

# Robust ASR Error Correction with Conservative Data Filtering

Takuma Udagawa, Masayuki Suzuki, Masayasu Muraoka, Gakuto Kurata  
IBM Research AI

takuma.udagawa@ibm.com, {szuk, mmuraoka, gakuto}@jp.ibm.com

## Abstract

Error correction (EC) based on large language models is an emerging technology to enhance the performance of automatic speech recognition (ASR) systems. Generally, training data for EC are collected by automatically pairing a large set of ASR hypotheses (as sources) and their gold references (as targets). However, the quality of such pairs is not guaranteed, and we observed various types of noise which can make the EC models brittle, e.g. inducing overcorrection in out-of-domain (OOD) settings. In this work, we propose two fundamental criteria that EC training data should satisfy: namely, EC targets should (1) improve linguistic acceptability over sources and (2) be inferable from the available context (e.g. source phonemes). Through these criteria, we identify low-quality EC pairs and train the models not to make any correction in such cases, the process we refer to as conservative data filtering. In our experiments, we focus on Japanese ASR using a strong Conformer-CTC as the baseline and finetune Japanese LLMs for EC. Through our evaluation on a suite of 21 internal benchmarks, we demonstrate that our approach can significantly reduce overcorrection and improve both the accuracy and quality of ASR results in the challenging OOD settings.

## 1 Introduction

Automatic speech recognition (ASR) is the task of transcribing human speech into readable text, which is of practical use in various applications. In contrast to the traditional hybrid approach (Sak et al., 2014), modern ASR systems are trained in an end-to-end manner using a large parallel corpus of acoustic speech paired with gold transcriptions (Prabhavalkar et al., 2017; Li et al., 2022). Despite their huge success, end-to-end ASR systems have limited linguistic knowledge due to the difficulty of leveraging unpaired text-only data which exist in abundance (Penedo et al., 2023).

Error correction (EC) is an effective strategy to correct linguistic errors produced by such ASR systems (Errattahi et al., 2018; Guo et al., 2019). Recently, large language models (LLMs) pretrained on massive text-only data have shown promising results for this purpose (Ma et al., 2023a; Chen et al., 2023). While several works explore the zero-shot or in-context learning capability of LLMs (Ma et al., 2023b; Yang et al., 2023a), finetuning LLMs with sufficient EC training data remains critical to impart the knowledge of ASR-specific error patterns and desired corrections (Mani et al., 2020; Leng et al., 2021; Radhakrishnan et al., 2023; Wang et al., 2023; Chen et al., 2024)

Generally, training data for EC are collected by automatically pairing the ASR hypothesis (source) and its gold transcription (target), and the task is formulated as sequence transduction from the source to target (Guo et al., 2019). However, the quality of such pairs is not guaranteed: in fact, we observed various types of noise which require incorrect, unnecessary, or uninferable corrections that are unreasonable to be predicted from the source. We show some illustrative examples in Table 1.

Training EC models on such noisy data can amplify *overcorrection*, which is a typical problem in current EC (Ma et al., 2023b; Leng et al., 2023). However, existing works largely overlook the existence of such noise and apply minimal data filtering, e.g. simply discard pairs with large edit distance (Hrinchuk et al., 2020; Zhao et al., 2021).

In this study, we propose two fundamental criteria that EC training data should satisfy in general. Specifically, we ensure that EC targets

- C1. improve linguistic acceptability over sources
- C2. are inferable from the available context (e.g. source phonemes)

Based on these criteria, we identify low-quality EC pairs and train the models to avoid making any correction on them. Such conservative behavior is

	ASR Hypothesis (Source $W^S$ )	Gold Reference (Target $W^T$ )
Clean	に雑音を蒸したもの [ni zatsuon o fuka shita mono] (to which noise is <i>steamed</i> )	に雑音を付加したもの [ni zatsuon o fuka shita mono] (to which noise is <i>added</i> )
	しかし一対一の場合ですと [shikashi ittaiichi no baai desuto] (but in case of <i>one-to-one</i> )	しかし一対一場合ですと [shikashi ittaiichi baai desuto] (but in case <i>one-to-one</i> )
Noisy	男の人はぐらいですかね [otokonohito wa gurai desukane] (would a male person be <i>about</i> )	男の人の方がいいですかね [otokonohito noho:ga ii desukane] (would a male person be <i>better</i> )
	Incorrect/ Unnecessary	

Table 1: Clean and noisy examples observed in our Japanese EC training data. Phonemes are shown in square brackets [] and English translation in round brackets (). Targets can be naturally inferred from the erroneous sources in the clean cases, while incorrect, unnecessary, or uninferable corrections are required in the noisy cases.

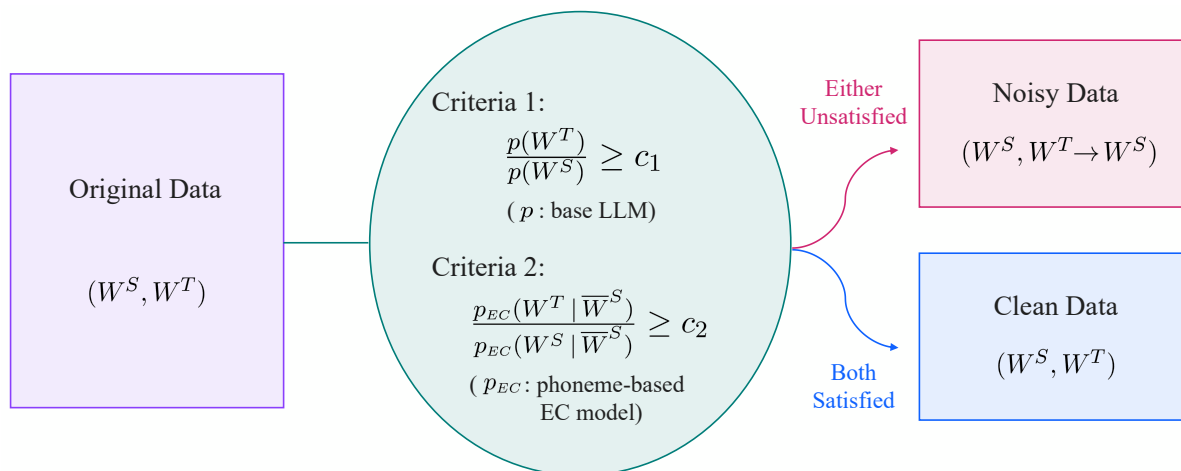


Figure 1: An illustration of our conservative data filtering. Precise details and terminologies are explained in §3.

often crucial to reduce overcorrection and improve robustness, esp. in the out-of-domain (OOD) settings (Li et al., 2024). The overall flow of our data filtering strategy is shown in Figure 1.

In our experiments, we focus on Japanese ASR using an internal Conformer-CTC as the baseline (Lee and Watanabe, 2021). For EC, we fine-tune open-source Japanese LLMs, namely Swallow-Mistral 7B<sup>1</sup> and Sarashina-2 7B<sup>2</sup>, and evaluate the performance on 21 internal benchmarks comprised of various domains. Through our experiments, we confirm that our approach can significantly reduce overcorrection and robustly improve ASR results in the most challenging OOD settings.

## 2 Related Work

In the existing literature, EC primarily focuses on the *in-domain* setup where models are trained and

evaluated over the same domain (Guo et al., 2019; Mani et al., 2020; Wang et al., 2020; Leng et al., 2021; Ma et al., 2023a). Recently, Li et al. (2024) proposed a *low-resource OOD* setup where EC models are finetuned on a limited amount of target domain data to generalize beyond in-domain data. However, target domains of EC are conceptually broad or even open-ended, so it is desirable that EC models work reliably in any target domain without prior knowledge or finetuning. In this study, we focus on the most challenging *zero-resource OOD* setup to develop general-purpose EC models which work out of the box in a variety of domains.

Despite the recent progress, overcorrection remains a major challenge in EC, esp. in the OOD setup. To alleviate this issue, constrained decoding (Zhao et al., 2021; Ma et al., 2023a,b) restricts or biases the correction towards retaining the original ASR hypotheses. Li et al. (2024) use a representative data source and partially train the models to copy the input to induce conservative behavior. In complementary to their approach, we focus on the

<sup>1</sup><https://huggingface.co/tokyotech-llm/Swallow-MS-7b-v0.1>

<sup>2</sup><https://huggingface.co/sbintuitions/sarashina2-7b>

quality of EC data and apply sophisticated data filtering, which is a novel aspect of our approach that works much more effectively than existing filtering based on edit distance (Hrinchuk et al., 2020; Zhao et al., 2021; Ma et al., 2023a).

Typically, ASR errors originate from confusing phonetically similar words and phrases. Therefore, supplementing EC models with phonetic/acoustic information can help improve their performance (Wang et al., 2020; Dutta et al., 2022; Higuchi et al., 2023; Li et al., 2024). In this study, we use the source phonemes as an additional input, which can be easily handled by the text-based LLMs. When available, the full N-best hypotheses can be used as input to provide richer clues on where the ASR systems are confused (Zhu et al., 2021; Ma et al., 2023a). However, for both simplicity and computational efficiency, we only use the 1-best hypothesis (i.e. top ASR prediction) in our experiments.

### 3 Methods

EC can be formulated as a sequence transduction task from the ASR hypothesis (source) to the gold reference (target). Formally, let  $(W^S, W^T)$  denote the source and target sequence pair. In the simplest setting, the EC model is trained to estimate  $p_{EC}(W^T | W^S)$  with the expectation of transforming an error-prone source into a clean target.

In this study, we also incorporate the source phonemes  $\overline{W}^S$  as an additional input, in which case the EC model estimates  $p_{EC}(W^T | W^S, \overline{W}^S)$ . Source phonemes are obtained from our ASR system (§4) and represented in hiragana, one of the Japanese syllabaries, which can be easily consumed by Japanese LLMs. Below is an example:

- $W^S$ : また海属性に関しては  
[mata kai zokuse: ni kanshitewa]  
(also in terms of sea attribute)
- $\overline{W}^S$ : またかいぞくせーにかんしてわ
- $W^T$ : また下位属性に関しては  
[mata kai zokuse: ni kanshitewa]  
(also in terms of subordinate attribute)

Generally, training data of EC can be collected at scale by automatically pairing the hypotheses and gold references in the ASR system’s training data.<sup>3</sup> However, not all source-target pairs are suitable

<sup>3</sup>Although the ASR systems are directly trained on these datasets, they usually make sufficient errors for EC models to learn from. One can virtually increase the amount of errors through noise injection (Zhao et al., 2021) or data partitioning to avoid training on each partition (Hrinchuk et al., 2020).

for training EC models, as we observed various types of noise (illustrated in Table 1). To address this issue, we propose two fundamental criteria that high-quality EC pairs should satisfy.

**Criteria 1: EC targets should improve linguistic acceptability over sources.** The main objective of EC is to resolve linguistic errors in the ASR system’s predictions and improve linguistic acceptability. While the gold reference usually contains cleaner text, this is not always the case, e.g. due to speaker disfluency in spontaneous speech or noisy transcriptions. In addition, Japanese is a language with rich orthographic variation where multiple valid spellings exist (Ohsugi et al., 2022; Karita et al., 2023). For instance, the correction is not necessary if the source transcribes a *bottle* as 瓶 [bin] while the target transcribes as ビン [bin], since both spellings are equally acceptable.

To improve robustness, EC models should only focus on apparent mistakes and resolve them accurately. One simple way to express this criteria is based on the following equation:

$$\frac{p(W^T)}{p(W^S)} \geq c_1 \quad (1)$$

Here,  $p(W^S)$  and  $p(W^T)$  denote the likelihoods of the source and target, which can be computed using any language model. In this study, we simply use the base Japanese LLM.  $c_1$  denotes the threshold, set to 1 by default, which can control the strength of the filter. Intuitively,  $(W^S, W^T)$  that do not satisfy this criteria do not sufficiently improve the linguistic acceptability, indicating the correction is incorrect or unnecessary.

**Criteria 2: EC targets should be inferable from the available context.** Existing works assume that EC targets are generally inferable from the source. However, this is not always the case: in fact, expert evaluation revealed that about one-third of the errors cannot be corrected from the source alone (Zhao et al., 2021). This is mainly attributed to the large phonetic discrepancy between the source and target, e.g. caused by environmental noise or incapability of the ASR system.

A robust EC model should only make the correction when it is inferable from the available context. To express this criteria, we quantify the degree of inferability from the source phonemes using the




Figure 2: Log-likelihood ratios for the two criteria, i.e.  $\log \frac{p(W^T)}{p(W^S)}$  for C1 and  $\log \frac{p_{EC}(W^T | \bar{W}^S)}{p_{EC}(W^S | \bar{W}^S)}$  for C2. Red line shows the default threshold ( $c_1 = c_2 = 1$ ).

following equation:

$$\frac{p_{EC}(W^T | \bar{W}^S)}{p_{EC}(W^S | \bar{W}^S)} \geq c_2 \quad (2)$$

Here,  $p_{EC}$  is a baseline EC model trained only using the source phonemes as input. In this study, we finetune the base Japanese LLM following the procedure described in §4. Again,  $c_2$  denotes the threshold which can be set to 1 by default. Intuitively,  $(W^S, W^T)$  that do not satisfy this criteria cannot be easily inferred from the available context, namely source phonemes in our case.

It is worth noting that edit distance is not a suitable measure of inferability. For instance, the uninferable example in Table 1 has a relatively small edit distance but is very difficult to be inferred. In contrast, the following example is quite dissimilar in terms of edit distance but can be more naturally inferred from the source phonemes:

- $W^S$ : そうか検出で [so:ka kensyutsu de]  
(based on the I see detection)
- $W^T$ : 相関係数で [so:kan ke:suu de]  
(based on the correlation coefficient)

Based on the above criteria (C1 and C2), we identify low-quality EC pairs and train the models to avoid making any correction on them by simply replacing the target with source ( $W^T \rightarrow W^S$ ): see Figure 1 for an illustration. We found this approach more effective than discarding the noisy pairs, since the model is explicitly trained to be conservative on noisy or otherwise ambiguous examples.

Note that both criteria are defined based on the likelihood ratio between the source and target (eq.

1, 2). In Figure 2, we plot the distribution of the (log-)likelihood ratio for each criteria in our training data, using Swallow-Mistral 7B as the LLM.<sup>4</sup> We can verify that a non-negligible portion of the pairs do not satisfy the criteria, suggesting noisy pairs are prevalent in EC training data. For additional examples of the filtered/non-filtered pairs, we refer the reader to Appendix A.

Out of the whole training data, our ASR baseline predicts the exact gold reference (i.e.  $W^S = W^T$ ) in about 34% of the cases. Therefore, the EC model effectively learns to make a correction (i.e.  $W^S \neq W^T$ ) in only 66% of the cases. Of these effective pairs, 34% are classified as noisy based on our C1 filter, 33% based on C2 filter, and 42% when combined. While this results in even fewer examples to learn from, we can expect the model to focus on clearer errors and improve OOD robustness. Our approach is also in line with the principle that data quality can be more important than quantity for LLM alignment (Zhou et al., 2023).

## 4 Experimental Setup

**ASR System** For the ASR baseline, we use an internal Conformer-CTC developed for commercial use cases. The acoustic model is a CTC (Graves et al., 2006) with 240-dimensional logmel-derived features every 40 milliseconds as input, consisting of 10 conformer layers (Gulati et al., 2020), followed by an output layer of 42 Japanese phonemes including the blank symbol. For inference, a static graph for graph decoding is created using a word n-gram model and a dictionary representing the mapping between words ( $W^S$ ) and their phoneme sequences ( $\bar{W}^S$ ). In total, our training data consists of 8000 hours of transcribed speech with little or no overlap between our benchmarking domains.

**EC Model** For EC, we finetune two Japanese LLMs, namely Swallow-Mistral 7B (Fujii et al., 2024; Okazaki et al., 2024) and a more recent Sarashina-2 7B, on a subset of the ASR training data ensuring a 1:1 mixture of read and spontaneous speech. For finetuning, we use LoRA (Hu et al., 2021) with rank  $r = 32$  and scaling factor  $\alpha = 16$ . The effective batch size is set to 128 on a single A100 GPU, and the learning rate is  $5e-4$  annealed with a cosine scheduler. All EC models are trained for a total of 1000 steps, since we observed more steps led to overfitting in the OOD setup.

<sup>4</sup>Statistics based on Sarashina-2 7B are provided in Appendix C, where we observed similar results.



Test	Orig.	No Filter			C1 Only			C2 Only			C1+C2			Inv. C1+C2		
	CER	CER	%EC	%LA	CER	%EC	%LA	CER	%EC	%LA	CER	%EC	%LA	CER	%EC	%LA
1	<b>6.66</b>	9.28	38.5	62.7	7.67	14.9	57.7	8.14	21.8	55.3	7.97	16.7	69.0	8.01	9.2	43.8
2	8.18	7.65	55.4	60.6	<b>7.13</b>	32.6	71.7	7.42	27.8	63.9	7.44	21.4	72.6	8.49	11.0	35.8
3	20.66	21.55	42.7	48.1	20.46	20.5	59.9	20.65	21.8	49.7	<b>20.15</b>	13.1	60.9	20.85	7.6	46.4
4	<b>18.74</b>	21.18	26.1	57.1	19.56	13.8	66.7	20.10	12.9	55.6	19.17	7.5	57.7	20.78	8.6	50.0
5	6.13	7.50	25.0	75.4	7.04	16.4	87.3	7.05	18.8	81.8	7.08	14.3	87.0	<b>5.96</b>	3.6	33.3
6	7.20	<b>6.89</b>	15.4	50.0	<b>6.89</b>	12.8	60.0	7.10	2.6	100.0	<b>6.89</b>	5.1	50.0	7.20	0.0	-
7	<b>12.50</b>	14.26	57.6	55.7	13.38	31.9	62.8	12.92	21.3	57.1	12.81	14.8	62.4	12.89	13.4	38.3
8	8.53	8.67	49.8	57.1	<b>8.39</b>	26.0	64.2	8.54	23.8	58.8	8.45	14.7	68.3	8.62	11.5	44.7
9	8.47	7.82	35.7	51.5	<b>6.91</b>	22.6	63.4	7.16	23.3	53.4	7.16	16.4	66.8	7.98	10.7	49.1
10	8.45	8.06	29.7	62.6	<b>7.34</b>	17.4	67.0	7.86	14.9	64.1	7.67	12.2	65.0	8.52	4.8	51.9
11	19.77	21.00	47.4	59.3	22.41	22.0	67.5	19.89	19.8	58.3	<b>19.70</b>	14.3	59.6	20.35	11.6	52.4
12	12.02	12.08	46.8	59.0	<b>11.34</b>	23.3	59.6	11.72	20.9	51.0	12.24	11.4	64.3	12.40	17.4	43.5
13	13.06	<b>12.64</b>	31.9	53.0	12.83	10.9	61.8	12.83	9.6	70.0	12.95	5.4	35.3	12.91	3.5	54.5
14	<b>26.10</b>	27.88	48.5	54.9	26.86	18.9	65.9	26.81	14.6	50.0	26.79	11.6	81.5	26.45	13.7	40.6
15	15.23	16.36	47.4	51.4	15.34	22.5	64.2	15.47	17.8	53.7	<b>15.20</b>	12.2	55.4	15.88	13.9	43.2
16	12.03	14.08	54.7	62.2	11.36	22.7	73.5	<b>10.74</b>	24.7	75.7	11.65	16.7	64.0	12.32	2.0	0.0
17	9.98	9.53	28.0	61.6	9.35	18.2	66.5	9.52	15.5	62.6	<b>9.31</b>	12.5	65.7	9.92	4.6	56.2
18	14.52	16.81	56.0	50.0	15.70	34.0	52.9	14.28	22.0	45.5	<b>14.03</b>	10.0	60.0	15.02	8.0	50.0
19	6.89	5.84	56.0	75.0	5.76	40.0	85.0	6.17	34.0	82.3	<b>5.69</b>	22.0	90.9	6.66	4.0	100.0
20	7.81	7.45	69.6	73.4	<b>7.33</b>	32.5	83.7	7.61	22.0	81.9	7.66	27.8	86.7	7.83	5.0	52.6
21	5.76	6.20	41.4	53.1	5.91	20.6	58.2	5.76	20.6	56.3	<b>5.39</b>	11.6	67.2	5.64	7.2	52.8
Avg.	11.84	12.51	43.0	58.8	11.86	22.6	66.7	11.80	19.5	63.2	<b>11.69</b>	13.9	66.2	12.13	8.2	47.0
<Orig.	-	38.1	-	-	57.1	-	-	57.1	-	-	<b>71.4</b>	-	-	28.6	-	-

Table 2: Results of EC using Swallow-Mistral 7B. Based on 21 internal test sets, we compute the macro average score for each metric (Avg.) and the ratio of test sets where CER is improved over the original ASR (<Orig.).

For inference, we use greedy decoding, which we found to be efficient yet effective.

For C2 filtering, we train the phoneme-based EC model to predict the target  $W^T$  only using the source phonemes  $\bar{W}^S$  as input. Otherwise, models are trained with both the source phonemes  $\bar{W}^S$  and the source hypothesis  $W^S$  as input.

As an ablation study, we compare the performance of EC models without any data filtering (No Filter), with C1 and C2 filtering applied independently (C1/C2 Only), and with both filtering applied in combination (C1+C2). In addition, to confirm that noisy pairs are less effective for EC training, we also experiment with an inverse filtering of C1+C2, considering the noisy pairs as clean and vice versa (Inv. C1+C2).

**Evaluation** We evaluate EC performance on 21 internal benchmarks comprised of various domains. Details of each benchmark are provided in Table 5. All EC models are evaluated in the zero-resource OOD setup without any domain adaptation.

As for the evaluation metrics, we primarily focus on character error rate ( $\text{CER}^\downarrow$ ) which is standardly used for ASR. To quantify the degree of overcorrection, we also measure the percentage of source hypotheses altered after EC ( $\text{\%EC}^\downarrow$ ). Fi-

nally, we measure the percentage of hypotheses where the linguistic acceptability is improved after EC ( $\text{\%LA}^\uparrow$ ). To measure  $\text{\%LA}$ , we compare the masked language modeling score (Salazar et al., 2020) of the hypothesis before and after EC using Japanese DeBERTa V2 large<sup>5</sup>. While DeBERTa is relatively small compared to recent LLMs, it can take into account the full (bidirectional) context of the hypothesis and effectively assess its linguistic acceptability (Udagawa et al., 2022).

## 5 Results and Discussion

In Table 2, we report the results of our experiments using Swallow-Mistral 7B as the Japanese LLM. Results based on the recent Sarashina-2 7B are provided in Appendix C, where we observed similar trends with even better performance.

Focusing on Swallow-Mistral 7B, there is no single approach which outperforms all others due to the diversity of the test sets. However, we can still draw several conclusions from the overall metrics. First, compared to the original ASR results, we can verify that EC without data filtering drastically worsens CER on average (11.84  $\rightarrow$  12.51). This is mainly attributed to the overcorrection problem, as

<sup>5</sup><https://huggingface.co/ku-nlp/deberta-v2-large-japanese>

we can see a large portion of the hypotheses (43.0% on average) are altered by EC. Such aggressive behavior can be helpful in some occasions (e.g. Test 13) but generally too risky in the OOD setup, leading to modest or even severe performance degradation (e.g. in Test 1, 4, 16, 18, to count a few).

In contrast, by applying our C1 filtering, we can substantially alleviate the degradation of CER (11.84  $\rightarrow$  11.86) by almost halving the frequency of corrections (22.1% on average). This shows that EC can be kept more accurate and conservative by training on cleaner pairs which improve linguistic acceptability. Our C2 filtering also has a similar benefit and makes the EC model more robust in the OOD setup, outperforming the original ASR results in 57.1% (12/21) of the test cases.

In addition, by combining our C1+C2 filtering, we can further cut down the frequency of corrections to 13.9% on average. Through this conservative behavior, we could significantly improve the OOD robustness of EC and reduce the original CER in 71.4% (15/21) of the test sets. This result demonstrates that both C1 and C2 filters help EC focus on clear and fixable ASR errors whilst ignoring more controversial ones.

To verify that clean (rather than noisy) portions of the data contribute to this improvement, we also experimented with the inverse filtering of C1+C2. Generally, we confirmed that inverse filtering worsens CER on average (11.84  $\rightarrow$  12.13) and only improves upon the original ASR in 28.6% (6/21) of the test sets. Therefore, noisy pairs are much less effective for accurate EC. While the frequency of correction is drastically suppressed (8.2% on average), this is largely attributed to the difficulty of learning from noisy examples and overlooking clear errors. In a few cases (e.g. Test 5), we found inverse filtering to be quite competitive, which suggests that noisy pairs still include useful examples for some domains. We expect that our filtering can be improved for such domains by appropriately tuning the thresholds (e.g. lowering  $c_1$  and  $c_2$ ) to include useful pairs of borderline quality.

Finally, in terms of the linguistic acceptability (%LA), we generally see improvement through EC: this indicates that EC is at least successful in resolving linguistic errors and improving ASR quality, even if by deviating from the ground truth (Zhao et al., 2021). Naturally, our C1 filtering consistently strengthens this desirable property by explicitly taking this criteria into account (eq. 1).

Test	Edit Dist. (0.5)			Edit Dist. (0.25)		
	CER	%EC	%LA	CER	%EC	%LA
Avg.	12.65	42.5	59.3	12.85	42.6	57.9
<Orig.	42.9	-	-	33.3	-	-

Table 3: Results of EC using Swallow-Mistral 7B with data filtering based on maximum edit distance.

As an additional experiment, we also evaluated the results of EC with data filtering based on maximum edit distance (Hrinchuk et al., 2020; Zhao et al., 2021; Ma et al., 2023a). In this approach, EC pairs with normalized edit distance above a certain threshold are simply discarded from the training data.<sup>6</sup> We chose the commonly used thresholds of 0.5 and 0.25, which discard 1% and 5% of the whole training data, respectively.

The results are shown in Table 3. We can confirm that filtering based on edit distance fails to improve CER and hardly reduces %EC. This demonstrates that such simple filtering is insufficient to improve the robustness of EC in the challenging OOD setup, regardless of its widespread usage.

Finally, to verify that our claims hold for a different Japanese LLM, we also experimented using Sarashina-2 7B. As discussed in Appendix C, we can draw similar conclusions with even better performance, achieving an average CER of 11.41 in the best case and outperforming the original ASR results in 85.7% (18/21) of the test sets. Therefore, our approach is generalizable using other LLMs and we can expect to further improve performance by leveraging more powerful LLMs.

## 6 Conclusion

EC is an emerging technology to boost the performance of ASR by harnessing the power of LLMs. However, current EC remains brittle, often degrading performance due to overcorrection in the OOD setup, which hinders its practical application.

In this study, we first focused on the quality of EC training data and proposed a method to identify noisy data based on two fundamental criteria. Second, we revealed that EC data contains a considerable proportion of such noisy pairs, which can be effectively handled through our conservative data filtering. Finally, we demonstrated that our approach can significantly alleviate the overcorrection problem and improve the robustness of EC in

<sup>6</sup>Before computing edit distance, we normalized source and target texts by converting them into hiragana using *pykakasi*: <https://github.com/miurahr/pykakasi>.

the challenging zero-resource OOD setup.

In contrast to the existing filtering methods (e.g. based on edit distance), we expect the quality of our data filtering to keep improving as the underlying LLMs become more powerful and accurate, which is a notable trend in the current literature. In future work, we also plan to control for other important factors of data quality, such as diversity and representativeness (Suzuki et al., 2023; Yang et al., 2023b), to further improve the robustness of EC. Overall, we expect our approach to be a foundational step towards developing general-purpose EC models applicable in any domain of interest, facilitating the utilization of LLM technology in the real-world scenarios.

## References

- Chen Chen, Yuchen Hu, Chao-Han Huck Yang, Sabato Marco Siniscalchi, Pin-Yu Chen, and EngSiong Chng. 2023. [Hyporadise: An open baseline for generative speech recognition with large language models](#). In [Thirty-seventh Conference on Neural Information Processing Systems Datasets and Benchmarks Track](#).
- Chen Chen, Ruizhe Li, Yuchen Hu, Sabato Marco Siniscalchi, Pin-Yu Chen, EngSiong Chng, and Chao-Han Huck Yang. 2024. [It’s never too late: Fusing acoustic information into large language models for automatic speech recognition](#). In [The Twelfth International Conference on Learning Representations](#).
- Samrat Dutta, Shreyansh Jain, Ayush Maheshwari, Souvik Pal, Ganesh Ramakrishnan, and Preethi Jyothi. 2022. Error correction in asr using sequence-to-sequence models. [arXiv preprint arXiv:2202.01157](#).
- Rahhal Errattahi, Asmaa El Hannani, and Hassan Ouahmane. 2018. Automatic speech recognition errors detection and correction: A review. [Procedia Computer Science](#), 128:32–37.
- Kazuki Fujii, Taishi Nakamura, Mengsay Loem, Hiroki Iida, Masanari Ohi, Kakeru Hattori, Hirai Shota, Sakae Mizuki, Rio Yokota, and Naoaki Okazaki. 2024. [Continual pre-training for cross-lingual LLM adaptation: Enhancing japanese language capabilities](#). In [First Conference on Language Modeling](#).
- Alex Graves, Santiago Fernández, Faustino Gomez, and Jürgen Schmidhuber. 2006. Connectionist temporal classification: labelling unsegmented sequence data with recurrent neural networks. In [Proceedings of the 23rd International Conference on Machine Learning](#), pages 369–376.
- Anmol Gulati, James Qin, Chung-Cheng Chiu, Niki Parmar, Yu Zhang, Jiahui Yu, Wei Han, Shibo Wang, Zhengdong Zhang, Yonghui Wu, and Ruoming Pang. 2020. [Conformer: Convolution-augmented Transformer for Speech Recognition](#). In [Proc. INTERSPEECH 2020](#), pages 5036–5040.
- Jinxi Guo, Tara N Sainath, and Ron J Weiss. 2019. A spelling correction model for end-to-end speech recognition. In [ICASSP 2019-2019 IEEE International Conference on Acoustics, Speech and Signal Processing \(ICASSP\)](#), pages 5651–5655. IEEE.
- Yosuke Higuchi, Tetsuji Ogawa, and Tetsunori Kobayashi. 2023. Harnessing the zero-shot power of instruction-tuned large language model in end-to-end speech recognition. [arXiv preprint arXiv:2309.10524](#).
- Oleksii Hrinchuk, Mariya Popova, and Boris Ginsburg. 2020. Correction of automatic speech recognition with transformer sequence-to-sequence model. In [ICASSP 2020-2020 IEEE International Conference on Acoustics, Speech and Signal Processing \(ICASSP\)](#), pages 7074–7078. IEEE.
- Edward J Hu, Yelong Shen, Phillip Wallis, Zeyuan Allen-Zhu, Yuanzhi Li, Shean Wang, Lu Wang, and Weizhu Chen. 2021. Lora: Low-rank adaptation of large language models. [arXiv preprint arXiv:2106.09685](#).
- Albert Q Jiang, Alexandre Sablayrolles, Arthur Mensch, Chris Bamford, Devendra Singh Chaplot, Diego de las Casas, Florian Bressand, Gianna Lengyel, Guillaume Lample, Lucile Saulnier, et al. 2023. Mistral 7b. [arXiv preprint arXiv:2310.06825](#).
- Shigeki Karita, Richard Sproat, and Haruko Ishikawa. 2023. [Lenient evaluation of Japanese speech recognition: Modeling naturally occurring spelling inconsistency](#). In [Proceedings of the Workshop on Computation and Written Language \(CAWL 2023\)](#), pages 61–70, Toronto, Canada. Association for Computational Linguistics.
- Taku Kudo, Kaoru Yamamoto, and Yuji Matsumoto. 2004. [Applying conditional random fields to Japanese morphological analysis](#). In [Proceedings of the 2004 Conference on Empirical Methods in Natural Language Processing](#), pages 230–237, Barcelona, Spain. Association for Computational Linguistics.
- Jaesong Lee and Shinji Watanabe. 2021. Intermediate loss regularization for ctc-based speech recognition. In [ICASSP 2021-2021 IEEE International Conference on Acoustics, Speech and Signal Processing \(ICASSP\)](#), pages 6224–6228. IEEE.
- Yichong Leng, Xu Tan, Wenjie Liu, Kaitao Song, Rui Wang, Xiang-Yang Li, Tao Qin, Ed Lin, and Tie-Yan Liu. 2023. Softcorrect: Error correction with soft detection for automatic speech recognition. In [Proceedings of the AAAI Conference on Artificial Intelligence](#), volume 37, pages 13034–13042.

- Yichong Leng, Xu Tan, Linchen Zhu, Jin Xu, Renqian Luo, Linquan Liu, Tao Qin, Xiang-Yang Li, Ed Lin, and Tie-Yan Liu. 2021. Fastcorrect: Fast error correction with edit alignment for automatic speech recognition. In *Advances in Neural Information Processing Systems* 35.
- Jinyu Li et al. 2022. Recent advances in end-to-end automatic speech recognition. *APSIPA Transactions on Signal and Information Processing*, 11(1).
- Yuanchao Li, Pinzhen Chen, Peter Bell, and Catherine Lai. 2024. Crossmodal asr error correction with discrete speech units. *arXiv preprint arXiv:2405.16677*.
- Rao Ma, Mark J. F. Gales, Kate M. Knill, and Mengjie Qian. 2023a. N-best T5: Robust ASR Error Correction using Multiple Input Hypotheses and Constrained Decoding Space. In *Proc. INTERSPEECH 2023*, pages 3267–3271.
- Rao Ma, Mengjie Qian, Potsawee Manakul, Mark Gales, and Kate Knill. 2023b. Can generative large language models perform asr error correction? *arXiv preprint arXiv:2307.04172*.
- Anirudh Mani, Shruti Palaskar, Nimshi Venkat Meripo, Sandeep Konam, and Florian Metze. 2020. Asr error correction and domain adaptation using machine translation. In *ICASSP 2020-2020 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 6344–6348. IEEE.
- Yasuhito Ohsugi, Itsumi Saito, Kyosuke Nishida, and Sen Yoshida. 2022. Japanese ASR-Robust Pre-trained Language Model with Pseudo-Error Sentences Generated by Grapheme-Phoneme Conversion. In *Proc. INTERSPEECH 2022*, pages 2688–2692.
- Naoaki Okazaki, Kakeru Hattori, Hirai Shota, Hiroki Iida, Masanari Ohi, Kazuki Fujii, Taishi Nakamura, Mengsay Loem, Rio Yokota, and Sakae Mizuki. 2024. Building a large japanese web corpus for large language models. In *First Conference on Language Modeling*.
- Guilherme Penedo, Quentin Malartic, Daniel Hesslow, Ruxandra Cojocaru, Hamza Alobeidli, Alessandro Cappelli, Baptiste Pannier, Ebtesam Almazrouei, and Julien Launay. 2023. The refinedweb dataset for falcon LLM: Outperforming curated corpora with web data only. In *Thirty-seventh Conference on Neural Information Processing Systems Datasets and Benchmarks Track*.
- Rohit Prabhavalkar, Kanishka Rao, Tara N. Sainath, Bo Li, Leif Johnson, and Navdeep Jaitly. 2017. A Comparison of Sequence-to-Sequence Models for Speech Recognition. In *Proc. INTERSPEECH 2017*, pages 939–943.
- Srijith Radhakrishnan, Chao-Han Yang, Sumeer Khan, Rohit Kumar, Narsis Kiani, David Gomez-Cabrero, and Jesper Tegnér. 2023. Whispering LLaMA: A cross-modal generative error correction framework for speech recognition. In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pages 10007–10016, Singapore. Association for Computational Linguistics.
- Haşim Sak, Andrew Senior, and Françoise Beaufays. 2014. Long short-term memory recurrent neural network architectures for large scale acoustic modeling. In *Proc. INTERSPEECH 2014*, pages 338–342.
- Julian Salazar, Davis Liang, Toan Q. Nguyen, and Katrin Kirchhoff. 2020. Masked language model scoring. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 2699–2712, Online. Association for Computational Linguistics.
- Jun Suzuki, Heiga Zen, and Hideto Kazawa. 2023. Extracting representative subset from extensive text data for training pre-trained language models. *Information Processing & Management*, 60(3):103249.
- Takuma Udagawa, Masayuki Suzuki, Gakuto Kurata, Nobuyasu Itoh, and George Saon. 2022. Effect and Analysis of Large-scale Language Model Rescoring on Competitive ASR Systems. In *Proc. INTERSPEECH 2022*, pages 3919–3923.
- Haoyu Wang, Shuyan Dong, Yue Liu, James Logan, Ashish Kumar Agrawal, and Yang Liu. 2020. ASR Error Correction with Augmented Transformer for Entity Retrieval. In *Proc. INTERSPEECH 2020*, pages 1550–1554.
- Yi-Wei Wang, Ke-Han Lu, and Kuan-Yu Chen. 2023. Hypr: A comprehensive study for asr hypothesis revising with a reference corpus. *arXiv preprint arXiv:2309.09838*.
- Chao-Han Huck Yang, Yile Gu, Yi-Chieh Liu, Shalini Ghosh, Ivan Bulyko, and Andreas Stolcke. 2023a. Generative speech recognition error correction with large language models and task-activating prompting. In *2023 IEEE Automatic Speech Recognition and Understanding Workshop (ASRU)*, pages 1–8. IEEE.
- Zhao Yang, Yuanzhe Zhang, Dianbo Sui, Cao Liu, Jun Zhao, and Kang Liu. 2023b. Representative demonstration selection for in-context learning with two-stage determinantal point process. In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pages 5443–5456, Singapore. Association for Computational Linguistics.
- Yun Zhao, Xuerui Yang, Jinchao Wang, Yongyu Gao, Chao Yan, and Yuanfu Zhou. 2021. BART Based Semantic Correction for Mandarin Automatic Speech Recognition System. In *Proc. INTERSPEECH 2021*, pages 2017–2021.
- Chunting Zhou, Pengfei Liu, Puxin Xu, Srinu Iyer, Jiao Sun, Yuning Mao, Xuezhe Ma, Avia Efrat, Ping Yu, LILI YU, Susan Zhang, Gargi Ghosh,



Mike Lewis, Luke Zettlemoyer, and Omer Levy. 2023. [LIMA: Less is more for alignment](#). In [Thirty-seventh Conference on Neural Information Processing Systems](#).

Linchen Zhu, Wenjie Liu, Linquan Liu, and Edward Lin. 2021. Improving asr error correction using n-best hypotheses. In [2021 IEEE Automatic Speech Recognition and Understanding Workshop \(ASRU\)](#), pages 83–89. IEEE.

## A Additional Data Examples

In Table 4, we show additional examples of the EC pairs filtered/non-filtered based on our criteria.

In the first example, the uncommon noun 被験者人図 (*human subject diagram*) is a transcription error and corrected into a more natural, similar-sounding phrase 被験者の人数 (*number of human subjects*). This is a perfectly valid example of EC and consequently assigned high log-likelihood ratios based on both criteria.

In the second example, the source (ASR hypothesis) is very unnatural and almost incomprehensible, significantly deviating from the target (gold reference). Therefore, while the target is more natural and acceptable, there is no sufficient context to make it inferable and the pair is reasonably rejected based on the second criteria (C2).

In the third example, the beginning of the source A [e:] is a filler in the speech and not included in the target. Such insertion errors are quite common in Conformer-CTC and hence regarded as inferable based on the baseline EC used in C2. However, this is not a linguistic error in a genuine sense and is properly regarded as an unnecessary correction based on the first criteria (C1).

In the last example, the source is a perfectly valid sentence and even more natural than the target with speaker disfluency: たか高く (*high-higher*). Therefore, it is unreasonable to expect an EC model to make such a correction, which can be safely ignored based on both criteria C1 and C2.

## B Benchmark Details

In Table 5, we provide a brief description of the benchmarks used in our experiments. To evaluate ASR from multiple aspects, our test sets encompass a wide range of domains with various difficulties and characteristics, which in turn introduces diverse ASR errors that need to be corrected through EC. While our training data inevitably contains some data similar to the benchmarking domains

(e.g. daily conversation and presentations), we consider the overlap to be sufficiently small to regard all of them as OOD.<sup>7</sup>

## C Experiments based on Sarashina-2 7B

While Swallow-Mistral 7B is a continuously pre-trained model built upon Mistral 7B (Jiang et al., 2023), Sarashina-2 7B is a recently opensourced Japanese LLM pretrained from scratch on a mixture of Japanese and English texts. To verify that our conclusions are generalizable to different LLMs, we also run the whole experimental pipeline (§3-§5) using Sarashina-2 7B.

In Figure 3, we plot the distribution of the log-likelihood ratio for each criteria in our training data based on Sarashina-2 7B. Out of the effective pairs (where  $W^S \neq W^T$ ), 33% are classified as noisy based on the C1 filter, 49% based on C2 filter, and 63% when combined. While the C2 filter removes a larger portion of the data, we generally observe similar trends as Swallow-Mistral 7B.

In Table 7, we report the results of our experiments using Sarashina-2 7B. Similar to Swallow-Mistral 7B, we found that EC without data filtering fails to improve CER on average (11.84  $\rightarrow$  11.84) due to overcorrection. By applying our C1 filtering, we could significantly improve the average CER (11.84  $\rightarrow$  11.41) whilst reducing the frequency of corrections. Our C2 filtering has a similar benefit, and by combining both filters, we could significantly mitigate overcorrection and improve the original CER in nearly all (85.7%; 18/21) of the test cases. As in the case of Swallow-Mistral 7B, we found that inverse filtering generally has a negative effect on EC performance.

In Table 6, we show the results of edit distance based filtering using Sarashina-2 7B. Again, we can confirm that simple filtering is much less effective compared to our sophisticated filtering which takes into account the pair-wise data quality and explicitly induces conservative behavior.

<sup>7</sup>In fact, we confirmed that EC performs much better on in-domain data, i.e. unseen samples from the ASR system’s training data, and keeps improving with more training steps.

ASR Hypothesis (Source $W^S$ )	Gold Reference (Target $W^T$ )	Log-likelihood Ratios	
		C1	C2
被験者人図を表しています [hikensyaninzu o arawashiteimasu] (it shows the <u>human subject diagram</u> )	被験者の人数を表わしています [hikensya no ninzu: o arawashiteimasu] (it shows the <u>number of human subjects</u> )	1.159	0.812
で高校右下ですね [de ko:ko: umoto desune] (and <u>high school lower right</u> )	でこういうモデルです [de ko: ko:yu: moderu desune] (and <u>it's a model like this- this</u> )	0.680	-0.407
A二の抽出方法ですが [e:ni no chu:syutsu ho:ho: desuga] (in terms of the <u>extraction method of A2</u> )	二の抽出方法ですが [ni no chu:syutsu ho:ho: desuga] (in terms of the <u>extraction method of 2</u> )	-0.511	0.174
暖かくなってまいりましてですね [atatakaku natte mairimashite desune] (it is getting <u>warmer</u> )	たか高くなってまいりましてですね [takatakaku natte mairimashite desune] (it is getting <u>high- higher</u> )	-1.187	-2.223

Table 4: Additional examples from the training data, along with their log-likelihood ratios for the two criteria: i.e.  $\log \frac{p(W^T)}{p(W^S)}$  for C1 and  $\log \frac{p_{EC}(W^T | \bar{W}^S)}{p_{EC}(W^S | \bar{W}^S)}$  for C2. Based on the default thresholds ( $c_1 = c_2 = 1$ ), both ratios must be above 0 to be considered clean (cf. §3 for further details).

Test	Domain	# Utterances	Avg. Length
1	business dialogue (spont.)	187	9.34
2	university lecture (spont.)	891	21.16
3	children stories (read)	1649	9.03
4	proper nouns (read)	350	3.86
5	agent-customer interactions (read)	1400	7.78
6	financial-domain dialogue (spont.)	39	13.23
7	presentation (spont.)	1046	16.08
8	presentation (spont.)	448	12.89
9	miscellaneous (read)	1600	8.04
10	daily conversation (read)	1118	7.90
11	interview (spont.)	378	25.79
12	presentation (spont.)	490	23.48
13	customer support (read)	332	13.34
14	financial-domain dialogue (spont.)	291	9.97
15	daily conversation (spont.)	586	10.56
16	addresses (read)	150	10.71
17	miscellaneous (read)	1050	10.39
18	miscellaneous (read)	50	18.10
19	news (read)	50	30.36
20	customer support (spont.)	379	32.72
21	miscellaneous (spont.)	500	12.40

Table 5: Benchmark details. Our test sets encompass a wide range of domains, including both monologues/dialogues and spontaneous/read speech. Average utterance lengths are computed with the Mecab tokenizer (Kudo et al., 2004).

Test	Edit Dist. (0.5)			Edit Dist. (0.25)		
	CER	%EC	%LA	CER	%EC	%LA
Avg.	11.87	40.6	58.8	12.01	39.7	62.5
<Orig.	61.9	-	-	47.6	-	-

Table 6: Results of EC using Sarashina-2 7B with data filtering based on maximum edit distance.

Test	Orig.	No Filter			C1 Only			C2 Only			C1+C2			Inv. C1+C2		
	CER	CER	%EC	%LA	CER	%EC	%LA	CER	%EC	%LA	CER	%EC	%LA	CER	%EC	%LA
1	<b>6.66</b>	9.25	33.9	62.7	7.97	13.8	58.3	7.87	14.4	72.0	7.30	6.3	54.5	7.60	10.3	55.6
2	8.18	8.09	54.4	58.9	<b>7.51</b>	33.7	68.2	7.57	25.5	67.4	7.68	20.2	72.0	8.37	8.5	50.0
3	20.66	20.85	41.8	49.9	<b>20.07</b>	22.0	59.9	20.22	20.3	55.7	20.16	10.8	60.1	20.57	7.5	49.2
4	18.74	<b>17.63</b>	25.3	68.2	18.52	11.2	69.2	18.42	8.9	64.5	18.20	5.2	66.7	18.56	8.3	51.7
5	6.13	7.11	21.4	78.3	6.57	9.8	81.0	6.98	14.0	81.6	6.29	5.1	80.3	<b>5.91</b>	3.9	50.9
6	7.20	<b>6.67</b>	12.8	60.0	<b>6.67</b>	7.7	66.7	6.99	7.7	33.3	6.89	5.1	50.0	6.99	2.6	100.0
7	12.50	13.80	55.3	54.5	12.90	26.6	59.5	12.74	20.5	49.2	<b>12.49</b>	11.4	59.6	12.85	15.1	55.6
8	8.53	8.40	47.1	59.4	<b>7.89</b>	25.0	64.7	8.27	23.8	59.8	8.10	12.8	69.2	8.45	9.8	47.5
9	8.47	7.92	33.6	52.1	<b>7.37</b>	18.9	56.4	7.48	18.2	52.9	7.68	11.4	64.3	8.14	10.1	34.2
10	8.45	7.99	26.6	64.3	<b>7.78</b>	14.7	67.7	7.81	12.0	70.9	7.85	7.1	72.2	8.35	4.7	39.6
11	19.77	19.59	50.1	58.8	19.64	24.8	58.9	<b>19.30</b>	18.2	59.1	19.70	12.4	60.0	21.74	11.6	59.5
12	12.02	11.94	46.0	56.4	11.65	24.1	60.2	<b>11.55</b>	20.4	52.0	12.01	9.6	48.9	11.95	11.4	50.0
13	13.06	12.99	27.5	60.5	13.01	12.5	69.2	<b>12.86</b>	8.0	56.0	12.92	3.8	33.3	13.00	4.8	66.7
14	<b>26.10</b>	28.14	46.8	63.3	26.44	16.7	66.7	26.70	13.3	61.3	26.19	7.3	52.9	26.81	12.9	70.0
15	15.23	16.12	45.9	53.5	15.28	21.2	54.9	15.32	16.9	45.6	<b>15.14</b>	9.0	56.2	15.45	12.9	46.4
16	12.03	9.06	41.3	77.4	<b>8.48</b>	38.7	86.2	9.90	30.7	71.7	10.07	25.3	76.3	11.61	3.3	100.0
17	9.98	9.49	26.5	63.0	9.42	17.2	71.3	<b>9.29</b>	15.1	68.3	9.58	10.6	71.2	9.97	5.2	61.1
18	14.52	14.96	62.0	51.6	14.65	34.0	64.7	14.34	14.0	57.1	<b>14.03</b>	8.0	100.0	14.65	2.0	100.0
19	6.89	<b>5.16</b>	58.0	79.3	5.46	40.0	85.0	6.10	44.0	86.4	5.91	30.0	86.7	6.78	6.0	66.7
20	7.81	7.38	66.7	69.4	<b>7.27</b>	31.2	71.2	7.46	20.9	78.5	7.49	10.3	84.6	7.88	9.0	70.6
21	5.76	6.17	39.6	56.1	<b>4.98</b>	21.0	70.5	5.26	17.6	62.5	5.19	12.4	71.0	5.60	5.8	48.3
Avg.	11.84	11.84	41.1	61.8	<b>11.41</b>	22.1	67.2	11.54	18.3	62.2	11.47	11.1	66.2	11.96	7.9	60.6
<Orig.	-	61.9	-	-	71.4	-	-	76.2	-	-	<b>85.7</b>	-	-	61.9	-	-

Table 7: Results of EC using Sarashina-2 7B. Based on 21 internal test sets, we compute the macro average score for each metric (Avg.) and the ratio of test sets where CER is improved over the original ASR baseline (<Orig.).




Figure 3: Log-likelihood ratio for the two criteria using Sarashina-2 7B. Red line shows the default threshold ( $c_1 = c_2 = 1$ ).