatbot](#). In *arXiv preprint arXiv:2001.09977*.
- Reina Akama, Kazuaki Inada, Naoya Inoue, Sosuke Kobayashi, and Kentaro Inui. 2017. [Generating Stylistically Consistent Dialog Responses with Transfer Learning](#). In *Proceedings of the 8th International Joint Conference on Natural Language Processing (IJCNLP)*, volume 2, pages 408–412.
- Sanjeev Arora, Yingyu Liang, and Tengyu Ma. 2017. [A Simple but Tough-to-Beat Baseline for Sentence Embeddings](#). In *5th International Conference on Learning Representations (ICLR)*.
- Ashutosh Baheti, Alan Ritter, Jiwei Li, and Bill Dolan. 2018. [Generating More Interesting Responses in Neural Conversation Models with Distributional Constraints](#). In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 3970–3980.
- Siqi Bao, Huang He, Fan Wang, Hua Wu, and Haifeng Wang. 2020. [PLATO: Pre-trained Dialogue Generation Model with Discrete Latent Variable](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics (ACL)*, pages 85–96.
- Piotr Bojanowski, Edouard Grave, Armand Joulin, and Tomas Mikolov. 2017. [Enriching Word Vectors with Subword Information](#). *Transactions of the Association for Computational Linguistics (TACL)*, 5:135–146.
- Gerlof Bouma. 2009. [Normalized \(Pointwise\) Mutual Information in Collocation Extraction](#). In *Proceedings of the International Conference of the German Society for Computational Linguistics and Language Technology (GSCL)*, pages 31–40.
- Hyundong Cho and Jonathan May. 2020. [Grounding Conversations with Improvised Dialogues](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics (ACL)*, pages 2398–2413.
- Alexis Conneau, Douwe Kiela, Holger Schwenk, Loic Barrault, and Antoine Bordes. 2017. [Supervised Learning of Universal Sentence Representations from Natural Language Inference Data](#). In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 670–680.
- Richárd Csáky, Patrik Purgai, and Gábor Recski. 2019. [Improving Neural Conversational Models with Entropy-Based Data Filtering](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics (ACL)*, pages 5650–5669.
- Emily Dinan, Stephen Roller, Kurt Shuster, Angela Fan, Michael Auli, and Jason Weston. 2019. [Wizard of Wikipedia: Knowledge-Powered Conversational Agents](#). In *7th International Conference on Learning Representations (ICLR)*.
- Chris Dyer, Victor Chahuneau, and Noah A Smith. 2013. [A Simple, Fast, and Effective Reparameterization of IBM Model 2](#). In *Proceedings of the 2013 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies (NAACL-HLT)*, pages 644–648.
- Michel Galley, Chris Brockett, Alessandro Sordoni, Yangfeng Ji, Michael Auli, Chris Quirk, Margaret Mitchell, Jianfeng Gao, and Bill Dolan. 2015. [deltaBLUE: A Discriminative Metric for Generation Tasks with Intrinsically Diverse Targets](#). In *Proceedings of the 53rd Annual Meeting of the Association for Computational Linguistics and the 7th International Joint Conference on Natural Language Processing (ACL-IJCNLP)*, volume 2, pages 445–450.
- Edouard Grave, Piotr Bojanowski, Prakhar Gupta, Armand Joulin, and Tomas Mikolov. 2018. [Learning Word Vectors for 157 Languages](#). In *Proceedings of the 11th International Conference on Language Resources and Evaluation (LREC)*, pages 3483–3487.
- Matthew Henderson, Paweł Budzianowski, Iñigo Casanueva, Sam Cope, Daniela Gerz, Girish Kumar, Nikola Mrkšić, Georgios Spithourakis, Pei-Hao Su, Ivan Vulić, and Tsung-Hsien Wen. 2019. [A Repository of Conversational Datasets](#). In *Proceedings of the First Workshop on NLP for Conversational AI*, pages 1–10.
- Marcin Junczys-Dowmunt. 2018. [Dual Conditional Cross-Entropy Filtering of Noisy Parallel Corpora](#). In *Proceedings of the 3rd Conference on Machine Translation: Shared Task Papers (WMT)*, pages 888–895.
- Philipp Koehn, Hieu Hoang, Alexandra Birch, Chris Callison-Burch, Marcello Federico, Nicola Bertoldi, Brooke Cowan, Wade Shen, Christine Moran Mit, Richard Zens, Rwth Aachen, Chris Dyer, Ondřej Bojar, and Evan Herbst Cornell. 2007. [Moses: Open Source Toolkit for Statistical Machine Translation](#). In *Proceedings of the 45th Annual Meeting of the Association for Computational Linguistics Companion (ACL) Volume Proceedings of the Demo and Poster Sessions*, pages 177–180.
- Philipp Koehn, Huda Khayrallah, Kenneth Heafield, and Mikel L Forcada. 2018. [Findings of the WMT 2018 Shared Task on Parallel Corpus Filtering](#). In

- Proceedings of the 3rd Conference on Machine Translation: Shared Task Papers (WMT)*, pages 726–739.
- Philipp Koehn and Rebecca Knowles. 2017. **Six Challenges for Neural Machine Translation**. In *Proceedings of the First Workshop on Neural Machine Translation*, pages 28–39.
- Jiwei Li, Michel Galley, Chris Brockett, Jianfeng Gao, and Bill Dolan. 2016. **A Diversity-Promoting Objective Function for Neural Conversation Models**. In *Proceedings of the 2016 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies (NAACL-HLT)*, pages 110–119.
- Rensis Likert. 1932. **A technique for the measurement of attitudes**. *Archives of psychology*, 22(140):1–55.
- Pierre Lison and Jörg Tiedemann. 2016. **OpenSubtitles2016: Extracting Large Parallel Corpora from Movie and TV Subtitles**. In *Proceedings of the 10th International Conference on Language Resources and Evaluation (LREC)*, pages 923–929.
- Pierre Lison, Jörg Tiedemann, and Milen Kouylekov. 2018. **OpenSubtitles2018: Statistical Rescoring of Sentence Alignments in Large, Noisy Parallel Corpora**. In *Proceedings of the 11th International Conference on Language Resources and Evaluation (LREC)*, pages 1742–1748.
- Chia-Wei Liu, Ryan Lowe, Iulian V Serban, Michael Noseworthy, Laurent Charlin, and Joelle Pineau. 2016. **How NOT To Evaluate Your Dialogue System: An Empirical Study of Unsupervised Evaluation Metrics for Dialogue Response Generation**. In *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 2122–2132.
- Ryan Lowe, Michael Noseworthy, Iulian V Serban, Nicolas A-Gontier, Yoshua Bengio, and Joelle Pineau. 2017a. **Towards an Automatic Turing Test: Learning to Evaluate Dialogue Responses**. In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (ACL)*, pages 1116–1126.
- Ryan Lowe, Michael Noseworthy, Iulian Vlad Serban, Nicolas Angelard-Gontier, Yoshua Bengio, and Joelle Pineau. 2017b. **Towards an Automatic Turing Test: Learning to Evaluate Dialogue Responses**. In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (ACL)*, volume 1, pages 1116–1126.
- Marco Marelli, Luisa Bentivogli, Marco Baroni, Raffaella Bernardi, Stefano Menini, and Roberto Zamparelli. 2014a. **SemEval-2014 Task 1: Evaluation of Compositional Distributional Semantic Models on Full Sentences through Semantic Relatedness and Textual Entailment**. In *Proceedings of the 8th International Workshop on Semantic Evaluation (SemEval)*, pages 1–8.
- Marco Marelli, Stefano Menini, Marco Baroni, Luisa Bentivogli, Raffaella Bernardi, and Roberto Zamparelli. 2014b. **A SICK cure for the evaluation of compositional distributional semantic models**. In *Proceedings of the 9th International Conference on Language Resources and Evaluation (LREC)*, pages 216–223.
- Shikib Mehri and Maxine Eskenazi. 2020. **USR: An Unsupervised and Reference Free Evaluation Metric for Dialog Generation**. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics (ACL)*, pages 681–707.
- Tomas Mikolov, Edouard Grave, Piotr Bojanowski, Christian Puhresch, and Armand Joulin. 2018. **Advances in Pre-Training Distributed Word Representations**. In *Proceedings of the 11th International Conference on Language Resources and Evaluation (LREC)*, pages 52–55.
- Makoto Morishita, Jun Suzuki, and Masaaki Nagata. 2018. **NTT’s Neural Machine Translation Systems for WMT 2018**. In *Proceedings of the 3rd Conference on Machine Translation: Shared Task Papers (WMT)*, pages 461–466.
- 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) (NAACL-HLT)*, pages 48–53.
- Kishore Papineni, Salim Roukos, Todd Ward, and Wei-Jing Zhu. 2002. **BLEU: a Method for Automatic Evaluation of Machine Translation**. In *Proceedings of the 40th Annual Meeting of the Association for Computational Linguistics (ACL)*, pages 311–318.
- Jiaxin Pei and Chenliang Li. 2018. **S2SPMN: A Simple and Effective Framework for Response Generation with Relevant Information**. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 745–750.
- Alan Ritter, Colin Cherry, and William B Dolan. 2011. **Data-Driven Response Generation in Social Media**. In *Proceedings of the 2011 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 583–593.
- Harvey Sacks. 1989. **Lecture One: Rules of Conversational Sequence**. *Human Studies*, 12(3/4):217–233.
- João Sedoc, Daphne Ippolito, Arun Kirubakaran, Jai Thirani, Lyle Ungar, and Chris Callison-Burch. 2019. **ChatEval: A Tool for Chatbot Evaluation**. In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics (NAACL) (Demonstrations)*, pages 60–65.

- Rico Sennrich, Barry Haddow, and Alexandra Birch. 2016. [Neural Machine Translation of Rare Words with Subword Units](#). In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (ACL)*, volume 1, pages 1715–1725.
- Rico Sennrich and Biao Zhang. 2019. [Revisiting Low-Resource Neural Machine Translation: A Case Study](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics (ACL)*, pages 211–221.
- Lifeng Shang, Zhengdong Lu, and Hang Li. 2015. [Neural Responding Machine for Short-Text Conversation](#). In *Proceedings of the 53rd Annual Meeting of the Association for Computational Linguistics and the 7th International Joint Conference on Natural Language Processing (ACL-IJCNLP)*, volume 1, pages 1577–1586.
- Mingyue Shang, Zhenxin Fu, Nanyun Peng, Yansong Feng, Dongyan Zhao, and Rui Yan. 2018. [Learning to Converse with Noisy Data: Generation with Calibration](#). In *Proceedings of the Twenty-Seventh International Joint Conference on Artificial Intelligence (IJCAI-18)*, pages 4338–4344.
- Xiaoyu Shen, Hui Su, Yanran Li, Wenjie Li, Shuzi Niu, Yang Zhao, Akiko Aizawa, and Guoping Long. 2017. [A Conditional Variational Framework for Dialog Generation](#). In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (ACL)*, volume 2, pages 504–509.
- Jack Sidnell. 2010. *Conversation Analysis: An Introduction*. Language in Society. John Wiley & Sons.
- Koustuv Sinha, Prasanna Parthasarathi, Jasmine Wang, Ryan Lowe, William L Hamilton, and Joelle Pineau. 2020. [Learning an Unreferenced Metric for Online Dialogue Evaluation](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 2430–2441. Association for Computational Linguistics.
- Sandeep Subramanian, Adam Trischler, Yoshua Bengio, and Christopher J Pal. 2018. [Learning General Purpose Distributed Sentence Representations via Large Scale Multi-task Learning](#). In *6th International Conference on Learning Representations (ICLR)*.
- Ilya Sutskever, Oriol Vinyals, and Quoc V. Le. 2014. [Sequence to Sequence Learning with Neural Networks](#). In *Advances in Neural Information Processing Systems 27 (NIPS)*, pages 3104–3112.
- Chongyang Tao, Lili Mou, Dongyan Zhao, and Rui Yan. 2018. [RUBER: An Unsupervised Method for Automatic Evaluation of Open-Domain Dialog Systems](#). In *Proceedings of the 32nd AAAI Conference on Artificial Intelligence (AAAI-18)*.
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N. Gomez, Łukasz Kaiser, and Illia Polosukhin. 2017. [Attention is All you Need](#). In *Advances in Neural Information Processing Systems 30 (NIPS)*, pages 5998–6008.
- Oriol Vinyals and Quoc Le. 2015. [A Neural Conversational Model](#). In *Proceedings of the 31st International Conference on Machine Learning (ICML) Deep Learning Workshop*.
- Chen Xing, Wei Wu, Yu Wu, Jie Liu, Yalou Huang, Ming Zhou, and Wei-Ying Ma. 2017. [Topic Aware Neural Response Generation](#). In *Proceedings of the 31st AAAI Conference on Artificial Intelligence*, pages 3351–3357.
- Xinnuo Xu, Ondře Dušek, Ioannis Konstas, and Verena Rieser. 2018a. [Better Conversations by Modeling, Filtering, and Optimizing for Coherence and Diversity](#). In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 3981–3991.
- Zhen Xu, Nan Jiang, Bingquan Liu, Wenge Rong, Bowen Wu, Baoxun Wang, Zhuoran Wang, and Xiaolong Wang. 2018b. [LSDSCC: a Large Scale Domain-Specific Conversational Corpus for Response Generation with Diversity Oriented Evaluation Metrics](#). In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies (NAACL-HLT)*, volume 1, pages 2070–2080.
- Tiancheng Zhao and Maxine Eskenazi. 2017. [Learning Discourse-level Diversity for Neural Dialog Models using Conditional Variational Autoencoders](#). In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics*, pages 654–664.

A Preliminary Experiment Settings

Dataset. For our preliminary experiment (Section 1), we use OpenSubtitles (Lison et al., 2018)¹⁹ in English, one of the largest corpora of movie scripts that are freely available and has been used in many data-driven dialogue response generations. We automatically obtained dialogue paired-data from the corpus which does not contain speaker annotations on the dialogue turns *following the previous methods* (Vinyals and Le, 2015; Li et al., 2016). Specifically, we extracted the consecutive two lines as an utterance pair based on the assumption that each line corresponds to a full-speaker’s turn. We collected pairs from the dataset in which the length of the utterance and response was 3-25 words each and obtained the dialogue dataset. For counting the number of words, we used SpaCy²⁰ to tokenize each utterance and response.

Evaluation settings. We used Amazon Mechanical Turk (MTurk) to evaluate the data manually. In our experiments, randomly sampled 100 utterance pairs were evaluated by native English speakers. We filtered out unreliable workers by integrating attention checks. We requested five workers to evaluate each pair with a five-point Likert scale (5: Strongly agree to 1: Strongly disagree) (Likert, 1932) as an answer to the following question: *Is the sequence of the two utterances acceptable as a dialogue?*

Result. As a result of our preliminary experiment, we discover that, out of all scores given for utterance pairs, 25% was unacceptable (scored as 1: Strongly disagree or 2: Disagree) and almost half was acceptable (scored as 5: Strongly agree or 4: Agree). The inter-annotator agreement (Krippendorff’s alpha) was 0.33.

B Dialogue Data Constructions

English Dataset. In our experiments (Section 6, 7), we used the OpenSubtitles as an example of the noisy million-scales English dialogue corpus. In addition to the previous method for the extraction of pair data (as described in Appendix A), we cleaned the data with some heuristic preprocesses. Some processings were inspired by the technique of noisy-parallel corpus filtering on NMT fields.

¹⁹<http://opus.nlpl.eu/OpenSubtitles-v2018.php>

²⁰<https://spacy.io/>

The additional preprocesses that we conducted are as follows:

- Using languid²¹, which is a tool that detects the language for given sentences, we removed the utterance pairs judged as any language other than the target language.
- Removed the parrot-back utterance pairs.
- Removed duplicate utterance pairs in order to remove the completely repeated conversational sequences, such as the opening scenes of serial dramas.

Eventually, we obtained 79,621,506 utterance pairs as our English dialogue corpus. For our experiments, we divided them into training/validation/test data. Table 9 shows the statistics of our English dataset. The “# pairs” indicates the number of utterance pairs obtained by the previous method described in Appendix A.

| Data | # works | # lines | # pairs | # our pairs |
|--------|---------|-------------|-------------|-------------|
| Corpus | 446,612 | 441,452,475 | 230,597,913 | 79,621,506 |
| Train | 442,433 | 441,065,310 | 230,392,431 | 79,445,453 |
| Valid | 200 | 195,297 | 104,007 | 90,317 |
| Test | 200 | 191,868 | 101,475 | 85,736 |

Table 9: The statistics of our English dataset.

Japanese Dataset. For our other experiment (Section 8), we prepared our Japanese corpus from OpenSubtitles. The data construction process, including preprocesses, is the same as those for English (as described in Appendix A and the previous paragraph). We used mecab²² to tokenize the Japanese utterances. Eventually, we obtained 1,917,721 utterance pairs as our Japanese dialogue corpus. For our experiments, we divided them into training/validation/test data. Table 10 shows the statistics of our Japanese dataset. The “# pairs” indicates the number of utterance pairs obtained by the previous method described in Appendix A.

| Data | # works | # lines | # pairs | # our pairs |
|---------|---------|-----------|-----------|-------------|
| Corpus: | 3,546 | 3,170,155 | 2,266,127 | 1,917,721 |
| Train | 3,506 | 3,135,812 | 2,240,847 | 1,893,477 |
| Valid | 20 | 15,489 | 11,939 | 11,486 |
| Test | 20 | 18,854 | 13,341 | 12,758 |

Table 10: The statistics of our Japanese dataset.

²¹<https://github.com/saffsd/langid.py>

²²<https://taku910.github.io/mecab/>

C Experimental Details of Proposed Method

C.1 Computing Connectivity as S_C

To compute alignment points with fastAlign, we set the null alignment probability to 0.5 and used the ‘grow-diag-final’ heuristics. To extract phrase pairs with Moses using the information of alignment points, we used the following settings: alignment=‘grow-diag-final-and’, reordering=‘msd-bidirectional-fe’, and first-step=4. Furthermore, we extended the standard phrase extraction algorithm (Algorithm 1) to only extract phrases that have at least one alignment point for every row and column when considering the matrix view of phrases (Algorithm 2). This is because unaligned words should not be positively dealt with in the evaluation of connectivity.

C.2 Computing Content relatedness as S_R

For SIF weighting in Japanese data, we obtained word frequency data from jawiki dataset²³ following English word frequency data provided in the author’s implementation.²⁴

D Training Details for Response Generation Model

To obtain the response generation model, we used a Transformer (Vaswani et al., 2017) based encoder-decoder model implemented in the fairseq toolkit²⁵ (Ott et al., 2019). We used ‘--arch transformer_wmt_en_de_big’ option with its default configuration, and set the number of maximum training steps to 100K. We used the byte pair encoding²⁶ (Sennrich et al., 2016) for token segmentation and set its vocabulary size to 16K. The numbers of parameters in our models were roughly 223M. We trained our models on 8 NVIDIA DGX-1 Tesla V100 GPUs. It took approximately 6 hours for training one model.

E Experimental Results on English

E.1 Distributions between Human and Automatic Scoring

Figure 4 shows that, for all the models including ablations, the distributions between human scores and automatically computed scores.

²³<https://dumps.wikimedia.org/jawiki/>

²⁴<https://github.com/PrincetonML/SIF>

²⁵<https://github.com/pytorch/fairseq>

²⁶<https://github.com/rsennrich/subword-nmt>

Algorithm 1 Phrase pair extraction

Input: word alignment \mathcal{A} for sentence pair (x, y)
Output: set of phrase pair $\overline{\mathcal{P}}$

```
1: for  $f_{\text{start}} \leftarrow 1, \dots, |x|$  do
2:   for  $f_{\text{end}} \leftarrow f_{\text{start}}, \dots, |x|$  do
3:      $e_{\text{start}} \leftarrow |y|$ 
4:      $e_{\text{end}} \leftarrow 0$ 
5:     for all  $(f, e) \in \mathcal{A}$  do
6:       if  $f_{\text{start}} \leq f \leq f_{\text{end}}$  then
7:          $e_{\text{start}} \leftarrow \min(e, e_{\text{start}})$ 
8:          $e_{\text{end}} \leftarrow \max(e, e_{\text{end}})$ 
9:       end if
10:    end for
11:    add extract( $e_{\text{start}}, e_{\text{end}}, f_{\text{start}}, f_{\text{end}}$ ) to set  $\overline{\mathcal{P}}$ 
12:  end for
13: end for
```

function extract($e_{\text{start}}, e_{\text{end}}, f_{\text{start}}, f_{\text{end}}$)

```
1: return {} if  $e_{\text{end}} = 0$ 
2: for all  $(e, f) \in \mathcal{A}$  do
3:   return {} if  $f < f_{\text{start}}$  OR  $f > f_{\text{end}}$ 
4: end for
5:  $F = \{\}$ 
6:  $e_s \leftarrow e_{\text{start}}$ 
7: repeat
8:    $e_e \leftarrow e_{\text{end}}$ 
9:   repeat
10:    add phrase pair  $(f_{\text{start}} \dots f_{\text{end}}, e_s \dots e_e)$  to set  $F$ 
11:     $e_e ++$ 
12:   until  $e_e$  aligned
13:    $e_s --$ 
14: until  $e_s$  aligned
15: return  $F$ 
```

Algorithm 2 Modified phrase pair extraction

Input: word alignment \mathcal{A} for sentence pair (x, y)
Output: set of phrase pair $\overline{\mathcal{P}}$

```
1: for  $f_{\text{start}} \leftarrow 1, \dots, |x|$  do
2:   for  $f_{\text{end}} \leftarrow f_{\text{start}}, \dots, |x|$  do
3:      $e_{\text{start}} \leftarrow |y|$ 
4:      $e_{\text{end}} \leftarrow 0$ 
5:      $\mathcal{F} = \{\}$ 
6:      $\mathcal{E} = \{\}$ 
7:     for all  $(f, e) \in \mathcal{A}$  do
8:       if  $f_{\text{start}} \leq f \leq f_{\text{end}}$  then
9:          $e_{\text{start}} \leftarrow \min(e, e_{\text{start}})$ 
10:         $e_{\text{end}} \leftarrow \max(e, e_{\text{end}})$ 
11:         $\mathcal{F} \leftarrow \mathcal{F} \cup \{f\}$ 
12:         $\mathcal{E} \leftarrow \mathcal{E} \cup \{e\}$ 
13:      end if
14:    end for
15:    add extract( $e_{\text{start}}, e_{\text{end}}, f_{\text{start}}, f_{\text{end}}, \mathcal{F}, \mathcal{E}$ ) to set  $\overline{\mathcal{P}}$ 
16:  end for
17: end for
```

function extract($e_{\text{start}}, e_{\text{end}}, f_{\text{start}}, f_{\text{end}}, \mathcal{F}, \mathcal{E}$)

```
1: return {} if  $e_{\text{end}} = 0$ 
2: for  $f \leftarrow f_{\text{start}}, \dots, f_{\text{end}}$  do
3:   return {} if  $f \notin \mathcal{F}$ 
4: end for
5: for  $e \leftarrow e_{\text{start}}, \dots, e_{\text{end}}$  do
6:   return {} if  $e \notin \mathcal{E}$ 
7: end for
8: return phrase pair  $(f_{\text{start}} \dots f_{\text{end}}, e_{\text{start}} \dots e_{\text{end}})$ 
```

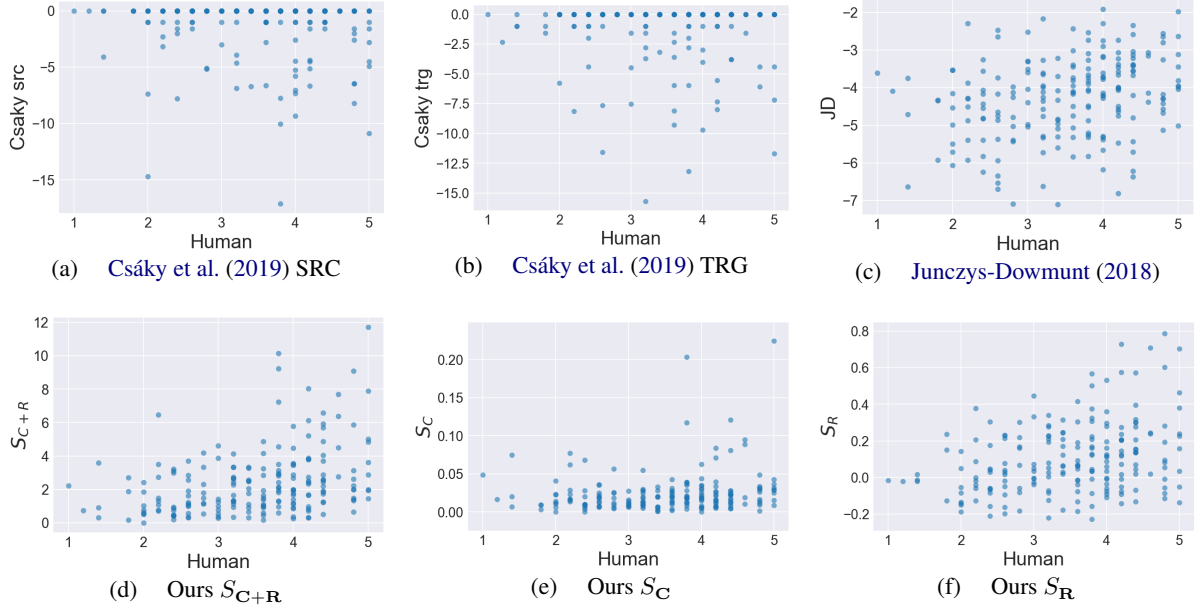


Figure 4: Distributions between human scores and automatically computed scores by each method (English).




Figure 5: Score distributions of our S_C , S_R , S_{C+R} across our training data (English).

| Utterance | Response | S_C | S_R | S_{C+R} | Human |
|--|---------------------------------------|-------|-------|-------------|-------|
| 1: What happened to your hand? | Just a scratch. | 1.38 | 0.00 | 1.38 | 4.8 |
| 2: But Carcharodontosaurus has the more lethal bite. | This time, the Spinosaurus triumphed. | 0.22 | 0.68 | 1.72 | 4.0 |
| 3: I'm right here with you. | Come on, boys. | 1.39 | 0.00 | 1.39 | 4.0 |
| 4: What Is It Officer Chan? | Brother Ho, I must leave now, | 1.04 | 0.68 | 1.72 | 4.4 |
| 5: Out on the balcony. | You shouldn't have come. | 0.30 | 0.00 | 0.30 | 4.2 |

Table 11: Samples of utterance pairs that cause low recall scored with our method and human judgements (English). The scores of S_C and S_R were normalized by α, β .

| Scored data | Utterance pair | | | Utterance (source-side) | | | Response (target-side) | | |
|---------------------|----------------|------------|------------|-------------------------|------------|------------|------------------------|------------|------------|
| | len | distinct-1 | distinct-2 | len | distinct-1 | distinct-2 | len | distinct-1 | distinct-2 |
| Top 50% (remained) | 18.06 | 0.018 | 0.313 | 9.05 | 0.028 | 0.474 | 9.02 | 0.028 | 0.472 |
| Worst 50% (removed) | 17.92 | 0.019 | 0.316 | 8.92 | 0.030 | 0.476 | 9.00 | 0.030 | 0.470 |
| Top 90% (remained) | 17.99 | 0.013 | 0.229 | 9.00 | 0.022 | 0.382 | 8.99 | 0.022 | 0.378 |
| Worst 10% (removed) | 17.99 | 0.047 | 0.654 | 8.85 | 0.068 | 0.872 | 9.14 | 0.067 | 0.873 |

Table 12: Comparison of the top and the worst utterance pairs in the training data scored by our method (English).

| English | # of pairs | len | dist1 | dist2 | B1 | bp | B1bp | MET | ROU | CID | EA | VE | GM |
|--------------------------|------------|-------|------------|------------|-----|------|------|------|------|-------|-------|-------|-------|
| non-filtered | 79,445,453 | 8.44 | 127/0.030 | 238/0.064 | 8.8 | 0.96 | 8.4 | 4.83 | 7.71 | 11.03 | 0.667 | 0.463 | 0.686 |
| Filtered out 10%: | | | | | | | | | | | | | |
| Csáky et al. (2019) SRC | 70,000,000 | 8.59 | 122/0.028 | 222/0.058 | 9.3 | 0.98 | 9.1 | 5.38 | 8.17 | 12.48 | 0.680 | 0.466 | 0.691 |
| Csáky et al. (2019) TRG | 70,000,000 | 16.73 | 194/0.023 | 507/0.064 | 6.0 | 1.00 | 6.0 | 5.63 | 7.25 | 4.11 | 0.699 | 0.440 | 0.683 |
| Junczys-Dowmunt (2018) | 70,000,000 | 8.91 | 126/0.028 | 225/0.057 | 8.9 | 1.00 | 8.9 | 5.12 | 7.68 | 8.55 | 0.673 | 0.466 | 0.688 |
| Ours S_{C+R} | 70,000,000 | 8.43 | 183/0.043 | 403/0.108 | 9.2 | 0.95 | 8.8 | 4.95 | 7.92 | 10.26 | 0.674 | 0.462 | 0.687 |
| Ours S_C | 70,000,000 | 8.60 | 130/0.030 | 231/0.061 | 9.1 | 0.99 | 9.0 | 5.11 | 7.95 | 10.53 | 0.682 | 0.467 | 0.688 |
| Ours S_R | 70,000,000 | 8.42 | 155/0.037 | 306/0.083 | 9.2 | 0.95 | 8.7 | 4.93 | 7.89 | 8.76 | 0.664 | 0.464 | 0.687 |
| Filtered out 50%: | | | | | | | | | | | | | |
| Csáky et al. (2019) SRC | 40,000,000 | 7.97 | 165/0.041 | 329/0.094 | 9.1 | 0.90 | 8.2 | 4.99 | 7.76 | 11.36 | 0.673 | 0.463 | 0.688 |
| Csáky et al. (2019) TRG | 40,000,000 | 18.25 | 213/0.023 | 591/0.069 | 5.4 | 1.00 | 5.4 | 5.15 | 6.86 | 3.33 | 0.701 | 0.453 | 0.682 |
| Junczys-Dowmunt (2018) | 40,000,000 | 8.63 | 206/0.048 | 478/0.125 | 9.4 | 0.98 | 9.2 | 5.16 | 8.32 | 10.25 | 0.668 | 0.463 | 0.688 |
| Ours S_{C+R} | 40,000,000 | 7.13 | 345/0.097 | 853/0.278 | 9.4 | 0.75 | 7.1 | 4.21 | 7.50 | 10.69 | 0.655 | 0.452 | 0.682 |
| Ours S_C | 40,000,000 | 7.31 | 201/0.055 | 466/0.148 | 9.2 | 0.80 | 7.3 | 4.38 | 7.56 | 13.54 | 0.674 | 0.463 | 0.685 |
| Ours S_R | 40,000,000 | 7.91 | 270/0.068 | 662/0.192 | 9.4 | 0.86 | 8.1 | 4.59 | 7.65 | 10.07 | 0.667 | 0.458 | 0.685 |
| reference | | 9.04 | 1301/0.288 | 3244/0.807 | - | - | - | - | - | - | - | - | - |

Table 13: Automatic evaluation results for generated responses (English). BLEU-1 (B1) and its brief penalty (bp), ROUGE (ROU) $\times 100$, METEOR $\times 100$ (MET), CIDEr $\times 100$ (CID). Embedding-based metrics: Embedding Average Cosine Similarity (EA), Vector Extrema Cosine Similarity (VE), Greedy Matching (GM).

E.2 Score Distributions

Figure 5 shows that the score distributions of our S_C , S_R , S_{C+R} across our training data. Note that, for computing S_{C+R} , S_R that less than 0 (pale-colored part in (c)) are treated as 0 (Equation 2).

E.3 Qualitative Analysis for Low Recall

Table 11 shows the samples of utterance pairs that cause low recall property. In the 1st and 2nd pairs, humans can observe the connectivity or the content relatedness, but S_C or S_R failed to provide high scores. For example, “Carcharodontosaurus” and “Spinosaurus” were unknown words for S_R . Other pairs cannot be correctly determined from only these two perspectives. The 3rd and 4th pairs are acceptable, although they have neither the connectivity nor the content relatedness. The 5th pair needs more external contexts and knowledge to determine whether it is acceptable as dialogue.

E.4 Comparison of Top versus Worst Data

Table 12 shows that the comparison of utterance pairs with a high score (i.e., remained as training data; top 50%) and a low score (i.e., removed from training data; worst 50%) in our S_{C+R} . We confirmed there is almost no difference in diversity between the top and the worst ones in terms of pairs, their utterances, and their responses, respectively.

E.5 Automatic Evaluation Results for Generated Responses

Table 13 shows automatic evaluation results for generated responses in our experiment (Section 7).

To calculate these scores, we used publicly available tools.^{27,28}

F Experimental Results on Japanese

F.1 Distributions between Human and Automatic Scoring

Figure 6 shows that, for all the models including ablations, the distributions between human scores and automatically computed scores.

F.2 The Distributions of Proposed Scores

Figure 7 shows that the score distributions of our S_C , S_R , S_{C+R} across our training data. Note that, for computing S_{C+R} , S_R that less than 0 (pale-colored part in (c)) are treated as 0 (Equation 2).

F.3 Comparison of Top versus Worst Data

Table 14 shows that the comparison of utterance pairs with a high score and a low score in our S_{C+R} . We confirmed there is almost no difference in diversity between the top and the worst ones in terms of pairs, their utterances, and their responses, respectively.

F.4 Automatic Evaluation Results for Generated Responses

Table 15 shows automatic evaluation results for generated responses in our experiment (Section 8.2). To calculate these scores, we used publicly available tools.²⁹

²⁷<https://github.com/moses-smt/mosesdecoder>

²⁸<https://github.com/Maluuba/nlg-eval>

²⁹For the embedding-based metrics, we used the pre-trained Japanese word embeddings (Grave et al., 2018).

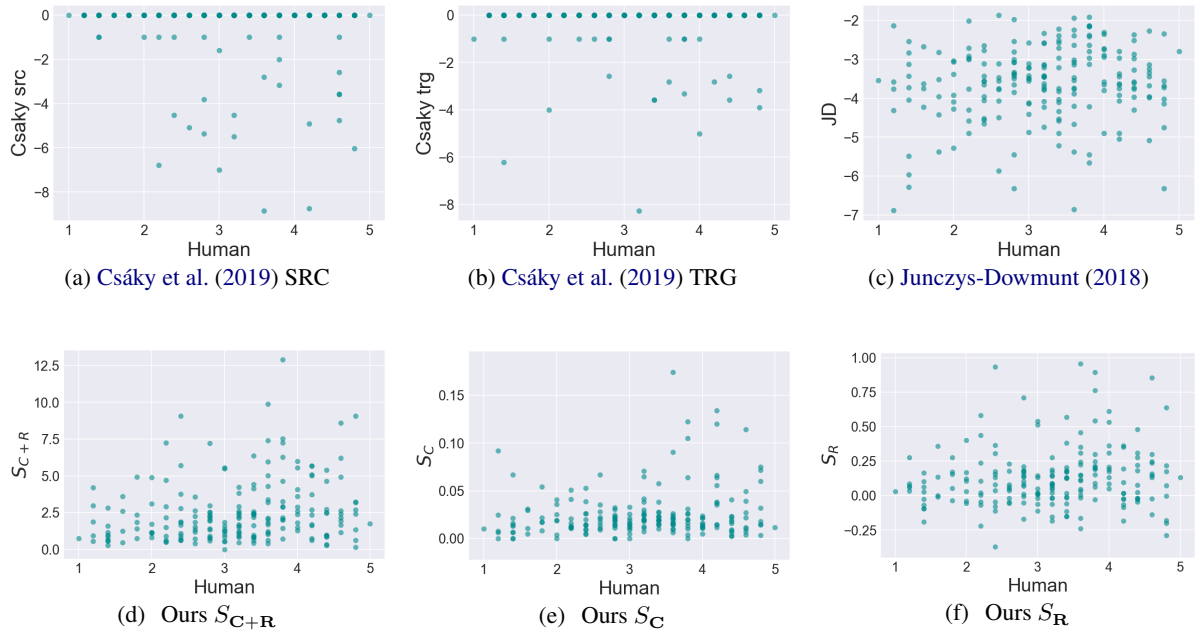


Figure 6: Distributions between human scores and automatically computed scores by each method (Japanese).




Figure 7: Score distributions of our S_C , S_R , S_{C+R} across training data (Japanese).

| Scored data | Utterance pair | | | Utterance (source-side) | | | Response (target-side) | | |
|---------------------|----------------|------------|------------|-------------------------|------------|------------|------------------------|------------|------------|
| | len | distinct-1 | distinct-2 | len | distinct-1 | distinct-2 | len | distinct-1 | distinct-2 |
| Top 50% (remained) | 15.00 | 0.023 | 0.245 | 7.51 | 0.041 | 0.366 | 7.49 | 0.041 | 0.367 |
| Worst 50% (removed) | 13.79 | 0.025 | 0.260 | 6.90 | 0.045 | 0.378 | 6.90 | 0.044 | 0.374 |
| Top 90% (remained) | 14.62 | 0.015 | 0.176 | 7.31 | 0.028 | 0.286 | 7.31 | 0.028 | 0.285 |
| Worst 10% (removed) | 12.71 | 0.091 | 0.639 | 6.40 | 0.146 | 0.851 | 6.31 | 0.147 | 0.852 |

Table 14: Comparison of the top and the worst utterance pairs in the training data scored by our method (Japanese).

| Japanese | # of pairs | len | dist1 | dist2 | B1 | bp | B1bp | MET | ROU | CID | EA | VE | GM |
|--------------------------|------------|------|-----------|------------|------|------|------|------|-------|-------|-------|-------|-------|
| non-filtered | 1,893,477 | 5.91 | 268/0.091 | 509/0.207 | 13.4 | 0.79 | 10.6 | 5.44 | 10.98 | 16.27 | 0.723 | 0.438 | 0.585 |
| Filtered out 10%: | | | | | | | | | | | | | |
| Csáky et al. (2019) SRC | 1,700,000 | 5.75 | 295/0.102 | 550/0.231 | 13.2 | 0.76 | 10.0 | 5.09 | 10.79 | 15.03 | 0.711 | 0.430 | 0.575 |
| Csáky et al. (2019) TRG | 1,700,000 | 7.06 | 336/0.095 | 662/0.219 | 11.6 | 0.96 | 11.1 | 5.75 | 9.91 | 11.23 | 0.730 | 0.434 | 0.581 |
| Junczys-Dowmunt (2018) | 1,700,000 | 5.31 | 284/0.107 | 516/0.240 | 12.6 | 0.68 | 8.5 | 4.87 | 9.84 | 14.89 | 0.711 | 0.441 | 0.574 |
| Ours S_{C+R} | 1,700,000 | 5.68 | 319/0.112 | 582/0.249 | 13.9 | 0.75 | 10.5 | 5.42 | 11.22 | 17.22 | 0.725 | 0.441 | 0.585 |
| Ours S_C | 1,700,000 | 5.51 | 264/0.096 | 492/0.218 | 13.7 | 0.72 | 9.8 | 5.28 | 10.74 | 15.43 | 0.724 | 0.447 | 0.586 |
| Ours S_R | 1,700,000 | 5.73 | 296/0.103 | 555/0.234 | 12.5 | 0.76 | 9.5 | 5.20 | 9.85 | 12.82 | 0.719 | 0.441 | 0.579 |
| Filtered out 50%: | | | | | | | | | | | | | |
| Csáky et al. (2019) SRC | 1,000,000 | 5.93 | 355/0.120 | 651/0.264 | 11.4 | 0.80 | 9.1 | 5.02 | 9.84 | 13.71 | 0.719 | 0.438 | 0.574 |
| Csáky et al. (2019) TRG | 1,000,000 | 6.94 | 405/0.117 | 811/0.273 | 12.2 | 0.95 | 11.5 | 5.89 | 10.77 | 13.77 | 0.719 | 0.420 | 0.574 |
| Junczys-Dowmunt (2018) | 1,000,000 | 5.99 | 421/0.140 | 802/0.321 | 11.2 | 0.79 | 8.9 | 5.02 | 9.25 | 16.41 | 0.706 | 0.421 | 0.561 |
| Ours S_{C+R} | 1,000,000 | 5.53 | 405/0.146 | 741/0.327 | 12.4 | 0.72 | 9.0 | 4.95 | 9.20 | 13.58 | 0.707 | 0.428 | 0.565 |
| Ours S_C | 1,000,000 | 5.48 | 318/0.116 | 599/0.267 | 11.9 | 0.71 | 8.5 | 4.94 | 9.14 | 16.25 | 0.714 | 0.429 | 0.570 |
| Ours S_R | 1,000,000 | 5.76 | 404/0.140 | 747/0.314 | 12.7 | 0.76 | 9.6 | 5.48 | 10.34 | 18.84 | 0.711 | 0.426 | 0.569 |
| reference | | 7.29 | 750/0.206 | 1446/0.460 | - | - | - | - | - | - | - | - | - |

Table 15: Automatic evaluation results for generated responses (Japanese). BLEU-1 (B1) and its brief penalty (bp), ROUGE (ROU) \times 100, METEOR \times 100 (MET), CIDEr \times 100 (CID). Embedding-based metrics: Embedding Average Cosine Similarity (EA), Vector Extrema Cosine Similarity (VE), Greedy Matching (GM).