

Diacritic Restoration for Low-Resource Indigenous Languages: Case Study with Bribri and Cook Islands Māori

Rolando Coto-Solano¹, Daisy Li¹, Manoela Teleginski Ferraz¹, Olivia Sasse¹,
Cha Krupka¹, Sharid Loáiciga² and Sally Akevai Tenamu Nicholas³

¹ Dartmouth College, ² University of Gothenburg, Department of Philosophy, Linguistics and Theory of Science

³ The University of Auckland (Waipapa Taumata Rau)

{rolando.a.coto.solano, daisy.li.26, manoela.e.teleginski.ferraz.27,
olivia.g.sasse.25, cha.j.krupka.25}@dartmouth.edu, sharid.loaiciga@gu.se,
ake.nicholas@auckland.ac.nz

Abstract

We present experiments on diacritic restoration, a form of text normalization essential for natural language processing (NLP) tasks. Our study focuses on two extremely under-resourced languages: Bribri, a Chibchan language spoken in Costa Rica, and Cook Islands Māori, a Polynesian language spoken in the Cook Islands. Specifically, this paper: (i) compares algorithms for diacritics restoration in under-resourced languages, including tonal diacritics, (ii) examines the amount of data required to achieve target performance levels, (iii) contrasts results across varying resource conditions, and (iv) explores the related task of diacritic correction. We find that fine-tuned, character-level Large Language Models (LLMs) perform best, likely due to their ability to decompose complex characters into their UTF-8 byte representations. In contrast, massively multilingual models perform less effectively given our data constraints. Across all models, reliable performance begins to emerge with data budgets of around 10,000 words. Zero-shot approaches perform poorly in all cases. This study responds both to requests from the language communities and to broader NLP research questions concerning model performance and generalization in under-resourced contexts.

Keywords: diacritic restoration, under-resourced languages, text normalization

1. Introduction

Building digital corpora is essential to preserve invaluable linguistic materials. Language documentation provides the source data for many subfields of linguistics, such as phonology, morphology, syntax, semantics, and pragmatics (Chelliah, 2021, p. 2). Over the past decade, many researchers, linguists, and consortiums have worked closely with native speakers and language communities to create corpora including digitized text, audio, transcriptions, translations, stories, etc. (Agarwal and Anastasopoulos, 2024, p. 89). In the best case scenario, these corpora can be used not only to preserve, but also revitalize endangered languages (Shi et al., 2021) and reclaim Indigenous sovereignty (Viatori and Ushigua, 2007; Leonard, 2023). High-quality corpora can also be used to train NLP language education tools for learners (Zhang et al., 2022).

In many Indigenous communities, writing is of very recent adoption, and there is considerable variation on how words are spelled. This is a natural process which has occurred in every written language (Pombo, 2012; Moessner, 2017), and it is not usually an issue for communication. However, it is an issue for the automatic processing of the data, since variation can hinder performance (Pettersson et al., 2013; Salloum and Habash, 2011).

This is especially complicated with diacritics and non-alphabetic characters. Diacritics are symbols

that are appended to consonants or vowels. They can indicate tone (e.g. the macron in the Mandarin *sān* ‘three’), stress (e.g. the acute accent in Spanish *caminé* ‘I walked’), nasality (e.g. the Polish *ogonek* in *kąt* ‘angle’), and different vocalic qualities (e.g. the German umlaut in *öl* ‘oil’) amongst their many uses. These can show even greater variation in Indigenous language data (Roberts et al., 2013). They also elicit a wide range of opinions of the users of these languages: While school teachers and linguists insist on using diacritics because they feel they bring precision to the text, poets and other writers sometimes think the diacritics make the text feel crowded or difficult to read (Bird, 1999; Roberts, 2009; Kadyamusuma and Kadenge, 2013; Coto-Solano, 2017), or simply do not use them because they feel the text can be understood without them (Kishindo, 1998; Bernard et al., 2002; Saynes-Vazquez, 2002; Koffi, 2014). This variation in opinions leads to variation in spelling, which compounds to the difficulties of building corpora.

Diacritics are not just important for a standard representation of text, they also make a difference in their processing using NLP tools. The same diacritic can be encoded in multiple ways in Unicode, resulting in variation. Most state-of-the-art large language models (LLMs) use preprocessing steps to strip diacritics that might be inconsistently encoded or inconsistently omitted in certain scripts. However, stripping diacritics (or failure to normal-

ize diacritics) leads to degradation in the performance of downstream tasks such as POS-tagging or dependency parsing. This degradation is exacerbated in scripts such as Hebrew, which rely heavily on diacritics to disambiguate. Thus, the preferred method of handling diacritic inconsistencies is to apply Unicode normalization using any standard library module (Gorman and Pinter, 2025).

In this paper, we present experiments on diacritic restoration for two extremely under-resourced languages: Bribri and Cook Islands Māori. We evaluate several approaches, including character-based translation models, fine-tuned LLMs and zero-shot LLMs for the task. Our results show that character-based models work best, and that data masses as low as 10000 words can produce acceptable restoration models. This paper not only benchmarks the well-established task of diacritic restoration for two extremely under-resourced languages, but also responds to the communities' request for the development of NLP tools. In addition, we provide a detailed linguistic analysis of the outputs, with particular attention to tones and glottal stops, known to be problematic for Deep Learning and critical to work of language documentarians.

Our contributions are as follows: We:

- compare the performance of different algorithms when restoring diacritics in low-resource languages;
- analyze how much data is needed to reach specific performance levels;
- contrast the diacritic restoration outcomes between high and low resource languages; and
- conduct a brief study of the related task of diacritic correction.

2. Related Work

Diacritic Restoration Some of the research on diacritic restoration originates from work on normalization of historical or non-standard varieties of a language (Muñoz-Ortiz et al., 2025; Ljubešić et al., 2016). Other studies focus on heavily-diacritized languages such as Arabic, where diacritics are often omitted in written text (Masmoudi et al., 2019; Alqahtani et al., 2020; Hifny, 2021). Broadly, two main approaches have been proposed: statistical or rule-based models, and deep learning models. Statistical models are typically more computationally efficient (Hifny, 2021), but recent research has found that deep learning models are increasing their accuracy in under-resourced language tasks (Náplava et al., 2018). Diacritic restoration has also been investigated in European languages such as Hungarian, Turkish, and South Slavic languages (Novák and Siklósi, 2015; Ozer et al., 2018; Hucko and Lacko, 2018; Stankevičius et al., 2022), and in tonal languages like Vietnamese (Nga et al., 2019;

Tran et al., 2021; Le-Hong, 2021), Igbo (Ezeani et al., 2016) and other African languages such as Yoruba and Kikuyu (De Pauw et al., 2007).

Interestingly, there is work on diacritic restoration in te reo Māori, a Polynesian language spoken in Aotearoa New Zealand. Cocks et al. (2011) used a corpus of 3.8 million words and a Naïve Bayes classifier to obtain a 99% accuracy when reconstructing diacritics, and Nakano (2023) used an RNN-based algorithm and a 463,000 word corpus to obtain a 96% accuracy.

Cook Islands Māori is a Polynesian language spoken in the Cook Islands by 12,500 people, plus 10,000 in the diaspora (Ministry of Finance and Economic Management, Government of the Cook Islands, 2021; Nicholas, 2018). It is an endangered language, which means that there is a decreasing number of children who speak it. It is closely related to other Polynesian languages like te reo Māori, Hawaiian and Tahitian. In this language, henceforth referred to as CIM, there is a partially accepted system for spelling which includes two diacritics (Nicholas, 2017), but most public text is written with few or no diacritics. Figure 1 shows an example of such text¹. In order to make it compatible with diacritically-marked text, the diacritics would need to be added, a time-consuming task that requires expert knowledge which few people possess.

E korero tupuna no Makirau.

Tana angaanga i rave i tona tuatau. Te vai ra te aua puaka o Makirau koia oki e aua kai tangata. Kia tae ki tetai tuatau kua sere mai te tamaine a Ra i te aravei i tona tungane, ia Itiao. Kia sere maira aia kua akatano

Figure 1: Example of typewritten text in CIM with long vowel macrons and glottal stop saltillos absent (Ponga, 2013). In the Nicholas (2017) orthography, the first words would be *Tāna* 'anga'anga 'her work'.

There is previous work on NLP for Cook Islands Māori, including work on speech recognition (Coto-Solano et al., 2022a, 2018), text-to-speech (James et al., 2024), parsing (Karnes et al., 2023; Gilbert

¹We want to emphasize that orthographic variation is not an 'error' or a 'mistake' of the part of the speaker of the Indigenous language. Our goal is not to eliminate variation, but to respond to two practical considerations: (i) the community's request for the development of NLP tools, and (ii) the existence of certain orthographic standards that the community itself has adopted. Orthographic variation should never be treated as a mistake, and we encourage writers to continue producing text freely, in their preferred forms. The responsibility for handling standardization lies with the developers of NLP tools, not with language users. At the same time, computer scientists and NLP practitioners must take care not to establish a single variant as the de facto norm, as doing so risks erasing valuable linguistic diversity, an issue that deserves explicit discussion within the field.

et al., 2025) and forced alignment (Nicholas and Coto-Solano, 2019; Coto-Solano et al., 2022b).

Orthographic Correction and Diacritics Another way in which variation is expressed is in divergent and idiosyncratic uses of orthography. This is common during the normalization of orthographic norms, but it adds difficulty to NLP processing.

Bribri is a Chibchan language, spoken by 7,000 people in southern Costa Rica (INEC, 2011). It is also vulnerable (Sánchez Avendaño, 2013), and it has at least three different orthographic standards, represented in Constenla et al. (2004), Jara (2018) and Margery (2005). Text produced by other authors and by online users shows considerable idiosyncratic variation, as shown in figure 2. This doesn't need to be corrected to be understood by humans, but it might need to be standardized to be used along other NLP tools.

Narraciones Bribriás

CHĀBECÖL DESETWÖ STKINA

C häbekölpä' detsëwö skina e alaköl amī aké. Pe' dör máñal alábusj dör ape' wén serke kamje ese sewa. Ie' wen akoana ie dawá miska yébalök haí tók kamúk erök. Be' ate' sa' panúk ie'

Figure 2: Example of spelling variation in Bribri (IETSAY, 2001, 41). Using Constenla's (2004) orthography, the first two lines would be: *Tkabèköl ditsèwö tskina: Tkabèköl ditsèwö tskina e' alaköl amī akè*. "The Snake People are Created: The Snake People who were created (came from) a woman, (her) mother (and her) brother". For example, the word *ditsèwö* 'clan, people' appears as *desetwö* and *detsëwö*.

There is previous work on NLP in Bribri, including work on speech recognition (Coto-Solano, 2021; Coto-Solano et al., 2024), machine translation (Feldman and Coto-Solano, 2020; Jones et al., 2023), parsing (Coto-Solano et al., 2021; Karson and Coto-Solano, 2024), morphological analysis (Anderson et al., 2025), word embeddings (Coto-Solano, 2022) and forced alignment (Solórzano and Coto-Solano, 2017; Coto-Solano and Solórzano, 2016).

Cook Islands Māori: Long vowels and glottal stops The Cook Islands Māori Revised New Testament orthography (henceforth CIMR) (Nicholas, 2017; Bible Society of South Pacific, 2014) for CIM uses one diacritic and one non-alphabetic symbol, shown on Table 1: A macron to mark long vowels, and the glyph saltillo (A78C) to indicate a glottal stop phoneme. This last one shows considerable variation, and regularly appears as an apostrophe (0027), a left quote (201C) or an 'okina (02BB).

Long vowel (macron)	ē
Glottal stop (saltillo)	'e

Table 1: Diacritics and special symbols in CIM

Bribri: Tones, nasals and glottal stops Bribri uses three diacritics and one non-alphabetic symbol to indicate tone, an underline diacritic for vowel nasality, and an umlaut for lax vowel quality. These are shown in Table 2. In this paper we will ignore consonantal diacritics and instead consider them part of the consonant, in particular the tilde in 'ñ'.

Low tone	e
High tone	è
Falling tone	é
Low rising tone	ê
Glottal (low rising tone)	e'
Nasal vowel	<u>e</u>
Umlaut (lax) vowel	ë

Table 2: Diacritics and special characters in Bribri

In addition to appearing individually, the diacritics can be combined. For example, the word *é* 'only' is /ɛ̃˥/, and it has a lax vowel with a falling tone. The word *én* 'liver' is /ɛ̃˥˥/, and it contains a tense, nasal 'e' with a falling tone.

Here we will use the Constenla et al. (2004) orthography for the experiment, but other orthographies use other diacritics. For example Jara (2018) uses a tilde for nasals, and Margery (2005) uses ogoneks. More generally, there is wide variation in the orthographic representations of Bribri. For example, the word *taí* 'much' has been found in writing in 14 different ways.

3. Evaluation of Different Approaches for Diacritic Restoration

3.1. Data Inputs

In the diacritic restoration task, the input is a string without diacritics or non-alphabetical characters (e.g. Bribri *Is be shkēna* 'How are you?'), and the output is the string with its diacritics according to the community's orthographic norms (e.g. Bribri *Īs be' shkēnā*). Examples for both languages are shown in Table 3. For the Bribri input we eliminated all diacritics and special characters from the vowels (tones, nasals, umlauts for lax vowels, apostrophe for glottal stops). For the CIM input we eliminated the diacritic macron from the vowels and the saltillo glottal stop consonant.

The Bribri dataset is made up of 10,962 sentences (78,784 words), which are part of the AmericasNLP dataset (AmericasNLP, 2021). The CIM

Bribri Input	Is be shkëna
Output	ìs be' shkèna
Meaning	'How are you?'
CIM Input	i toku apii
Output	i tōku 'āpi'i
Meaning	'From my school'

Table 3: Examples of inputs

dataset contains 5,439 sentences (226,933 words) from the Te Vairanga Tuatua collection (Nicholas, 2012) in the Paradisec corpus (Thieberger, 2014). We made 5 random train/valid/test partitions with 80%, 10% and 10% of the words respectively.

3.2. Models

We used three families of algorithms for our tests: character-based statistical machine translation (SMT) models, fine-tuned LLMs, and zero-shot LLMs. The SMT system Moses (Koehn et al., 2007) was used as a baseline for all our experiments. SMT has been shown to be a strong baseline for diacritic restoration, and statistical methods remain useful in under-resourced languages settings (Kuparinen et al., 2023). We also implemented a statistical transformation of the text by training KenLM models (Heafield, 2011) to go from a form without diacritics to a form with them. This has proven successful in languages like Hawaiian (Kapali et al., 2025).

In principle, a fully rule-based system might appear to be the least resource-intensive alternative. However, in the context of language documentation, this is not a realistic option. Developing a comprehensive rule-based system that accounts for all contexts in which diacritics are used would require substantial time and expert knowledge. Language documenters who are not members of the speech community may lack the linguistic depth needed to construct such a system, while community-based documenters are often primarily engaged in revitalization efforts and cannot devote significant time to NLP development. Although expanding NLP training within language communities is an important long-term goal, current documentation scenarios call for technological solutions that reduce, rather than increase, the human workload.

For these reasons, we use statistical machine translation. Specifically, we used the Moses distribution by Ljubešić et al. (2016); Scherrer and Ljubešić (2016) which runs Moses character-based models.² All Moses systems use a n-gram language model of size six.

The main LLM selected for fine-tuning was ByT5 (Xue et al., 2022), chosen for two key reasons. First,

²<https://github.com/clarinsi/csmtiser>

it is a character-based model that operates directly on UTF-8 bytes, a feature that has proven effective for processing low-resource datasets. Second, it is multilingual, pretrained on the mC4 corpus (Raffel et al., 2020a), which covers a wide range of languages. To enable a comparison with more conventional, token-based architectures, we also included the monolingual T5 (Raffel et al., 2020b) and the multilingual mT5 (Xue et al., 2021). We had to extend T5 and mT5’s tokenizers to recognize Bribri diacritics by retraining SentencePiece (Kudo and Richardson, 2018). All three models were fine-tuned on our Bribri and CIM datasets and evaluated on their ability to restore diacritics.

For zero-shot LLM testing, we chose Claude Sonnet 4.5 2025-09-29 (Anthropic, 2025). We performed two types of experiments: First, we attempted a zero-shot diacritic restoration, where we gave the model the test sentences and asked it to add the diacritics. In the second experiment, we uploaded a portion of the training set (the equivalent of 150,000 tokens) and asked it to add the diacritics to the test sentences.

3.3. Benchmarking Model Effectiveness

In order to evaluate performance, we trained, fine-tuned or prompted the algorithms using the training and validation sets, and evaluated using the word error rate (WER) by predicting the output from the test set inputs. We did this five times (for each of the randomly generated sets based on our total data) and calculated the average WER.

In addition to comparing the expected output with the predictions of each algorithm, we need to calculate the accuracy when trying to predict specific diacritics. In order to do this, we transform the output strings into strings that focus on a single type of diacritic. Table 4 shows an example for Bribri. In the *Tones* transcription, each syllable is transcribed solely as its tone. Here, the first word *is* 'how' has a high tone, and is therefore represented as H. The second word has a glottal low rising tone, and is represented as G. In the *Nasal* condition, the syllables are either oral (O) or nasal (N), and in the *umlaut lax vowel* condition, the syllables either have a tense (T) or a lax (L) vowel. These transformed strings allow for two calculations: (i) a measure of how many tonal, nasal or umlaut diacritics are wrong overall, and (ii) a measurement of how accurate is the marking for each algorithm. The first is reported using WER, and the second is reported using the F1 for each tone and diacritic.

These transformations were also performed on the CIM text. In the *Long/short vowel* condition, each syllable is represented by whether it has a (L) or short (S) syllable. The *Consonants and glottals* condition has three possibilities for each syllable: it has a (G), a (C), or it is (N). A more simple version

	Transcription
Translation	How are you doing, sir/ma'am?
All	ìs be' shkèna, akékèpa
Tones	H G HL LFLL
Nasal vowels	O O ON OOOO
Umlaut (lax)	T T TT TLLT

Table 4: Different transcriptions of the Bribri phrase *How are you?* asked to an elder. These are used to analyze the different types of errors in the model. Tones distinguishes {L:low, H:high, F:falling, R:rising, G:glottal} tones. Nasal distinguishes between {O:oral, N:nasal} vowels. Umlaut distinguishes {T:tense, L:lax} vowels.

of this is the *Presence of glottals* condition, which describes an entire word by the number of glottals it has (one or more G), or as having no glottals (N).

	Transcription
Translation	From my school
All	i tōku 'āpi'i
Long and short vowels	S LS LSS
Consonants and glottals	N CC GCG
Presence of glottals	N N GG

Table 5: Different transcriptions of the CIM phrase *the school* to analyze the different types of errors. Long/short distinguishes {L: long, S: short} vowels. Consonants has {G: glottals, C: other consonants, N: words without glottals or consonants}. Presence of glottals has {G: glottal, N: word without glottals}.

Results The main baseline for our calculations is the “no restoration” condition. These are the results if we compare the original text to a text without any diacritics at all, and it gives us an idea of the actual improvement for each algorithm. Table 6 contains this and the results of the diacritic restoration for each algorithm. The baseline error rate is 85 for Bribri and 33 for CIM.

The statistical and zero-shot LLM methods have generally worse results, and fine-tuning the character-based model has the best results. The statistical transformation with KenLM had the largest word error (Bribri: 100; CIM 75). Following this is the other statistical method, Moses (Bribri: 88; CIM: 38). The LLM-based zero-shot using Claude also had very high error rates (Bribri: 85; CIM: 30), but showing examples to the LLM cuts the WER by half (Bribri: 40; CIM: 11). As for the T5 fine-tuning, the mT5 does not offer improvements in performance (Bribri: 43; CIM: 24), but the T5 does decrease the error (Bribri: 32; CIM:

	Bribri	CIM
No restoration	85	33
KenLM (n=4)	100	75
Moses	88	38
Claude Zero-Shot	85	30
Claude Few-Shot	40	11
T5 Fine-tuning	32	9
mT5 Fine-tuning	43	24
ByT5 Fine-tuning	16	7

Table 6: WER for statistical, fine-tuned LLM and zero-shot LLM approaches.

9). The best result comes from fine-tuning ByT5, with WER of 16 for Bribri and 7 for CIM. These results are comparable to the ones for te reo Māori, while using a much smaller mass of data. Given the low-performance of the KenLM, it will be not be included in subsequent calculations.

Table 7 shows the overall error for tonal, nasal and lax vowel diacritics in Bribri. Moses and the zero-shot Claude produced the results with the highest error. The few-shot Claude cut the error in half compared to the Zero-shot (72 versus 32). The mT5 and T5 fine-tuning had similar error rates; and again the ByT5 had the lowest error, with tones having only WER=13. Overall, tones are the most difficult diacritics to mark; the error rates for nasal and umlauts are 5 and 3 respectively.

	Tone	Nasal	Umlaut
No restoration	69	34	24
Moses	52	35	10
Claude Zero-Shot	72	40	30
Claude Few-Shot	32	17	15
T5 Fine-tuning	29	20	17
mT5 Fine-tuning	38	28	24
ByT5 Fine-tuning	13	5	3

Table 7: WER for tonal, nasal and umlaut transcriptions in Bribri, as described in Table 4.

The patterns observed in Bribri also held true for CIM. As can be seen in Table 8, ByT5 has the lowest error rate. Interestingly, the few-shot Claude was closer in error reduction compared to the ByT5.

Confusion patterns in complete restoration

We used the strings described in tables 4 and 5 to determine how often there are tonal replacements (e.g. how often a correct high tone is replaced by an incorrect falling tone). Using these strings allows us to treat this problem as a classification problem and to calculate precision, recall, and F1. Table 9 shows the F1 for specific diacritics in Bribri. Tonal diacritics that go directly on the vowels (i.e.

	Macron (long vowel)	Presence of glottals
No restoration	21	15
Moses	8	15
Claude Zero-Shot	19	15
Claude Few-Shot	8	2.8
T5 Fine-tuning	7	4
mT5 Fine-tuning	19	9
ByT5 Fine-tuning	6	2.5

Table 8: WER for macron and glottal transcriptions in CIM, as described in Table 5.

high, falling and rising) have the lowest F1, and are therefore the most difficult for the computer to learn. Non-tonal diacritics, such as those that indicate nasal and lax vowels have relatively high performance, and vowels with no explicit tonal diacritics (low tone) have the best performance for all of the algorithms. The low performance of the zero-shot solution is noteworthy: all tonal and nasal diacritics have values below F1:23, which indicates that this method might not be suitable for this task.

Table 10 shows the F1 for diacritics in CIM. The ByT5 has F1 results of 89 or above. The few-shot Claude is the runner-up, but its performance is 23 points lower for glottals. In general, glottals appear to be difficult for all algorithms except for the ByT5.

3.4. Amounts of Data and Performance

Once we determined the algorithm with the best performance, we created a series of smaller subsets of the training and validation data, using randomly selected sentences from the main train/validation sets. We then calculated the WER when the algorithm was trained on these lower masses of data. The objective of this experiment is to determine how much data is required to reach certain levels of error rates. This can be used by researchers and community members to determine whether training these diacritic restoration algorithms is possible with the available data for their languages.

For Bribri, we created five subsets with {1000, 5000, 10000, 25000, 50000} words each. For CIM, we created five subsets with {1000, 5000, 10000, 25000, 50000, 100000} words each. We used these to train and evaluate ByT5 models.

Results Figure 3 shows the relationship between number of words in the dataset and word error rate. There is significant variation with low data masses, e.g., 1000 words; approximately 2 pages of single-spaced text in English. At this point, the expected performance is about WER=80-90 for both languages. A performance of about WER=25 is reached with a mass of approximately ten thou-

sand words (about 20 pages of single-spaced text). After ten thousand words, the decrease in error rate becomes slower and more linear.

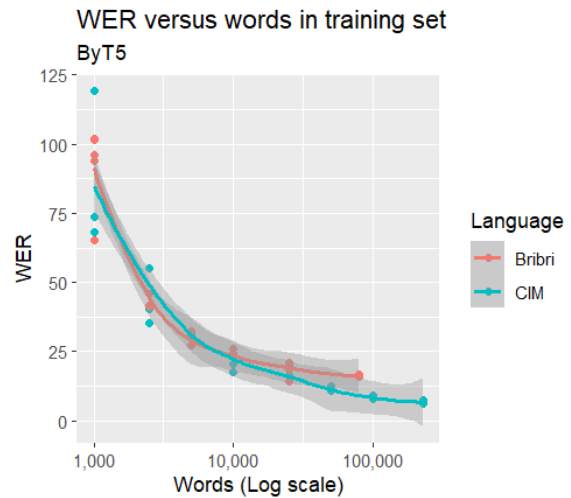


Figure 3: WER of ByT5 outputs as a function of word presence in the training set.

3.5. Comparison with High-Resource Languages

In order to better understand the performance of the best algorithm on low-resource languages, we also trained models using datasets from three higher-resourced languages: Spanish, Vietnamese and Pinyin-transcribed Mandarin Chinese.

We selected Spanish as a comparison with CIM because Spanish vowels also have two diacritics (the acute accent \acute{e} and the umlaut \ddot{u}). We used the MLS dataset (Pratap et al., 2020) and extracted a dataset of 6325 sentences (226929 words), similar in size to the CIM sample.

In order to create a comparison to Bribri, we selected Vietnamese and Mandarin transliterated into Pinyin. Vietnamese was used because it is the most directly comparable high-resource language to Bribri: It is the only language that is higher-resourced and has tonal diacritics in the Roman alphabet. In addition to this, Vietnamese is a language that is included in the mC4 dataset, and Vietnamese and Bribri have similar diacritic densities: In the Vietnamese corpus, 27% of the characters are diacritics, while in the Bribri corpus, 31% of the characters are diacritics. There are some differences: Vietnamese has nine diacritics for different tones and vowel qualities, higher than the number of diacritics in Bribri (6). We selected a set of 7716 sentences (78793 words) from the VIVOS corpus (Luong and Vu, 2016) and split it into five random train/valid/test sets.

As for Mandarin, its Pinyin representation is

	Tone						
	Low	High	Falling	Rising	Glottal	Nasal	Umlaut
Moses	80	55	52	54	0	0	80
Claude Zero-Shot	66	20	18	2	22	23	53
Claude Few-Shot	86	59	53	42	83	76	75
T5 Fine-tuning	81	56	50	48	74	66	64
mT5 Fine-tuning	79	53	44	43	64	64	62
ByT5 Fine-tuning	96	85	80	78	93	94	95

Table 9: Average F1 for individual diacritics for Bribri.

	Macron	Glottal
Moses	77	26
Claude Zero-Shot	40	12
Claude Few-Shot	78	68
T5 Fine-tuning	71	58
mT5 Fine-tuning	39	37
ByT5 Fine-tuning	89	91

Table 10: Average F1 for CIM diacritics

also comparable to Bribri in several ways. It has four tonal diacritics and one umlaut to indicate a rounded vowel. Like Bribri, it is not included in pre-training sets, as Mandarin is usually represented using Chinese characters. Also, Pinyin has a lower but similar diacritic density (22% of the characters are diacritics, compared to Bribri’s 30%). We selected a set of 7244 sentences (78779 words) from the Aishell-3 corpus (Shi et al., 2020) and again split it into five random train/valid/test sets.

Results Table 11 shows the WER for the low-resource languages (i.e. Bribri, CIM) and the high-resource languages chosen as comparison points (i.e. Spanish, Vietnamese and Pinyin-transliterated Mandarin Chinese).

	No restoration	ByT5 Fine-tuning
Bribri	85	16
Mandarin (Pinyin)	94	22
Vietnamese	87	16
CIM	33	7
Spanish	12	3

Table 11: Average WER diacritic restoration in high and low-resource languages

ByT5 reached similar WER values in Vietnamese as it did in Bribri, but this might not mean that the learning process was the same. On one hand, Vietnamese does appear in the ByT5 pretraining set, which might give it a leg up. On the other hand, the

Vietnamese diacritic system is more complex. Not only does Vietnamese have more diacritics (9 versus Bribri’s 6), it also has a higher tonal density and the tones themselves have a higher functional load (Hockett, 1967). This means that they distinguish more minimal pairs, making them more difficult to tag correctly. In contrast, Bribri has fewer diacritics per word, only one per morpheme, and its root morphemes tend to be polysyllabic. Therefore, even if both Vietnamese and Bribri have a WER=16, the results for Vietnamese are better because it had to solve a more difficult problem. Notice also that, due to the low mass of Vietnamese data it is seeing, these results are naturally lower than those for previous work (Le-Hong, 2021; Stankevičius et al., 2022).

As for Pinyin, its performance was lower than the one for Bribri and Vietnamese. This is to be expected, as (i) it has numerous diacritics, (ii) it has a high density of text with diacritics, and the diacritics carry a high functional load, and (iii) the Mandarin text that ByT5 is pretrained with possibly includes very little Pinyin, and even less text with Pinyin tones, as this text rarely occurs in public spaces (Mathias and Kennedy, 1980).

In the case of Spanish, the restoration of diacritics is aided by the knowledge in the pretraining. Both CIM and Spanish see their error reduced to a quarter by the fine-tuning, but the error rate for Spanish is very low, WER=3.

3.6. Diacritic Correction

In our final experiment, we replicate spell checking, where some diacritics might be correct but some might not be. We create synthetic “wrong” inputs by introducing probabilistic constraints on the appearance of diacritics, based on common replacements observed in naturalistic text. For example, in Bribri, if the original string has a vowel with a high tone, we set a 15% probability that the tone will be represented as a falling tone, 15% probability that it will stay as a high tone, and a 70% probability that the tone diacritic will be absent.

This procedure was also used to create CIM strings with synthetic spelling mistakes. For exam-

ple, if a word starts with the causative prefix {‘aka-}, we set a 20% probability that the glottal stop will appear. If the word ends with the nominalizer suffix {-‘anga}, there’s a 50% probability that the glottal stop will appear. In other cases, if the string has a glottal stop, there is a 20% chance that this will appear in the synthetic text.

We used this transformation to modify the train-/valid/test sets for each language. These were used to train models using the best-performing algorithm. The ultimate objective of this experimental condition is to replicate spell checking, and study how these systems might behave when correcting real texts with spelling variation.

Results The experiment for the simulated spelling correction had results of WER=7 for CIM and WER=13 for Bribri. The results for Bribri are slightly better than those for the complete restoration condition (WER=16), but they were the same for CIM. More research is needed, in particular, the building of a corpus of natural Bribri and its corrected version in order to further study this problem.

4. Discussion

4.1. Classical versus LLM-based Work on Under-Resourced Languages

Whether classical machine learning or LLM-based approaches are “better” for under-resourced languages depends greatly on the definition of “under-resourced” and the goals of the project. For extremely low-resource languages that need immediate work from language documenters, fine-tuning an LLM or using an adequately controlled commercial LLM might make sense for some tasks. When there is a similar or related language in the pretraining data, for example, this might help. This was observed for the zero and few shot conditions: CIM had much lower error rates than Bribri, possibly because the models include data for other Polynesian languages. In addition to this, if the task can be accommodated to the point where specialized knowledge of the language is not required (e.g., diacritic restoration), these solutions will be more efficient than creating a rule-based system, once a certain mass of data is reached. On the other hand, maintaining control of hallucinations made the task complicated. For example, only small amounts of data could be processed at a time; giving the system large lists would create confusion in its outputs.

Classical approaches, however, remain more accessible from a computational perspective. For instance, our Moses models ran entirely on CPUs. Nevertheless, we observed substantially lower WER and higher F1 performance with byT5, which is also a character-level model. This indicates

that knowledge encoded during pre-training can be leveraged for under-resourced languages. While Moses did not achieve high overall performance, it did particularly well with umlaut and macron transcription. This suggests a strong sensitivity to context-dependent phenomena, as expected from an n-gram language model, and indicates that larger data amounts may be required for the SMT model to reach its full potential.

ByT5 clearly outperformed all other models, reinforcing the understanding that tokenization continues to pose challenges for NLP. Its description paper actually positions it as a step towards tokenizer-free approaches. Because ByT5 decomposes each character into separate UTF-8 bytes, it can learn fine-grained character-level patterns such as those found in Bribri. Off-the-shelf tokenizers, in contrast, are trained mainly on high-resource languages, needing re-training in the context of under-resourced languages like those presented here. This process can introduce biases from vocabulary size or BPE merges when working with a new language. The trade-off is that ByT5 has higher memory and compute requirements, and it is an open question how well this approach scales up to typical NLP tasks beyond the character-level transduction experiments reported here.

One interesting result is that we could not replicate the results in [Kapali et al. \(2025\)](#). In that paper, KenLM models produced the best results for Hawaiian. However, the models produced the worst results in the languages presented here. This might be because their Hawaiian corpus is much larger: 77 thousand sentences versus 4300 sentences in the CIM corpus used here.

4.2. Multilingual Models

Multilingual LLMs are trained under the assumption that exposure to multiple languages helps generalization, an approach that has shown success in other tasks such as machine translation ([Team et al., 2022](#)). However, under-resourced languages vary widely in their amount of available digital data ([Joshi et al., 2020](#)). As a result, although under-resourced languages may benefit from the transfer of knowledge from the large-resourced ones, this is not guaranteed. A typologically related language may not be present in the pretraining data, and the quality of Indigenous language data might be low because it has been gathered without the collaboration or consent of the community.

In our experiments, we observe that multilingual word-level models did not outperform their monolingual counterparts. When benefits from multilingual pretraining emerged, they were limited to the high-resource languages in its collection (e.g., Spanish). The multilingual pretraining itself did not seem to be aiding Bribri or CIM, given the low performance

of the mT5 models and the Claude experiments. These findings suggest that, while LLMs can be useful for fine-tuning, the most effective base models may not be massively multilingual ones. Instead, models specialized or adapted for a single under-resourced language may yield better results.

4.3. Tones

Tones are very difficult to transcribe reliably, even for highly qualified and experienced linguists and community members. This also appears to be true of computer systems. The tonal languages had the lowest overall performance, and the tonal diacritics for Bribri were its lowest performing ones. A tool that can aid with tonal transcription is highly necessary for documentary linguists, and this is an underexplored area in the literature. We encourage more researchers to explore this problem.

4.4. Data Sovereignty Concerns

Not only did we have issues with the computing power used for the fine-tuning of LLMs, we also had to be careful with data sovereignty issues while using our data. While all of the raw data we used is available through public sources, its labeled form is not. We used Claude because it allowed privacy options so that the datasets would not be incorporated in future models, an option that wasn't available with models from OpenAI (2025), for example.

While the use of this tool might be desirable to specialists and people working on language revitalization, there are potential risks in releasing them to the public. Releasing it to the public might increase interest in learning the language amongst community members, but it could also result in an unintentional endorsement of a certain way of writing these languages. Not only might this be biased towards one language variety or another, but it might be biased towards LLM hallucinations. These might be construed to be more authoritative given the elevated status of computer systems in 21st century life (Eubanks, 2017), and might be considered the "correct" form of these languages, intimidating actual speakers who speak correctly but who feel that their correct speech and writing diverges from the computer's recommendations.

5. Conclusions

In this paper we have explored the task of diacritic restoration on two extremely under-resourced languages. We found that some Deep Learning methods, i.e., character-based LLMs, are beneficial in a variety of sub-tasks in this setting, even with very small amounts of data. We also explored how these algorithms work with tonal languages, whose transcription is vexing and time-consuming for linguists

and community members who work on language documentation.

6. Acknowledgments

SL has been supported by a grant from the Swedish Research Council (VR project 2014-39) for the establishment of the Centre for Linguistic Theory and Studies in Probability (CLASP) at the University of Gothenburg. We thank Yves Scherrer for helpful discussions and comments.

7. Bibliographical References

- Milind Agarwal and Antonios Anastasopoulos. 2024. *A Concise Survey of OCR for Low-Resource Languages*. In *Proceedings of the 4th Workshop on Natural Language Processing for Indigenous Languages of the Americas (AmericasNLP 2024)*, pages 88–102, Mexico City, Mexico. Association for Computational Linguistics.
- Sawsan Alqahtani, Ajay Mishra, and Mona Diab. 2020. A Multitask Learning Approach for Diacritic Restoration. *arXiv preprint arXiv:2006.04016*.
- AmericasNLP. 2021. AmericasNLP 2021 Shared Task on Open Machine Translation. <https://github.com/AmericasNLP/americasnlp2021>. [Accessed 25-10-2025].
- Carter Anderson, Mien Nguyen, and Rolando Coto-Solano. 2025. Unsupervised, semi-supervised and llm-based morphological segmentation for bribri. In *Proceedings of the Fifth Workshop on NLP for Indigenous Languages of the Americas (AmericasNLP)*, pages 63–76.
- Anthropic. 2025. Introducing Claude Sonnet 4.5. <https://www.anthropic.com/news/claude-sonnet-4-5>.
- H Russell Bernard, George N Mbeh, and W Penn Handwerker. 2002. Does marking tone make tone languages easier to read? *Human organization*, 61(4):339–349.
- Bible Society of South Pacific. 2014. *Cook Islands Māori Revised New Testament: Digital Publication*. Life.Church / YouVersion.
- Steven Bird. 1999. When marking tone reduces fluency: an orthography experiment in Cameroon. *Language and Speech*, 42(1):83–115.
- Shobhana L. Chelliah. 2021. *Impact of Language Documentation*, pages 91–94. Springer International Publishing, Cham.

- John Cocks et al. 2011. A word-based approach for diacritic restoration in Māori. In *Proceedings of the Australasian Language Technology Association Workshop 2011*, pages 126–130.
- Adolfo Constenla, Feliciano Elizondo, and Francisco Pereira. 2004. *Curso Básico de Bribri*. Editorial de la Universidad de Costa Rica.
- Rolando Coto-Solano. 2017. *Tonal reduction and literacy in Me'ph aa Vátháá*. The University of Arizona.
- Rolando Coto-Solano. 2021. Explicit tone transcription improves asr performance in extremely low-resource languages: A case study in bribri. In *Proceedings of the first workshop on natural language processing for Indigenous languages of the Americas*, pages 173–184.
- Rolando Coto-Solano. 2022. Evaluating word embeddings in extremely under-resourced languages: A case study in bribri. In *Proceedings of the 29th International Conference on Computational Linguistics*, pages 4455–4467.
- Rolando Coto-Solano, Tai Wan Kim, Alexander Jones, and Sharid Loáiciga. 2024. Multilingual models for asr in chibchan languages. In *Proceedings of the 2024 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies (Volume 1: Long Papers)*, pages 8521–8535.
- Rolando Coto-Solano, Sharid Loáiciga, and Sofía Flores-Solórzano. 2021. Towards universal dependencies for bribri. In *Proceedings of the Fifth Workshop on Universal Dependencies (UDW, SyntaxFest 2021)*, pages 16–29.
- Rolando Coto-Solano, Sally Akevai Nicholas, Samiha Datta, Victoria Quint, Piripi Wills, Emma Ngakuravaru Powell, Liam Koka'ua, Syed Tanveer, and Isaac Feldman. 2022a. Development of automatic speech recognition for the documentation of cook islands māori. In *Proceedings of the Thirteenth Language Resources and Evaluation Conference*, pages 3872–3882.
- Rolando Coto-Solano, Sally Akevai Nicholas, Britany Hoback, and Gregorio Tiburcio Cano. 2022b. Managing data workflows for untrained forced alignment: examples from costa rica, mexico, the cook islands, and vanuatu. *The Open Handbook of Linguistic Data Management*, 35.
- Rolando Coto-Solano, Sally Akevai Nicholas, and Samantha Wray. 2018. Development of natural language processing tools for cook islands māori. In *Proceedings of the Australasian Language Technology Association Workshop 2018*, pages 26–33.
- Rolando Coto-Solano and Sofía Flores Solórzano. 2016. Alineación forzada sin entrenamiento para la anotación automática de corpus orales de las lenguas indígenas de costa rica. *Kánina*, 40(4):175–199.
- Guy De Pauw, Peter W Wagacha, and Gilles-Maurice De Schryver. 2007. Automatic Diacritic Restoration for Resource-Scarce Languages. In *International Conference on Text, Speech and Dialogue*, pages 170–179. Springer.
- Virginia Eubanks. 2017. *Automating inequality: How high-tech tools profile, police, and punish the poor*. Picador.
- Ignatius Ezeani, Mark Hepple, and Ikechukwu Onyenwe. 2016. Automatic Restoration of Diacritics for Igbo Language. In *International Conference on Text, Speech, and Dialogue*, pages 198–205. Springer.
- Isaac Feldman and Rolando Coto-Solano. 2020. Neural machine translation models with back-translation for the extremely low-resource indigenous language bribri. In *Proceedings of the 28th International Conference on Computational Linguistics*, pages 3965–3976.
- Gabriel H Gilbert, Rolando Coto-Solano, Sally Akevai Nicholas, Lauren Houchens, Sabrina Barton, and Trinity Pryor. 2025. Crosslingual dependency parsing of hawaiian and cook islands māori using universal dependencies. In *Proceedings of the 18th International Conference on Parsing Technologies (IWPT, SyntaxFest 2025)*, pages 40–50.
- Kyle Gorman and Yuval Pinter. 2025. [Don't Touch My Diacritics](#). In *Proceedings of the 2025 Conference of the Nations of the Americas Chapter of the Association for Computational Linguistics: Human Language Technologies (Volume 2: Short Papers)*, pages 285–291, Albuquerque, New Mexico. Association for Computational Linguistics.
- Kenneth Heafield. 2011. [KenLM: Faster and smaller language model queries](#). In *Proceedings of the Sixth Workshop on Statistical Machine Translation*, pages 187–197, Edinburgh, Scotland. Association for Computational Linguistics.
- Yasser Hifny. 2021. Recent advances in Arabic syntactic diacritics restoration. In *ICASSP 2021-2021 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 7768–7772. IEEE.

- Charles F Hockett. 1967. The quantification of functional load. *Word*, 23(1-3):300–320.
- Andrej Hucko and Peter Lacko. 2018. Diacritics restoration using deep neural networks. In *2018 World Symposium on Digital Intelligence for Systems and Machines (DISA)*, pages 195–200. IEEE.
- Instituto de Estudios de las Tradiciones Sagradas de Abia Yala IETSAY. 2001. *Narraciones Bribris: Talamanca y Cabagra, II Edición*, 2nd edition. Fundación Coordinadora de Pastoral Aborígen, San José, Costa Rica.
- INEC. 2011. [X Censo Nacional de Población y VI de Vivienda 2011 - Territorios Indígenas - Principales Indicadores Demográficos y Socioeconómicos](#).
- Jesin James, Rolando Coto-Solano, Sally Akevai Nicholas, Joshua Zhu, Bovey Yu, Fuki Babasaki, Jenny Tyler Wang, and Nicholas Derby. 2024. Development of community-oriented text-to-speech models for māori ‘avaiki nui (cook islands māori). In *Proceedings of the 2024 Joint International Conference on Computational Linguistics, Language Resources and Evaluation (LREC-COLING 2024)*, pages 4820–4831.
- Carla Victoria Jara. 2018. *Gramática de la lengua bribri*. E-Digital ED.
- Alex Jones, Rolando Coto-Solano, and Guillermo González Campos. 2023. Talamt: Multilingual machine translation for cabécar-bribri-spanish. In *Proceedings of the 3rd Workshop on Multi-lingual Representation Learning (MRL)*, pages 106–117.
- Pratik Joshi, Sebastin Santy, Amar Budhiraja, Kalika Bali, and Monojit Choudhury. 2020. The state and fate of linguistic diversity and inclusion in the NLP world. *arXiv preprint arXiv:2004.09095*.
- McLoddy R Kadyamusuma and Maxwell Kadenge. 2013. The potential effect of non-tone marking on homograph processing in Shona: a preliminary exploration. *Per Linguam*, 29(2):85–94.
- Jaden Kapali, Keaton Williamson, and Winston Wu. 2025. [Statistical and neural methods for Hawaiian orthography modernization](#). In *Proceedings of the 2025 Conference on Empirical Methods in Natural Language Processing*, pages 35149–35155, Suzhou, China. Association for Computational Linguistics.
- Sarah Karnes, Rolando Coto-Solano, and Sally Akevai Nicholas. 2023. Towards universal dependencies in cook islands māori. In *Proceedings of the Sixth Workshop on the Use of Computational Methods in the Study of Endangered Languages*, pages 124–129.
- Jessica Karson and Rolando Coto-Solano. 2024. Morphological tagging in bribri using universal dependency features. In *Proceedings of the 4th Workshop on Natural Language Processing for Indigenous Languages of the Americas (AmericasNLP 2024)*, pages 56–66.
- PJ Kishindo. 1998. On the standardization of Citembuka and Ciyao orthographies: Some observations. *South African Journal of African Languages*, 18(4):85–91.
- Philipp Koehn, Hieu Hoang, Alexandra Birch, Chris Callison-Burch, Marcello Federico, Nicola Bertoldi, Brooke Cowan, Wade Shen, Christine Moran, Richard Zens, Chris Dyer, Ondřej Bojar, Alexandra Constantin, and Evan Herbst. 2007. [Moses: Open Source Toolkit for Statistical Machine Translation](#). In *Proceedings of the 45th Annual Meeting of the Association for Computational Linguistics Companion Volume Proceedings of the Demo and Poster Sessions*, pages 177–180, Prague, Czech Republic. Association for Computational Linguistics.
- Ettien N Koffi. 2014. Towards an Optimal Representation of Tones in the Orthographies of African Languages. *Linguistic Portfolios*, 3(1):12.
- Taku Kudo and John Richardson. 2018. [SentencePiece: A simple and language independent subword tokenizer and detokenizer for neural text processing](#). In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing: System Demonstrations*, pages 66–71, Brussels, Belgium. Association for Computational Linguistics.
- Olli Kuparinen, Aleksandra Miletić, and Yves Scherrer. 2023. [Dialect-to-Standard Normalization: A Large-Scale Multilingual Evaluation](#). In *Findings of the Association for Computational Linguistics: EMNLP 2023*, pages 13814–13828, Singapore. Association for Computational Linguistics.
- Phuong Le-Hong. 2021. Diacritics generation and application in hate speech detection on Vietnamese social networks. *Knowledge-Based Systems*, 233:107504.
- Wesley Y Leonard. 2023. Refusing “endangered languages” narratives. *Daedalus*, 152(3):69–83.
- Nikola Ljubešić, Tomaž Erjavec, and Darja Fišer. 2016. [Corpus-Based Diacritic Restoration for South Slavic Languages](#). In *Proceedings of the Tenth International Conference on Language Resources and Evaluation (LREC’16)*, pages 3612–3616, Portorož, Slovenia. European Language Resources Association (ELRA).

- Hieu-Thi Luong and Hai-Quan Vu. 2016. A non-expert Kaldi recipe for Vietnamese speech recognition system. In *Proceedings of the Third International Workshop on Worldwide Language Service Infrastructure and Second Workshop on Open Infrastructures and Analysis Frameworks for Human Language Technologies (WLSI/OIAF4HLT2016)*, pages 51–55.
- Enrique Margery. 2005. *Diccionario fraseológico bribri-español español-bribri*, second edition. Editorial de la Universidad de Costa Rica.
- Abir Masmoudi, Salima Mdhaffar, Rahma Sellami, and Lamia Hadrich Belguith. 2019. Automatic diacritics restoration for Tunisian dialect. *ACM Transactions on Asian and Low-Resource Language Information Processing (TALLIP)*, 18(3):1–18.
- Jim Mathias and Thomas L Kennedy. 1980. *Computers, Language Reform, and Lexicography in China: A Report*. Washington State University Press.
- Ministry of Finance and Economic Management, Government of the Cook Islands. 2021. **Census 2021: Key findings**. <https://www.mfem.gov.ck/statistics/census-and-surveys/census/267-census-2021>.
- Lilo Moessner. 2017. Standardization. *The History of English*, 4:167–187.
- Alberto Muñoz-Ortiz, Verena Blaschke, and Barbara Plank. 2025. **Evaluating Pixel Language Models on Non-Standardized Languages**. In *Proceedings of the 31st International Conference on Computational Linguistics*, pages 6412–6419, Abu Dhabi, UAE. Association for Computational Linguistics.
- Masato Nakano. 2023. Using Neural Network Models for Diacritics Restoration in Māori. Master's thesis, Georgetown University.
- Jakub Náplava, Milan Straka, Pavel Straňák, and Jan Hajič. 2018. **Diacritics Restoration Using Neural Networks**. In *Proceedings of the Eleventh International Conference on Language Resources and Evaluation (LREC 2018)*, Miyazaki, Japan. European Language Resources Association (ELRA).
- Cao Hong Nga, Nguyen Khai Thinh, Pao-Chi Chang, and Jia-Ching Wang. 2019. Deep Learning based Vietnamese Diacritics Restoration. In *2019 IEEE international symposium on multimedia (ISM)*, pages 331–3313. IEEE.
- Sally Akevai Nicholas. 2012. **Te Vairanga Tuatua o te Te Reo Māori o te Pae Tonga: Cook Islands Māori (Southern dialects)**. Collection SN1 at catalog.paradisec.org.au [Open Access]. Collector: Sally Akevai Nicholas.
- Sally Akevai Nicholas. 2017. *Ko te Karāma o te Reo Māori o te Pae Tonga o Te Kuki Airani: A Grammar of Southern Cook Islands Māori*. Unpublished PhD thesis, University of Auckland, Auckland N.Z. <http://hdl.handle.net/2292/32929> [Accessed 2017-07-14].
- Sally Akevai Nicholas and Rolando Coto-Solano. 2019. Glottal variation, teacher training and language revitalization in the cook islands. In *Proceedings of the 19th International Congress of Phonetic Sciences, University of Melbourne, Australia*, pages 3602–3606.
- Sally Akevai Te Namu Nicholas. 2018. Language contexts: Te Reo Māori o te Pae Tonga o te Kuki Airani also known as Southern Cook Islands Māori. *Language Documentation and Description*, 15:64.
- Attila Novák and Borbála Siklósi. 2015. Automatic diacritics restoration for Hungarian. In *Proceedings of the 2015 Conference on Empirical Methods in Natural Language Processing*, pages 2286–2291.
- OpenAI. 2025. ChatGPT (GPT-5) [Large language model]. <https://chat.openai.com/>. Accessed: 2025-10-25.
- Zeynep Ozer, Ilyas Ozer, and Oguz Findik. 2018. Diacritic restoration of Turkish tweets with word2vec. *Engineering Science and Technology, an International Journal*, 21(6):1120–1127.
- Eva Pettersson, Beáta Megyesi, and Joakim Nivre. 2013. Normalisation of historical text using context-sensitive weighted Levenshtein distance and compound splitting. In *Proceedings of the 19th Nordic conference of computational linguistics (Nodalida 2013)*, pages 163–179.
- Elena Llamas Pombo. 2012. Spanish. In Susan Baddeley and Anja Voeste, editors, *Orthographies in Early Modern Europe*, pages 15–62. De Gruyter Mouton.
- Akaiti Ponga. 2013. **Genealogy book of Akaiti Ponga**. Retrieved from <https://familysearch.org/search/catalog/2017778?availability=Family%20History%20Library> accessed January 16, 2013. International Genealogical Index. Index. FamilySearch. . Index based upon data collected by the Genealogical Society of Utah.
- Vineel Pratap, Qiantong Xu, Anuroop Sriram, Gabriel Synnaeve, and Ronan Collobert. 2020. MLS: A large-scale multilingual

- dataset for speech research. *arXiv preprint arXiv:2012.03411*.
- Colin Raffel, Noam Shazeer, Adam Roberts, Katherine Lee, Sharan Narang, Michael Matena, Yanqi Zhou, Wei Li, and Peter J. Liu. 2020a. Exploring the limits of transfer learning with a unified text-to-text transformer. *J. Mach. Learn. Res.*, 21(1).
- Colin Raffel, Noam Shazeer, Adam Roberts, Katherine Lee, Sharan Narang, Michael Matena, Yanqi Zhou, Wei Li, and Peter J. Liu. 2020b. [Exploring the Limits of Transfer Learning with a Unified Text-to-Text Transformer](#). *Journal of Machine Learning Research*, 21(140):1–67.
- David Roberts. 2009. Visual crowding and the tone orthography of African languages. *Written Language & Literacy*, 12(1):140–155.
- David Roberts, SR Borgwaldt, and T Joyce. 2013. A tone orthography typology. *Typology of writing systems*, 51:85.
- Wael Salloum and Nizar Habash. 2011. Dialectal to standard Arabic paraphrasing to improve Arabic-English statistical machine translation. In *Proceedings of the first workshop on algorithms and resources for modelling of dialects and language varieties*, pages 10–21.
- Carlos Sánchez Avendaño. 2013. Lenguas en peligro en Costa Rica: vitalidad, documentación y descripción. *Revista Káñina*, 37(1):219–250.
- Floria Edaena Saynes-Vazquez. 2002. *Zapotec language shift and reversal in Juchitan, Mexico*. Ph.D. thesis, The University of Arizona.
- Yves Scherrer and Nikola Ljubešić. 2016. Automatic normalisation of the Swiss German Archi-Mob corpus using character-level machine translation. In *Proceedings of the 13th Conference on Natural Language Processing (KONVENS 2016)*, pages 248–255.
- Jiatong Shi, Jonathan D. Amith, Xuankai Chang, Siddharth Dalmia, Brian Yan, and Shinji Watanabe. 2021. [Highland Puebla Nahuatl Speech Translation Corpus for Endangered Language Documentation](#). In *Proceedings of the First Workshop on Natural Language Processing for Indigenous Languages of the Americas*, pages 53–63, Online. Association for Computational Linguistics.
- Yao Shi, Hui Bu, Xin Xu, Shaoji Zhang, and Ming Li. 2020. Aishell-3: A multi-speaker mandarin tts corpus and the baselines. *arXiv preprint arXiv:2010.11567*.
- Sofía Flores Solórzano and Rolando Coto-Solano. 2017. Comparison of two forced alignments systems for aligning bribri speech. *CLEI Electronic Journal*, 20(1):2–1.
- Lukas Stankevičius, Mantas Lukoševičius, Jurgita Kapočiūtė-Dzikienė, Monika Briedienė, and Tomas Krilavičius. 2022. Correcting diacritics and typos with a ByT5 transformer model. *Applied Sciences*, 12(5):2636.
- NLLB Team, Marta R. Costa-jussà, James Cross, Onur Çelebi, Maha Elbayad, Kenneth Heafield, Kevin Heffernan, Elahe Kalbassi, Janice Lam, Daniel Licht, Jean Maillard, Anna Sun, Skyler Wang, Guillaume Wenzek, Al Youngblood, Bapi Akula, Loic Barrault, Gabriel Mejia Gonzalez, Prangthip Hansanti, John Hoffman, Smerley Jarrett, Kaushik Ram Sadagopan, Dirk Rowe, Shannon Spruit, Chau Tran, Pierre Andrews, Necip Fazil Ayan, Shruti Bhosale, Sergey Edunov, Angela Fan, Cynthia Gao, Vedanuj Goswami, Francisco Guzmán, Philipp Koehn, Alexandre Mourachko, Christophe Ropers, Safiyyah Saleem, Holger Schwenk, and Jeff Wang. 2022. [No Language Left Behind: Scaling Human-Centered Machine Translation](#).
- Nick Thieberger. 2014. PARADISEC: Building methods for preserving ethnographic fieldwork recordings and providing long term access. *Australasian Sound Archive, The*, 39:43–53.
- Quang-Linh Tran, Gia-Huy Lam, Van-Binh Duong, and Trong-Hop Do. 2021. A Study on Diacritic Restoration Problem in Vietnamese Text using Deep Learning based Models. In *2021 IEEE International Conference on Communication, Networks and Satellite (COMNETSAT)*, pages 306–310. IEEE.
- Maximilian Viatori and Gloria Ushigua. 2007. [Speaking Sovereignty: Indigenous Languages and Self-Determination](#). *Wicazo Sa Review*, 22(2):7–21.
- Linting Xue, Aditya Barua, Noah Constant, Rami Al-Rfou, Sharan Narang, Mihir Kale, Adam Roberts, and Colin Raffel. 2022. [ByT5: Towards a Token-Free Future with Pre-trained Byte-to-Byte Models](#). *Transactions of the Association for Computational Linguistics*, 10:291–306.
- Linting Xue, Noah Constant, Adam Roberts, Mihir Kale, Rami Al-Rfou, Aditya Siddhant, Aditya Barua, and Colin Raffel. 2021. [mT5: A Massively Multilingual Pre-trained Text-to-Text Transformer](#). In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 483–498, Online. Association for Computational Linguistics.

Shiyue Zhang, Ben Frey, and Mohit Bansal. 2022.
How can NLP Help Revitalize Endangered Languages? A Case Study and Roadmap for the Cherokee Language.