

# ERROR RATE IS NOT ENOUGH: COSTS OF AUGMENTATION IN LOW-RESOURCE ASR

Katsumi Ibaraki      David Chiang

University of Notre Dame  
Dept. of Computer Science and Engineering  
{kibaraki, dchiang}@nd.edu

## ABSTRACT

Data augmentation is standard practice for low-resource automatic speech recognition (ASR), but its benefits are almost always reported as aggregate error rates, leaving open *why* and *when* a method helps. On four low-resource languages and English at five data sizes, we cross two augmentations that modify existing audio (speed perturbation and SpecAugment) with two that modify text transcriptions and synthesize audio from it. We analyze errors at the phoneme level and find that they help for complementary reasons. SpecAugment has the most reliable aggregate improvements, but increases tone and palatalization errors, whereas text-based methods introduce consonant insertions; in combination, each reduces the errors the other introduces. This complementarity is predictable. Tone and palatalization are each signaled by a single acoustic property and occur least frequently in training (lowest training exposure), and text-based methods that flatten word frequency raise their frequency by 37–38%, reversing those errors while further improving aggregate error. The improvement from combining the methods exceeds the sum of their individual improvements; the superadditivity vanishes by 100 hours, limiting the complementarity to the low-resource regime. Where specific error types matter, we recommend selecting augmentation by training exposure and acoustic properties, both measurable before any training.

**Index Terms**— Automatic Speech Recognition (ASR), data augmentation, low-resource languages, language documentation

## 1. INTRODUCTION

For some endangered and under-documented languages, transcribed speech can be measured in minutes rather than thousands of hours. Data augmentation is the standard response, and both modification of existing audio [1, 2] and conversion of text to new audio [3, 4, 5] are reported to help. However, when methods are compared by a single aggregate error rate, it establishes *that* they help without establishing *why* or *when* they help. For language documentation, error *type* matters as much as error rate [6]. For a linguist that studies (say) tone, an augmentation method that lowers overall error while degrading tone marking may be worse than nothing. In this paper, we look more deeply into the types of errors that different augmentation methods make. We find that the method that delivers the best aggregate error rate, SpecAugment [2], actually increases some error types, namely, tone and palatalization. Yet these increases are reversible while still lowering aggregate error; making rare words more frequent in training reverses them, because tone and palatalization sit disproportionately in rare words.

We fine-tune wav2vec2-large-xlsr-53 [7] on four low-resource languages and on LibriSpeech subsets from one to one hundred

hours, crossing two modification-based methods with two text-based methods—to our knowledge, the first per-category account of what augmentation does to a phoneme inventory. We make three contributions:

- While SpecAugment has the most reliable improvements in WER/CER of any method tested, it increases errors in recognizing tone and palatalization. A text-based augmentation, Random Replacement, reverses these errors while further improving WER/CER.
- These error categories are predictable: SpecAugment increases errors on categories that are rare in training and carried by a single acoustic property, because these categories are more easily masked out. Random Replacement, which replaces words uniformly over types, increases the word frequencies of rare categories and reduces errors.
- A frequency-preserving replacement supports the mechanism: Tag-Based Replacement, which leaves category frequencies unchanged, fails to reverse the tone errors—with SpecAugment, it instead significantly *increases* them (+11%).

## 2. RELATED WORK

Standard text augmentation relies on external lexical resources such as WordNet [8, 9] or predefined semantic frames [10, 11], often unavailable for low-resource languages. Audio synthesis via TTS [3, 4, 5] usually assumes unpaired text; we follow Robinson et al. [12] by using a high-resource *pivot* TTS system instead of an in-language TTS model [13]. While Rossenbach et al. [14] found TTS and SpecAugment contribute largely independently, we revisit this under low-resource constraints. SpecAugment [2] and speed perturbation [1] are the *de facto* acoustic standards, as they require no external data; speed perturbation resamples the waveform and therefore shifts F0. The IARPA BABEL program [15, 16] targeted 10–80 hours; we target 1–5 hours, where glossaries and audio sufficient for pseudo-labeling [17, 18] are unavailable. Boundary-aware masking [19] keeps masks inside phoneme boundaries to protect tonal contours, but requires forced alignments our setting lacks.

Across this literature, augmentation is evaluated almost exclusively by aggregate error rate, so the phonological effects each method introduces go unmeasured. Akavarapu et al. [20] relate phoneme accuracy to training frequency across models; augmentation, however, alters the training distribution itself, and its per-phoneme effects remain unexamined.

### 3. METHOD

#### 3.1. Augmentation Methods

We select two methods that modify existing audio and two that generate new text–audio pairs, chosen to differ maximally in what they do to the signal. The text-based methods extend template-based substitution [10, 11] to a setting where no external text resources exist; they draw only on the annotations already present in the corpus, and synthesize the result through TTS.

##### 3.1.1. Text Modification

**Tag-Based Replacement** treats the sequence of annotation tags (interlinear glosses, or POS and dependency tags for English) as a fixed template, replacing each word with one sampled from the training word types with the same tag. It does not guarantee semantic coherence, but it preserves the tag sequence, and with it most of the sentence’s grammatical structure. **Random Replacement**, a more extreme variant, replaces each position with a word sampled from the training vocabulary; all words remain in-vocabulary, but sentences need not be grammatical. The replacement ratio  $r$  (proportion of positions substituted) is tuned per language and method over  $\{0.0, 0.2, 0.4, 0.6, 0.8, 1.0\}$ . At  $r = 0.0$ , it re-synthesizes the original text, testing whether the benefit comes from novel text or merely additional audio. Text is synthesized with Kokoro,<sup>1</sup> an open-weight TTS accepting IPA and orthographic input, using five fixed voices across all languages; per-speaker audio was insufficient for voice cloning, and voice conversion [21] did not help.

##### 3.1.2. Audio Modification

**Speed perturbation** [1] applies warp factors in  $[0.9, 1.1]$ ; application probability was tuned from 10–50%, with 30% selected for all languages. **SpecAugment** [2] applies time and frequency masking to the feature encoder outputs following [22]. We tuned the mask length between 3 and 10 and the masking probability between 5% and 50% by grid search on the validation sets; a length of 5 and a probability of 30% were chosen uniformly for all languages.

#### 3.2. Fine-tuning Wav2Vec2

We fine-tune the pre-trained wav2vec2-large-xlsr-53 model [7] using a Connectionist Temporal Classification (CTC) loss [23]. We use a learning rate of  $1.0 \times 10^{-4}$  and tune training epochs for each language over  $\{5, 10, 15, 20, 30, 40, 45, 50, 55, 60, 80, 100\}$ . To optimize the data augmentation pipeline, we tuned the ratio of original to synthetic training data for each language across 1:1, 1:2, 1:3, 1:4, and 1:5. The selected hyperparameters, optimized on the validation set, are detailed in Table 1. To mitigate tag sparsity in our Tag-Based Replacement, we relaxed the exact-match constraint, falling back to the nearest tag by edit distance so a replacement is always found and the target ratio is met. Model training was conducted on a single NVIDIA A10 GPU, with runtimes ranging from 3 to 160 hours depending on the dataset size and augmentation ratio.

## 4. EXPERIMENT SETUP

#### 4.1. Data

We tested on four low-resource languages, all with five hours or less of audio. We allowed speaker overlap between splits follow-

**Table 1:** Corpus statistics and selected hyperparameters: minutes of transcribed audio, speakers, training word tokens and types, tag types, replacement ratio  $r$  and original:synthetic mix (Tag/Random), and epochs. LibriSpeech (LS) splits are subsets of train-clean-100 [24].

|                       | Min. | Spk. | Train Words |       | Tag<br>Typ. | Repl. $r$<br>(T/R) | Mix<br>(T/R) | Ep. |
|-----------------------|------|------|-------------|-------|-------------|--------------------|--------------|-----|
|                       |      |      | Tok.        | Typ.  |             |                    |              |     |
| Vatlongos             | 286  | 60   | 23315       | 2823  | 422         | .2/6               | 1:3/1:3      | 40  |
| Nashta                | 48   | 4    | 3982        | 886   | 637         | .6/2               | 1:2/1:1      | 60  |
| S. Buryat             | 97   | 4    | 7895        | 2276  | 2242        | .4/2               | 1:3/1:1      | 45  |
| Kakabe                | 99   | 34   | 12512       | 1525  | 259         | .2/1.0             | 1:2/1:3      | 60  |
| English (LibriSpeech) |      |      |             |       |             |                    |              |     |
| LS-1h                 | 60   | 169  | 6967        | 2149  | 383         | .2/4               | 1:1/1:1      | 60  |
| LS-2h                 | 120  | 222  | 13528       | 3568  | 488         | .2/4               | 1:1/1:1      | 50  |
| LS-5h                 | 301  | 249  | 34992       | 6566  | 646         | .2/4               | 1:1/1:1      | 30  |
| LS-20h                | 1207 | 251  | 138790      | 14333 | 888         | .2/4               | 1:1/1:1      | 10  |
| LS-100h               | 6036 | 251  | 693914      | 29598 | 1184        | .2/4               | 1:1/1:1      | 5   |

ing [25, 26], and as the target application—assistive transcription of a field linguist’s unlabeled collection—inherently features such overlap. While the level of transcription detail varies across our data sources, we treat all IPA-based transcriptions as phonemic representations and refer to the units as *phonemes*.

Vatlongos (Vanuatu [27]), Nashta (a South Slavic variety of the Balkans [28]), Shinekhen Buryat (Mongolic, Inner Mongolia [29]) and Kakabe (Mande, Guinea [30]) are drawn from the Pangloss Collection and the Tokyo University of Foreign Studies repository. All four are annotated with IPA transcriptions, translations, and interlinear glosses; an example from Shinekhen Buryat is shown in (1). Kakabe marks tone, Nashta marks stress, and Nashta and Shinekhen Buryat mark palatalization. LibriSpeech is a corpus of read English speech [24], to which we add spaCy [31] POS and syntactic dependency tags, emulating a documentation-style annotation.

- (1) *iimerxuu buleg buleg-ee-ee bai-dag*  
*suchlike group group-INS-REFL be-PTCP.HBT*  
 ‘They lived in groups like this.’

#### 4.2. Evaluation

We report Word Error Rate (WER) with Phoneme Error Rate (PER) for the four IPA-transcribed low-resource corpora and Character Error Rate (CER) for orthographic English. Significance is assessed against the baseline with paired bootstrap resampling [32, 33], 10,000 samples at  $\alpha = 0.05$ ; the same resampling is applied to per-category error counts in Table 3, with a Holm–Bonferroni correction [34] across categories. Because the repair claims of Section 5.3 compare SpecAugment alone with SpecAugment plus a text-based method, rather than either with the baseline, those are additionally tested head-to-head. To distinguish which errors a method introduces, we align each hypothesis to its reference at the phoneme level using Levenshtein distance, and we label every substitution, deletion, and insertion with a phonological category (those of Table 3).

Finally, we rank every individual unit in a corpus (each vowel, consonant, tone, stress, and palatalized segment) by its training token count and report the median rank of each category’s members as a percentile of the full inventory, which we call the *training exposure percentile*. This makes exposure comparable across corpora of different sizes and inventories. We also report raw token counts where magnitude matters.

<sup>1</sup><https://huggingface.co/hexgrad/Kokoro-82M>

**Table 2:** Results (%). \*: significant improvement over the baseline at  $\alpha = .05$ , paired bootstrap resampling [32, 33]; bold is best before rounding. SP = speed perturbation, SA = SpecAugment.

| Method | Vatlongos    |              | Nashta       |              | S. Buryat    |             | Kakabe       |              |
|--------|--------------|--------------|--------------|--------------|--------------|-------------|--------------|--------------|
|        | PER          | WER          | PER          | WER          | PER          | WER         | PER          | WER          |
| Base   | 13.4         | 36.8         | 18.4         | 59.1         | 14.5         | 45.3        | 23.6         | 47.8         |
| + SP   | 13.1         | 36.5         | 18.5         | 60.4         | 14.6         | 46.7        | 23.3         | 46.7         |
| + SA   | 12.3*        | 34.7*        | 17.6*        | 56.0*        | 13.3*        | 44.4        | 21.6*        | 45.7*        |
| Tag    | 12.6         | 34.8         | 18.1         | 56.5         | 14.3         | 45.2        | 24.2         | 47.3         |
| + SP   | 13.0         | 35.9         | 18.7         | 58.9         | 13.9         | 44.2        | 23.9         | 47.9         |
| + SA   | 11.5*        | 31.9*        | 16.5*        | 53.2*        | 13.4*        | 43.5        | 22.3*        | 46.3         |
| Rand.  | 12.9         | 36.0         | 18.6         | 57.6         | 14.5         | 45.2        | 23.5         | 46.5         |
| + SP   | 12.8         | 35.4         | 18.0         | 55.4         | 14.2         | 44.7        | 23.4         | 46.9         |
| + SA   | <b>11.5*</b> | <b>31.6*</b> | <b>16.1*</b> | <b>51.1*</b> | <b>13.2*</b> | <b>42.9</b> | <b>21.2*</b> | <b>43.1*</b> |

  

| Method | English (LibriSpeech) |              |             |              |             |             |             |             |             |             |
|--------|-----------------------|--------------|-------------|--------------|-------------|-------------|-------------|-------------|-------------|-------------|
|        | 1h                    |              | 2h          |              | 5h          |             | 20h         |             | 100h        |             |
| Method | CER                   | WER          | CER         | WER          | CER         | WER         | CER         | WER         | CER         | WER         |
|        | CER                   | WER          | CER         | WER          | CER         | WER         | CER         | WER         | CER         | WER         |
| Base   | 6.5                   | 24.4         | 4.8         | 17.3         | 2.9         | 10.1        | 3.0         | 9.0         | 1.9         | 6.0         |
| + SP   | 6.0*                  | 22.0*        | 4.1*        | 14.9*        | 3.0         | 10.4        | 2.4*        | 7.9*        | 1.9         | 5.9         |
| + SA   | 5.0*                  | 18.7*        | 4.4*        | 16.5*        | 3.0         | 10.9        | 2.0*        | 7.0*        | 1.6*        | 5.5*        |
| Tag    | 4.3*                  | 15.0*        | 3.7*        | 12.6*        | 2.7*        | 8.8*        | 2.0*        | 6.6*        | 1.9         | 5.7         |
| + SP   | 4.4*                  | 15.4*        | 3.7*        | 12.7*        | 2.7*        | 8.9*        | 2.1*        | 6.6*        | 2.1         | 6.3         |
| + SA   | 4.1*                  | 14.9*        | <b>3.4*</b> | <b>11.9*</b> | <b>2.3*</b> | <b>7.8*</b> | <b>1.7*</b> | <b>5.8*</b> | <b>1.4*</b> | <b>4.9*</b> |
| Rand.  | 4.7*                  | 16.4*        | 3.9*        | 13.1*        | 3.0         | 9.7*        | 2.5*        | 7.6*        | 2.0         | 6.0         |
| + SP   | 4.8*                  | 16.8*        | 3.8*        | 12.9*        | 2.9         | 9.6*        | 2.7*        | 8.2*        | 2.0         | 5.9         |
| + SA   | <b>4.0*</b>           | <b>13.9*</b> | 3.5*        | 12.2*        | 2.6*        | 8.7*        | 1.9*        | 6.6*        | 1.6*        | 5.5*        |

## 5. RESULTS AND DISCUSSION

### 5.1. Overall Performance

In Table 2, no text-based method alone reaches significance on either metric across the four low-resource languages; the largest single improvement is 2.6% WER for Tag-Based Replacement on Nashta.  $r = 0.0$  was never selected: re-synthesized audio of the original text never beat novel text, so the value of text-based augmentation lies in the text, not merely in additional audio. Speed perturbation is also not significant in any of its twelve comparisons. SpecAugment is the exception, improving PER and WER in all combinations by an average of 1.5% PER and 2.8% WER absolute; a practitioner testing methods individually against a baseline would adopt it, as the only single method that improves both metrics in every low-resource comparison.

Combinations change the picture. Random Replacement with SpecAugment gives the lowest PER and WER in all four languages, reducing WER by 2.4–8.0% absolute (5.3–14.1% relative). On Nashta, Random Replacement alone improves WER by 1.5% and SpecAugment alone by 3.1%, yet the combination improves it by 8.0%. This super-additivity holds in all four languages, but vanishes by 100 h in LibriSpeech, where the methods contribute largely independently, matching findings of Rossenbach et al. [14].

### 5.2. Phoneme-level Analysis

Analyzing by phonological category shows that while SpecAugment decreases overall errors, it causes some error types to increase. On Kakabe, SpecAugment alone improves both PER and WER significantly and is the strongest single method by either metric, but increases tone errors by 4%; Random Replacement, by itself significant on neither metric, reduces them by 2%. Neither change is significant against the baseline on its own, but adding Random Replacement to SpecAugment reduces tone errors significantly relative

to SpecAugment alone: by 9% ( $p = .043$ ; Section 5.3). Aggregate error rate shows none of this: neither the tone errors SpecAugment adds nor their reversal by Random Replacement.

See Table 3. SpecAugment’s improvements are strongest in vowel and consonant errors (−18%), while the categories in which SpecAugment appreciably increases errors are tone (+4%) and palatalization (+7%). Neither reaches significance on its own, but the tone increase does reach significance in one combination, Tag-Based Replacement with SpecAugment (+11%, 95% CI [+2, +22],  $p = .025$ ), for reasons we return to in Section 5.3.

One possible explanation is **exposure**, or how often these categories occur. The vowel and consonant categories are among the most frequent, while tone and palatalization are the two scarcest categories.

A second possible explanation is what we call **signaling**, or how many different kinds of acoustic cues a category has. When SpecAugment masks out one out of several cues, the model must learn to combine evidence from the others. Vowels and consonants differ in periodicity, duration, and gross spectral shape, so masking one of these leaves the others intact. Stress is typically signaled by several acoustic properties, including duration, intensity, and F0, and improves slightly. By contrast, both tone and palatalization are signaled by essentially one acoustic property: tone by F0, and palatalization by formant transitions, unlike the segmental categories above.

A third possible reason is that tone and palatalization are both non-segmental (that is, they are realized on or across vowels and consonants). But stress is also non-segmental, and it improves under SpecAugment. The difference is that tone and palatalization have low exposure and single-cue signaling, while stress has neither: it sits at the 98th exposure percentile, appearing in 87% of Nashta word types, and is carried by several cues. So exposure and signaling, not segmentality, are the more likely reasons SpecAugment increases tone and palatalization errors.

An alternative explanation of the increase in tone errors would be mask size: a mask longer than the tone-bearing segment hides the whole tone, so longer masks should hurt tone more. To test this, we swept the mask length on Kakabe from 60 to 200 ms. Aggregate performance peaked at 100 ms (PER 22.7, 21.6, 22.2, 22.5), but tone errors increased at every width and did not grow with width (+9%, +5%, +1%, +5%). So mask size does not explain the increase.

Speed perturbation reduces insertions (−11% vowels, −6% consonants) but increases deletions and substitutions, producing a near-zero net effect. The errors it introduces fall on the dimensions it manipulates—it shifts F0, which carries tone, and it warps duration, affecting stress and consonants.

### 5.3. Method Complementarity

The combinations’ advantage runs in two directions: each method repairs the errors the other introduces. Random Replacement raises consonant insertions by 18% while leaving deletions largely unchanged; the model outputs too many symbols, and its total labeled errors rise significantly (+4%,  $p = .004$ )—yet it forms the best-performing combination. One possible explanation is that in synthetic audio, every frame is a clean phoneme or silence—synthetic speech measurably differs in distribution from real recordings [35]—so the model learns no mapping from real-world non-speech frames (breath, noise) to blank, and inserts symbols instead. SpecAugment fixes this. Masking contiguous regions prevents the model from depending on any single local feature and pushes it to integrate longer context [2, 36], suppressing exactly these locally-triggered inser-

**Table 3:** Relative change (%) in error count vs. the baseline; negative is improvement. Segmental rows pool the four languages; tone is Kakabe-only, stress Nashta-only, palatalization Nashta and Shinekhen Buryat. Tone, stress, and palatalization rows pool substitution (sub), insertion (ins), and deletion (del) channels; a mark on an inserted or deleted segment counts in both its segmental and non-segmental row, so rows overlap and Total sums the listed channels. **Exp.** is the median training exposure percentile (Section 4.2); rows are ordered by it, with the rule separating high from low. **Bold:**  $p < .05$ , paired bootstrap; †: survives Holm correction across categories. SA = SpecAugment, SP = speed perturbation.

| Error Type           | Kind           | Exp. | Errors | SA           | SP           | Tag  |              |     | Rand         |              |     |
|----------------------|----------------|------|--------|--------------|--------------|------|--------------|-----|--------------|--------------|-----|
|                      |                |      |        |              |              | Base | +SA          | +SP | Base         | +SA          | +SP |
| stress               | suprasegmental | 98   | 79     | -5           | +15          | +10  | +22          | +15 | +6           | +4           | +3  |
| sub: vowel           | segmental      | 89   | 981    | -6           | +2           | +4   | -6           | +1  | +0           | <b>-10</b> † | -4  |
| ins: vowel           | segmental      | 89   | 961    | <b>-14</b> † | <b>-11</b> † | -7   | <b>-15</b> † | -4  | +4           | <b>-10</b> † | -0  |
| del: vowel           | segmental      | 89   | 870    | +1           | +5           | +6   | -2           | +5  | +1           | -4           | +3  |
| sub: consonant       | segmental      | 51   | 762    | <b>-15</b> † | +1           | +3   | <b>-11</b> † | -1  | +2           | <b>-17</b> † | -0  |
| ins: consonant       | segmental      | 51   | 865    | <b>-10</b>   | -6           | +3   | <b>-13</b> † | +3  | <b>+18</b> † | <b>-11</b> † | -0  |
| del: consonant       | segmental      | 51   | 843    | -7           | +4           | -0   | <b>-11</b> † | +1  | +1           | <b>-14</b> † | -1  |
| sub: vowel↔consonant | segmental      | —    | 673    | <b>-18</b> † | +5           | -1   | <b>-21</b> † | +3  | +4           | <b>-15</b> † | +1  |
| tone                 | suprasegmental | 26   | 488    | +4           | +3           | +2   | <b>+11</b>   | +3  | -2           | -5           | +0  |
| palatalization       | diacritic      | 23   | 69     | +7           | -3           | -14  | -19          | -12 | -9           | -14          | -6  |
| Total                |                |      | 6591   | <b>-8</b>    | +0           | +1   | <b>-9</b>    | +1  | <b>+4</b>    | <b>-11</b>   | -0  |

tions. If SpecAugment worked by simply making the model output fewer symbols overall, deletions should rise as insertions fall; instead, consonant deletions drop 11% and 14% when Tag-Based and Random Replacement are combined with SpecAugment, beyond SpecAugment’s own 7%. Insertions and deletions fall together, so masking makes emissions more accurate, not just rarer.

The repair in the other direction—Random Replacement reducing the errors masking introduces—is not coincidental. SpecAugment’s errors increase precisely where exposure is lowest. Tone in Kakabe and palatalization in Nashta and Shinekhen Buryat all fall below the 30th training exposure percentile (25.7th, 18.4th and 27.0th), while vowels fall above the 80th (82.9th, 89.8th, 89.2nd); an individual Kakabe tone value occurs in a median of 81 training tokens against 2,120 for an individual vowel. Random Replacement draws uniformly from the training vocabulary, so its synthetic text is flatter in word frequency than the real corpus, and tone and palatalization, concentrated in less frequent words, gain token share. In Kakabe, flattening raises the tone-bearing token share by 38% (palatalized tokens by 37% in Nashta), while vowel and consonant shares move by at most 7%.

Tag-Based Replacement also samples word types uniformly, but only among words sharing a tag (or a close tag, under the fallback of Section 3.2). The sequence of tags itself is kept, with its natural frequencies, and since tone and palatalization are in some tag classes and not others, their token shares are preserved almost exactly (tone at 1.00, palatalization at 0.97).

If exposure drives the errors, flattening word frequencies should reverse them and tag-based substitution should not. Tested directly against SpecAugment rather than the baseline, adding Random Replacement reduces tone errors by 9% (95% CI  $[-16, -0.4]$ ,  $p = .043$ ) and palatalization errors by 20% ( $p = .061$ ); no segmental category significantly worsens, and the addition significantly improves even the total labeled errors over SpecAugment alone ( $-2.7\%$ ,  $p = .013$ ). Tag-Based Replacement behaves as predicted. It leaves exposure unchanged, fails to reverse the tone errors, and with SpecAugment worsens them (+11% against the baseline,  $p = .025$ ). One result falls outside the exposure account: Tag-Based Replacement also reduces palatalization errors ( $-24\%$ ,  $p = .041$ ) despite changing exposure by nothing. Whatever drives this repair, it is not token frequency; identifying the channel is left to future work.

## 5.4. Limitations

Training on randomly replaced text exposes the model to invalid syntax, and text-based augmentation does not address out-of-vocabulary words. English differs from the target languages: read speech, orthographic transcription, and presence in both models’ pre-training. Three links in the mechanism remain untested directly: signaling is inferred from phonetics, low exposure and single-property signaling coincide for tone and palatalization, and Tag-Based Replacement’s reduction of palatalization errors, achieved without any exposure change, is not explained by our account. We mask frames and replace words; doing either at the phoneme level would need boundaries that neither CTC alignment nor an aligner trained on 1–5 h reliably gives. The tone results also rest on a single language (Kakabe), and palatalization on two. Finally, five fixed pivot-TTS voices likely underestimate what text-based augmentation can achieve, since speaker diversity drives synthetic data’s value [4, 5].

## 6. CONCLUSION

A method’s aggregate improvement says nothing about what it preserves; here the method with the most reliable aggregate improvements is the one that increases tone and palatalization errors, and no method we test is free of such trades. In the low-resource regime, SpecAugment with Random Replacement is the default: it gave the lowest error rates in all four languages, improving aggregate error over SpecAugment alone while repairing the cost of masking. Before adopting any augmentation, a practitioner can rank everything the target transcription distinguishes by training exposure and signaling; a masking method will increase errors on what is low on both. The same inspection suggests the repair: making the rare words that carry these categories more frequent reverses the errors. Where transcripts serve documentation, augmentation choices should be checked against an analysis of error types, not the total error rate alone.

## 7. ACKNOWLEDGEMENTS

The authors have no relevant financial or nonfinancial interests to disclose.

## 8. COMPLIANCE WITH ETHICAL STANDARDS

This research was conducted retrospectively using human subject data made available in open access by the Pangloss Collection (Vatlongos<sup>2</sup>: CC BY-NC-ND 3.0, Nashta<sup>3</sup>: CC BY-NC 2.5, Kakabe<sup>4</sup>: CC BY-NC-ND 3.0), Tokyo University of Foreign Studies<sup>5</sup> (CC BY-SA 4.0), and Panayotov et al. [24] (CC BY 4.0). Ethical approval was not required as confirmed by the license attached with the open access data.

## 9. REFERENCES

- [1] T. Ko, V. Peddinti, D. Povey, and S. Khudanpur, "Audio augmentation for speech recognition," in *Proceedings of Interspeech*, 2015, pp. 3586–3589.
- [2] D. S. Park, W. Chan, Y. Zhang, C. Chiu, B. Zoph, E. D. Cubuk, and Q. V. Le, "SpecAugment: A simple data augmentation method for automatic speech recognition," in *Proceedings of Interspeech 2019*, 2019, pp. 2613–2617.
- [3] G. Zhong, H. Song, R. Wang, L. Sun, D. Liu, J. Pan, X. Fang, J. Du, J. Zhang, and L. Dai, "External text based data augmentation for low-resource speech recognition in the constrained condition of OpenASR21 challenge," in *Proceedings of Interspeech*, 2022, pp. 4860–4864.
- [4] C. Du and K. Yu, "Speaker augmentation for low resource speech recognition," in *Proceedings of ICASSP*, 2020, pp. 7719–7723.
- [5] G. Yang, F. Yu, Z. Ma, Z. Du, Z. Gao, S. Zhang, and X. Chen, "Enhancing low-resource ASR through versatile TTS: Bridging the data gap," in *Proceedings of ICASSP*, 2025, pp. 1–5.
- [6] O. Adams, T. Cohn, G. Neubig, H. Cruz, S. Bird, and A. Michaud, "Evaluating phonemic transcription of low-resource tonal languages for language documentation," in *Proceedings of LREC 2018*, 2018.
- [7] A. Conneau, A. Baevski, R. Collobert, A. Mohamed, and M. Auli, "Unsupervised cross-lingual representation learning for speech recognition," *arXiv preprint arXiv:2006.13979*, 2020.
- [8] G. A. Miller, "WordNet: A lexical database for English," *Communications of the ACM*, vol. 38, no. 11, pp. 39–41, Nov. 1995.
- [9] J. Wei and K. Zou, "EDA: Easy data augmentation techniques for boosting performance on text classification tasks," in *Proceedings of EMNLP-IJCNLP*, 2019, pp. 6382–6388.
- [10] R. Zevallos, N. Bel, G. Cámbara, M. Farrús, and J. Luque, "Data augmentation for low-resource Quechua ASR improvement," in *Proceedings of Interspeech*, 2022, pp. 3518–3522.
- [11] Y. Hou, Y. Liu, W. Che, and T. Liu, "Sequence-to-sequence data augmentation for dialogue language understanding," in *Proceedings of the 27th International Conference on Computational Linguistics*, 2018, pp. 1234–1245.
- [12] N. R. Robinson, P. Ogayo, S. R. Gangu, D. R. Mortensen, and S. Watanabe, "When is TTS augmentation through a pivot language useful?," in *Proceedings of Interspeech*, 2022, pp. 3538–3542.
- [13] M. Bartelds, N. San, B. McDonnell, D. Jurafsky, and M. Wieling, "Making more of little data: Improving low-resource automatic speech recognition using data augmentation," in *Proceedings of ACL*, 2023, pp. 715–729.
- [14] N. Rossenbach, A. Zeyer, R. Schlüter, and H. Ney, "Generating synthetic audio data for attention-based speech recognition systems," in *Proceedings of ICASSP*. IEEE, 2020, pp. 7069–7073.
- [15] A. Ragni, K. M. Knill, S. P. Rath, and M. J. F. Gales, "Data augmentation for low resource languages," in *Proceedings of Interspeech*, 2014, pp. 810–814.
- [16] M. J. F. Gales, K. M. Knill, A. Ragni, and S. P. Rath, "Speech recognition and keyword spotting for low-resource languages: Babel project research at CUED," in *4th Workshop on Spoken Language Technologies for Under-Resourced Languages (SLTU 2014)*, 2014, pp. 16–23.
- [17] L. Lugosch, T. Likhomanenko, G. Synnaeve, and R. Collobert, "Pseudo-labeling for massively multilingual speech recognition," in *Proceedings of ICASSP*, 2022, pp. 7687–7691.
- [18] K. S. Bhogale, D. Mehendale, N. Parasa, T. Javed, P. Kumar, M. M. Khapra, et al., "Empowering low-resource language ASR via large-scale pseudo labeling," in *Proceedings of Interspeech*, 2024, pp. 2519–2523.
- [19] L. Zhang, S. Wu, and Z. Wang, "Phoneme-aware augmentation for robust Cantonese ASR under low-resource conditions," *Symmetry*, vol. 17, no. 9, 2025.
- [20] V. S. D. S. M. Akavarapu, M. Daniel, and G. Jäger, "Hard to be heard: Phoneme-level ASR analysis of phonologically complex, low-resource endangered languages," in *Findings of ACL*, 2026, pp. 3014–3028.
- [21] M. Baas and H. Kamper, "Voice conversion can improve ASR in very low-resource settings," in *Proceedings of Interspeech*, 2022, pp. 3513–3517.
- [22] A. Baevski, H. Zhou, A. Mohamed, and M. Auli, "wav2vec 2.0: a framework for self-supervised learning of speech representations," in *Proceedings of NeurIPS*, 2020.
- [23] A. Graves, S. Fernández, F. Gomez, and J. Schmidhuber, "Connectionist temporal classification: Labelling unsegmented sequence data with recurrent neural networks," in *Proceedings of ICML*, 2006, pp. 369–376.
- [24] V. Panayotov, G. Chen, D. Povey, and S. Khudanpur, "LibriSpeech: An ASR corpus based on public domain audio books," in *Proceedings of ICASSP*, 2015, pp. 5206–5210.
- [25] X. Wei, C. Cucchiari, R. van Hout, and H. Strik, "Automatic speech recognition and pronunciation error detection of Dutch non-native speech: Cumulating speech resources in a pluricentric language," *Speech Communication*, vol. 144, pp. 1–9, 2022.
- [26] Z. Liu, J. Spence, and E. Prud'hommeaux, "Investigating data partitioning strategies for crosslinguistic low-resource ASR evaluation," in *Proceedings of EACL*, 2023, pp. 123–131.
- [27] E. Ridge, "Language contexts: Vatlongos, southeast Ambrym (Vanuatu)," *Language Documentation and Description*, vol. 15, 2018.
- [28] E. Adamou, "Nashta grammatical sketch with examples linked to online corpus," *Slavische Mikrosprachen im absoluten Sprachkontakt*, 2022.
- [29] Y. Yamakoshi, "Three folktales in Shinekhen Buryat," *Asian and African Languages and Linguistics*, vol. 6, pp. 109–136, 2011.
- [30] A. Vydrina, *A corpus-based description of Kakabe, a Western Mande language: prosody in grammar*, Ph.D. thesis, Université Sorbonne Paris Cité, 2017.
- [31] M. Honnibal, I. Montani, S. Van Landeghem, and A. Boyd, "spaCy: Industrial-strength natural language processing in python," 2020.
- [32] M. Bisani and H. Ney, "Bootstrap estimates for confidence intervals in ASR performance evaluation," in *Proceedings of ICASSP*, 2004, vol. 1, pp. 1–409.
- [33] P. Koehn, "Statistical significance tests for machine translation evaluation," in *Proceedings of EMNLP*, 2004, pp. 388–395.
- [34] S. Holm, "A simple sequentially rejective multiple test procedure," *Scandinavian Journal of Statistics*, vol. 6, no. 2, pp. 65–70, 1979.
- [35] C. Minixhofer, O. Klejch, and P. Bell, "Evaluating and reducing the distance between synthetic and real speech distributions," in *Proceedings of Interspeech*, 2023, pp. 2078–2082.
- [36] P. Vieting, M. Kannen, B. Hilmes, R. Schlüter, and H. Ney, "Regularizing Learnable Feature Extraction for Automatic Speech Recognition," in *Proceedings of Interspeech*, 2025, pp. 4943–4947.

<sup>2</sup><https://pangloss.cnrs.fr/corpus/Vatlongos>

<sup>3</sup><https://pangloss.cnrs.fr/corpus/Nashta>

<sup>4</sup><https://pangloss.cnrs.fr/corpus/Kakabe>

<sup>5</sup>[https://tufs.repo.nii.ac.jp/search?search\\_type=2&q=1729497608274](https://tufs.repo.nii.ac.jp/search?search_type=2&q=1729497608274)