
Composing RNNs and FSTs for Small Data: Recovering Missing Characters in Old Hawaiian Text

Oiwi Parker Jones*^{1,2} and Brendan Shillingford*^{2,3}
Jesus College¹, University of Oxford², Google DeepMind³
oiwi.parkerjones@jesus.ox.ac.uk,
brendan.shillingford@cs.ox.ac.uk

Abstract

In contrast to the older writing system of the 19th century, modern Hawaiian orthography employs characters for long vowels and glottal stops. These extra characters account for about one-third of the phonemes in Hawaiian, so including them makes a big difference to reading comprehension and pronunciation. However, transliterating between older and newer texts is a laborious task when performed manually. We introduce two related methods to help solve this transliteration problem automatically, given that there were not enough data to train an end-to-end deep learning model. One method is implemented, end-to-end, using finite state transducers (FSTs). The other is a hybrid deep learning approach which approximately composes an FST with a recurrent neural network (RNN). We find that the hybrid approach outperforms the end-to-end FST by partitioning the original problem into one part that can be modelled by hand, using an FST, and into another part, which is easily solved by an RNN trained on the available data.

1 Introduction

From 1834 to 1948, more than 125,000 newspaper pages were published in the Hawaiian language (Nogelmeier, 2010). Yet by 1981, many expected this once flourishing language to die (Benton, 1981). Hawaiian has since defied expectations and experienced the beginnings of a remarkable recovery (Warner, 2001; Wilson and Kamanā, 2001). However much of the literary inheritance that is contained in the newspapers has become difficult for modern Hawaiians to read, since the newspapers were written in an orthography that failed to represent about one-third of the language’s phonemes. This orthography, which we will refer to as the *missionary orthography*, excluded Hawaiian phonemes that did not have equivalents in American English (see Schütz, 1994), namely long vowels /i: e: a: o: u:/ and glottal stop /ʔ/. By contrast, the *modern Hawaiian orthography*, an innovation of Pukui and Elbert’s Hawaiian dictionary (Pukui and Elbert, 1957), presents a nearly perfect, one-to-one mapping between graphemes and phonemes. The process of manual transliteration from missionary to modern Hawaiian orthography is extremely labor intensive. Yet the cultural benefits are so great that hundreds of pages of newspaper-serials have already been transliterated by hand, such as Nogelmeier’s new edition of the epic tale of *Hi‘iakaikapoliopele*, the volcano goddess’s sister (Ho‘oulumāhiehie, 2007). Critically important as such efforts are to the continued revitalization of this endangered language, they are still only a small sample of the material that could be made available to a modern Hawaiian audience.

In this paper, we propose to automate, or semi-automate, the transliteration of old Hawaiian texts into the modern orthography. Following a brief review of related work (Section 2), we begin by describing a dataset of modern Hawaiian (Section 3). In Section 4, we present two methods for recovering missing graphemes (and hence phonemes) from the missionary orthography. The first composes a

* Authors contributed equally.

series of weighted FSTs; the second approximately composes an FST with a recurrent neural network language model (RNNLM) using a beam search procedure. Both approaches require only modern Hawaiian texts for training, which are much more plentiful than parallel corpora. Section 5 reports the results of our transliteration experiments using a simulated parallel corpus, as well as two 19th century newspaper articles for which we also have modern Hawaiian transcriptions. Being based on FSTs, both approaches are modular and extensible. We observe useful and promising results for both of our methods, with the best results obtained by the hybrid FST-RNNLM. These results showcase the strength of combining established hand-engineering methods with deep learning in a smaller data regime, with practical applications for an endangered language.

2 Related work

Many of the themes that we address relate to existing literature. For example, Hajič et al. (2000) and Scannell (2014) have written on machine translation (MT) for closely related languages and on multilingual text normalization, respectively. Though language-relatedness makes MT easier (Kolovratník et al., 2010), state-of-the-art techniques such as neural machine translation (NMT) have not performed well for languages with little data (Östling and Tiedemann, 2017). So while the Hawaiian transliteration problem could be cast as an instance of MT or of NMT, we chose to sidestep the scarcity of parallel data by not considering such approaches.

Hybrid approaches, which combine expert knowledge for well-understood structures with deep learning for data-plentiful subproblems, offer rich opportunities for data-efficient modelling. In particular, prior work has explored ways to combine FSTs and RNNs. For example, Sproat and Jaitly (2016) used an FST to restrict the search space when decoding from an RNN; Rastogi et al. (2016) incorporated RNN information into an FST. However, these approaches differ from the approximate FST-to-RNN composition algorithm that we introduce here (in Section 4.2).

3 Data

3.1 Phonemes & the modern orthography

Ignoring case, there is a neat mapping between the modern Hawaiian orthography and the Hawaiian phonemic inventory. The phonemic inventory contains eight consonants $\langle h k l m n p v \int \rangle$ and ten vowels, of which five are short $\langle a e i o u \rangle$ and five are long $\langle a: e: i: o: u: \rangle$ (Parker Jones, 2018).

The consonants map onto the orthographic symbols $\langle H h K k L l M m N n P p W w \int \rangle$, where we give the upper- and lower-case variants in adjacent pairs: $\langle H h \rangle$ for $\langle h \rangle$, $\langle K k \rangle$ for $\langle k \rangle$, \dots , $\langle W w \rangle$ for $\langle v \rangle$. An exception, the symbol $\langle \int \rangle$ has only one variant which maps to $\langle \int \rangle$. The vowels map onto the symbols: $\langle A a E e I i O o U u \bar{A} \bar{a} \bar{E} \bar{e} \bar{I} \bar{i} \bar{O} \bar{o} \bar{U} \bar{u} \rangle$. Note that vowel length is denoted by the absence or presence of a macron (e.g. $\langle A \rangle$ and $\langle a \rangle$ map onto short $\langle a \rangle$ and $\langle \bar{A} \rangle$ and $\langle \bar{a} \rangle$ map onto long $\langle a: \rangle$).

The Hawaiian conventions for capitalization, numbering, and punctuation are analogous to those in English, except again there is no upper-case variant of $\langle \int \rangle$, so the following vowel is capitalized instead (e.g. ‘*Okakopa* ‘October’). In foreign words, such as *kolorofolorokalapona* ‘chlorofluorocarbon’, one can find the additional consonants: $\langle B b C c D d F f G g J j Q q R r S s T t V v X x Y y Z z \rangle$.

3.2 Modern & missionary orthographies

The primary difference between the missionary and modern Hawaiian orthographies is that the missionary orthography does not encode long vowels or the glottal stop. For example, the following Hawaiian phrases were recorded by a 19th-century German traveller in the missionary orthography: *Ua oia au, E ue ae oe ia li, E ao ae oe ia ia* (Chamisso, 1837, p. 7). In the modern orthography these become: *Ua ‘ō ‘ia au* ‘I am speared’, *E uē a’e ‘oe iā ‘Ī‘ī* ‘You must weep for ‘Ī‘ī (a person)’, and *E a’o a’e ‘oe iā ia* ‘You teach him’ (Elbert and Pukui, 1979, p. 3).

We can convert text in the modern Hawaiian orthography *backward* chronologically to an approximate missionary orthography by mapping each glottal stop $\langle \int \rangle$ to the empty string ϵ , and each long vowel, e.g. $\langle \bar{a} \bar{e} \bar{i} \bar{o} \bar{u} \rangle$, to its corresponding short vowel, $\langle a e i o u \rangle$. As a first approximation, we may treat mappings from the modern-to-missionary orthographies as unambiguously many-to-one; thus there is information loss. We will return to secondary differences between the orthographies in Section 6.

Source	Chars	Words
Ulukau(160 texts)	6,518,451	1,334,451
Hi‘iakaikapoliopole	1,272,935	259,947
Wikipedia	577,794	10,221
<i>Total</i>	8,369,180	1,604,619

Table 1: Modern data sources and their sizes.

To illustrate, the following four words in the modern orthography all map to the same missionary string *aa*: *a‘a* (root), *‘a‘a* (brave), *‘a‘ā* (crumbly lava rock), and *‘ā‘ā* (stutter).

The *forward* mapping from missionary-to-modern orthographies is one-to-many. Thus the missionary string *aa* could map to *a‘a*, *‘a‘a*, *‘a‘ā*, or *‘ā‘ā*. The *transliteration problem* we address here seeks to discover how we can use context to recover the information not present in the missionary orthography that modern Hawaiian orthography retains.

3.3 Data sources

We draw on three sources for modern Hawaiian text: the main text of *Hi‘iakaikapoliopole* (Ho‘oulumāhie, 2007), 160 short texts from *Ulukau: The Hawaiian Electronic Library*, and the full Hawaiian Wikipedia (see Table 1).²

For evaluation, we simulate a missionary-era version of the modern texts using the backward mapping described above. In addition, we evaluated our models on a couple of 19th century newspaper samples for which we have parallel missionary-era and modern text. Both simulated and real parallel corpora will be described in Section 5.

4 Models

We can frame the task of transliterating from missionary-to-modern Hawaiian orthographies as a sequence transduction problem. Many deep learning approaches (e.g. Sutskever et al., 2014; Graves, 2012) are not easily applicable to this task since we do not have a sufficiently large dataset of parallel texts. Instead, we focus on approaches that mix hand-designed FSTs with trained language models, including deep learning approaches like RNNLMs (Mikolov et al., 2010).

4.1 End-to-end FSTs

Our initial approach represents the mapping from missionary to modern orthography using a composition of (weighted) FSTs; for a thorough review, see Mohri (1997).

First, we construct a finite state acceptor, I , from the input text. Here we construct a trivial chain-shaped acceptor that accepts only the input text. Each symbol in the input text is represented by a state which emits this symbol on a single transition that moves to the next state. The transition emitting the final symbol in the string leads to the sole accepting state.

Second, we construct an FST called C , which models potential orthography changes that can occur when transliterating from the missionary to modern Hawaiian orthography (see Figure 1). For example, two non-deterministic transitions introduce an optional long-vowel map for ⟨a⟩: (a : a) and (a : ā). Another transition inserts glottal stops: (ε : ‘). By capturing the orthographic changes we know to occur, the composition $I \circ C$ produces a large set of candidates to be narrowed using the language model.

Third, we use the modern Hawaiian text from Section 3.3 to construct and evaluate a number of character-level n-gram language models, of various combinations of order and Katz backoff and Kneser-Ney (KN) smoothing (Katz, 1987; Kneser and Ney, 1995) (see Section 5.1 for a list of models that we trained). N-gram language models can be expressed as weighted FSTs. We denote the n-gram

²*Ulukau: The Hawaiian Electronic Library*: <http://ulukau.org/>, Hawaiian Wikipedia: <https://haw.wikipedia.org/>. Both accessed 19 May 2018.

or weighted FST language model as G . Character-level models are used as we wanted to generalize to out-of-vocabulary words, which we expect to occur frequently in a relatively small corpus like the one we have for Hawaiian.

Finally, we use this model to infer modern orthography given a piece of text in missionary orthography as input, then compose the FSTs to form the *search graph* FST: $S = I \circ C \circ G$. The minimum cost path through S gives the predicted modern orthography. Of these n-gram-based approaches, we found the Kneser-Ney-based models to perform best. These approaches will be referred to as FST- C -NGRAM-KN and FST- C_{wb} -NGRAM-KN.

We circumvent the lack of a large, non-simulated parallel corpus by training the language model exclusively on text in the modern Hawaiian orthography. In turn, the orthographic transliteration FST C produces candidates which are disambiguated by the language model. The result is finally evaluated against the ground-truth modern text.

Although the orthographic transliteration model is an approximation, and thus not exhaustive, it embodies an explicit and interpretable representation that can be extended independently of the rest of the model. To illustrate this, we constructed a variant C_{wb} (where wb stands for word boundary). C_{wb} optionally inserts a space after each vowel using an additional arc that maps ($\epsilon : space$) (again see Figure 1). This variant is able to model some changes in Hawaiian’s word-boundary conventions (Wilson, 1976), such as *alaila* becoming *a laila* which demarcates the preposition *a* ‘until’ and noun *laila* ‘then’. We report on the use of C_{wb} to predict modern equivalents of 19th century newspaper samples in Section 5. An example prediction can be found in Section 6, with more in Appendix A.




Figure 1: Two FSTs. The first, C , transduces between missionary and modern orthographies. The second, C_{wb} , introduces optional spaces (or word boundaries) after a vowel. In each FST, s_1 serves as both the initial and end state, while labelled arrows denote arcs. In the labelled transitions, V and \bar{V} are variables for short and long vowels, respectively, and C can be any consonant other than ‘ ϵ ’. Some arcs, for upper-case letters, numbers, and punctuation, have been omitted for brevity.

4.2 Hybrid approach: Composing FSTs and RNNLMs

As an alternative approach, we tried combining the FST C in the previous section with an RNNLM, since RNNLMs often generalize better than n-gram language models (Mikolov et al., 2010).

An RNN is a neural network that models temporal or sequential data, by iterating a function mapping a state and input to a new state and output. These can be stacked to form a deep RNN. For language modelling, each step of the final RNN layer models a word or character sequence via $p(w_1, \dots, w_n) = \prod_{i=1}^n p(w_i | w_{1:i-1})$ and can be trained by maximum likelihood. Recent language modeling work has typically used the long short-term memory (LSTM) unit, because of its favorable gradient propagation properties (Hochreiter and Schmidhuber, 1997). All RNNs in this paper are LSTMs.

Our goal in the hybrid approach is to replace the n-gram language model in the end-to-end FST with an RNNLM. While the minimum cost path through an FST can be computed exactly as done in the previous section, it is not straightforward to compose the relation defined by an FST with an arbitrary one like that defined by an RNNLM. Nonetheless, a minimum cost path through the composition of

the FST and the RNNLM can be defined as a path (i.e. label sequence) that minimizes the sum of the FST path cost and the RNNLM cost.

We can approximately find a minimum cost path for the composition of the two models by a breadth-first search over the FST graph, using a beam search, as follows. At any particular iteration, consider a single beam element. The beam element holds the current FST and RNN states, and the path taken through the FST so far. We follow each possible arc from the current FST state, each producing a new child beam element, and feed the output symbol into the RNN (unless it is ϵ). We note that there may be duplicate beam elements due the nondeterminicity of the FST, in which case the lower cost edge wins. We sort by the sum of the FST and RNN costs, keep the lowest-cost K , and then proceed to the next iteration. If a beam element is on an accepting state of the FST, it is kept as-is between iterations. Detailed pseudocode is provided in Figure 2.

```

B ← {(fst-state = s0, rnn-state = 0, tokens = [], score = 0)} ▷ single initial beam element
while ∃b ∈ B : b.fst-state is not final do
  B' = {} ▷ child beam elements
  for each beam element b ∈ B that is not in a final state do
    for each arc leaving b.fst-state (new state s', output c, weight w) do
      ▷ Note that we do not need to use the input symbol for our purposes, but in
      some settings it may be useful to keep track of it.
      if b.tokens[-1] ≠ ε then
        (h', p(y)) = run RNN one timestep with input (b.rnn-state, c)
        rnncost = -log p(b.tokens[-1])
        ▷ denotes the negative log probability of the RNN's discrete distribution
        output at this timestep
      else
        h' = b.rnn-state
        rnncost = 0
      end if
      b' = copy(b)
      b'.score = b'.score + rnncost + w
      b'.tokens.append(c)
      b'.rnn-state = h'
      B' = B' ∪ {b'}
    end for
  end for
  B' = B' ∪ {b ∈ B : b is in a final state}
  Sort B' by score, put best (lowest score) K into B
end while

```

Figure 2: Algorithm for approximately composing an FST with a RNN LM.

Conceptually, this algorithm performs the same operation as the n-gram language model, except that we replace the n-gram language model with an RNN language model, and then we search over the FST graph, producing scores from FST weights and RNN's outputs. Incidentally, we note that our implementation of the algorithm is slightly different than how we present it in the pseudocode, as we grouped RNN operations into batches for computational efficiency.

5 Results

5.1 Summary of language models trained

In this paper, we will refer to the hybrid models as FST-RNNLM in general, and as FST-RNNLM- C and FST-RNNLM- C_{wb} if we want to distinguish which FST was used. The end-to-end FST models will similarly be referred to as FST- C and FST- C_{wb} , with suffixes denoting what kind of n-gram and smoothing were used. For example, FST- C -7GRAM-KN denotes an end-to-end FST with a 7-gram language model and Kneser-Ney smoothing. Here is a list of n-gram language models that we considered (with perplexity scores in parentheses):

- 7-gram Katz backoff (3.15)
- 7-gram Kneser-Ney (3.07)
- 7-gram Kneser-Ney with backoff (3.15)
- 9-gram Katz backoff (3.15)
- **9-gram Kneser-Ney (2.95)**
- 9-gram Kneser-Ney with backoff (3.24)
- 11-gram Katz backoff (3.39)
- **11-gram Kneser-Ney (2.94)**
- 11-gram Kneser-Ney with backoff (3.67)

We note that 13- and higher n-gram language models performed far worse. The numbers in parentheses show character-level perplexities produced using a validation set that was drawn from a synthetic parallel corpus, which will be explained in Section 5.3.

We also trained two character-level RNNLMs with the following configurations:

- 2 layers \times 200 unit LSTM (2.70)
- **3 layers \times 200 unit LSTM (2.65)**

Both RNNLMs were trained with plain SGD, using a batch size of 30, a learning rate of 10, truncated backpropagation through time with 45 unrolling steps, and the gradient renormalized to norm 1 when it exceeded 1. Both RNNLMs used a dropout rate of 0.2 at the input, between RNN layers, and after the last RNN layer. The best n-gram and RNNLMs are highlighted in **bold**. Overall, the best results were produced with the hybrid approach, using RNNLMs.

5.2 Evaluation

Because we were unable to find a sufficiently large corpus of parallel texts in the missionary and modern Hawaiian orthographies, we instead trained the n-gram and RNN language models on a corpus of modern Hawaiian texts (*ground-truth*) (see Section 3.3). Parallel corpora were only required to test predictions from missionary-era to modern texts, which were produced by composing one of the FSTs, C or C_{wb} , with either an n-gram or RNN language model. To evaluate the accuracy of our approaches, we first derived a synthetic parallel corpus from our collection of modern Hawaiian texts. We also used a small but real parallel corpus, based on two 19th century newspaper texts and their hand-edited modern equivalents. Results based on these parallel corpora are reported in the following subsections.

5.3 Simulated parallel corpus (modern texts)

To produce a simulated parallel corpus (*input-missionary*), we systematically reduced the orthography in the modern texts (Section 3.3) using the backward mapping described above (Section 3.2). We then applied the end-to-end FST and hybrid FST-RNNLM models (Section 4), with the aim of learning a forward mapping between orthographies that recovers the lost information.

We evaluated the predicted modern text (*predictions*) by computing

$$\text{CERR} = \frac{d(\text{prediction}, \text{ground-truth})}{d(\text{input-missionary}, \text{ground-truth})},$$

where d denotes character-level edit distance. This is a modification of character error rate, normalized by the distance of the input and target rather than by the length of the target. We note that CERR may be high even when the predictions are very accurate as $d(\text{input-missionary}, \text{ground-truth})$ is small when the text is similar in both orthographies.

Table 2 gives more details about the strongest models from both approaches. Out of the Kneser-Ney n-gram models, we found that the FST- C -9GRAM-KN and the version modelling word boundaries (FST- C_{wb} -9GRAM-KN) performed best on the synthetic parallel corpus and newspapers, respectively. C_{wb} was not applied to the synthetic parallel corpus as the synthetic parallel corpus did not model word splitting. However, the hybrid models (FST-RNNLM) outperformed all end-to-end FSTs.

Transliteration model	LM perplexity		Transliteration performance (%CERR)		
	Valid.	Test	Corpus	Newspaper 1	Newspaper 2
FST-(C/C_{wb})-9GRAM-KN	2.95	3.02	26.6%	50.7% / 39.3%	52.5% / 47.2%
FST-(C/C_{wb})-11GRAM-KN	2.94	3.02	27.8%	53.9% / 41.3%	54.1% / 48.7%
FST-RNNLM-(C/C_{wb})	2.65	2.69	16.3%	47.2% / 34.3%	49.8% / 41.2%

Table 2: Performance (%CERR). Slash-separated pairs denote FSTs incapable/capable of inserting word boundaries, respectively; see Section 4. The -KN suffix denotes Kneser-Ney smoothing. The data from Section 3.3 is used for evaluating the modern-orthography language model perplexity, and “Corpus” evaluates test-set transliteration performance from the synthetic missionary text back to the original modern text.

5.4 Real parallel corpus (newspaper texts)

Not content to evaluate the model on simulated missionary orthography, we also evaluated it on two newspaper texts, using selections originally published in 1867 and 1894 for which we had 19th century and manually-edited modern equivalents. The newspaper selections discuss *Kahahana*, one of the last kings of O’ahu (Kamakau and Perreira, 2002), and *Uluhaimalama*, a garden party and secret political gathering, held after the deposition of Hawai’i’s last queen (Pukui et al., 2006). Unlike the synthetic missionary corpus evaluation where we did not model word splitting, we found that replacing C with C_{wb} on the newspaper texts significantly improved the output, especially on the FST-RNNLM model. Again, we found a hybrid model (FST-RNNLM- C_{wb}) to be the best performing model overall (Table 2).

6 Conclusions and future work

With this paper we introduced a new transliteration problem to the field, that of mapping between old and new Hawaiian orthographies—where the modern Hawaiian orthography represents linguistic information that is missing from older missionary-era texts. One difficulty of this problem is that there is a limited amount of Hawaiian data, making data-hungry solutions like end-to-end deep learning unlikely to work. To solve the transliteration problem, we therefore proposed two models: the first was implemented end-to-end using weighted FSTs; the second was a hybrid deep learning approach that combined an FST and an RNNLM. Both models gave promising results, but the hybrid approach performed best. It allowed us to use a more powerful recurrent neural network-based language model, despite our dataset’s small size. Factoring a problem like ours into one part that can be modelled exactly using expert domain knowledge and into another part that can be learned directly from data using deep learning is not novel; however it is a promising research direction for data-efficient modelling. To our knowledge, this paper is the first to describe a procedure to compose an FST with an RNN by approximately performing beam search over the FST.

While the role of the RNNLM part of the hybrid approach may be obvious, the FST component plays an important role too. For example, the hand-designed FST can be replaced without needing to re-train the RNNLM. We tried to showcase this modularity by constructing two FSTs which we referred to as C and C_{wb} , where only the latter allowed the insertion of spaces. Future work could extend the FST to model orthographic changes suggested by an error analysis of the current model’s predictions.

<i>Input:</i>	Weheia ka Malapua Alii a Kanuia na Uluwehi no ia Wao.
<i>Prediction:</i>	Wehe_’ia ka Māla_pua Ali’i a Kanu_’ia na Uluwehi nō ia Wao.
<i>Ground-truth:</i>	Wehe ‘ia ka Māla Pua Ali’i a Kanu ‘ia nā Uluwehi no ia Wao.

Table 3: Sample prediction. Data from from a 19th century Hawaiian newspaper (Pukui et al., 2006). Correct predictions are **green and bold**. Characters omitted by the model as compared to the ground-truth are denoted by *blue italics*; erroneous insertions or substitutions are denoted by **red underline**. To make white spaces explicit, we represent them with the symbol ‘_’. More sample predictions can be found in Appendix A.

An example of the current model’s predictions (i.e. missionary input, predicted modern text, modern ground-truth) is given in Table 3. In this example, we see the model correctly predicting some word boundaries, glottal stops and long vowels; however, we note that the model could not predict uppercase *Pua* (correct), because the input text contained lowercase *pua* (incorrect), and no (p : P) transitions were included in C or C_{wb} . Similar observations (see Appendix A) motivate new mappings for consonant substitutions like (r : l) and (s : k) that occur in loanword adaptations (e.g. *rose* \Rightarrow *loke*). The error analysis also motivates mappings to delete spaces ($_ : \epsilon$) and to handle contractions, like *na’lii* \Rightarrow *nā ali’i*. We could further incorporate linguistic knowledge of Hawaiian into the FST, which tells us, for example, about expected sequences of vowels (Parker Jones, 2010). Additional improvements to the hybrid model might be obtained by increasing the amount of modern Hawaiian text used to train the RNNLM. One way to do this would be to accelerate the rate at which missionary-era Hawaiian texts are modernized. To this end, we hope that the present models will be used within the Hawaiian community to semi-automate, and thereby accelerate, the modernization of old Hawaiian texts.

Acknowledgments

Mahalo to M. Puakea Nogelmeier for providing an electronic copy of *Hi’iakaikapoliopole* (Ho’oulumāhie, 2007). We gratefully acknowledge the support of the NVIDIA Corporation for donating the Titan X(p) GPUs used in this research. We also gratefully acknowledge the ACL for allowing us to reproduce parts of a shorter (5 page) version of this work (Shillingford and Parker Jones, 2018).

References

- Richard A Benton. 1981. *The flight of the Amokura: Oceanic languages and formal education in the South Pacific*. New Zealand Council for Educational Research, Wellington.
- Adelbert von Chamisso. 1837. *Über die Hawaiische Sprache, Vorgelegt der Königlichen Academie der Wissenschaften zu Berlin am 12, Januar, 1837*. Weidmann, Leipzig.
- Samuel H. Elbert and Mary Kawena Pukui. 1979. *Hawaiian Grammar*. University of Hawai’i Press, Honolulu.
- Alex Graves. 2012. Sequence transduction with recurrent neural networks. *arXiv preprint arXiv:1211.3711*.
- Jan Hajič, Jan Hric, and Vladislav Kuboň. 2000. Machine translation of very close languages. In *ANLC ’00 Proceedings of the sixth conference on Applied natural language processing*, pages 7–12.
- Sepp Hochreiter and Jürgen Schmidhuber. 1997. Long short-term memory. *Neural computation*, 9(8):1735–1780.
- Ho’oulumāhie. 2007. *Ka Mo’olelo o Hi’iakaikapoliopole*. Awaiaulu Press, Honolulu. Edited by M. Puakea Nogelmeier.
- Samuel Manaiakalani Kamakau and Hiapo Perreira. 2002. Ka mo’olelo o Kahahana, māhele 1. *Ka Ho’oilina*, 1(1):102–121.
- Slava Katz. 1987. Estimation of probabilities from sparse data for the language model component of a speech recognizer. *IEEE transactions on acoustics, speech, and signal processing*, 35(3):400–401.
- Reinhard Kneser and Hermann Ney. 1995. Improved backing-off for m-gram language modeling. In *Acoustics, Speech, and Signal Processing, 1995. ICASSP-95., 1995 International Conference on*, volume 1, pages 181–184. IEEE.
- David Kolovratník, Natalia Klyueva, and Ondřej Bojar. 2010. Statistical machine translation between related and unrelated languages. In *Proceedings of the Conference on Theory and Practice of Information Technologies (ITAT-09)*, pages 31–36.

- Tomas Mikolov, Martin Karafiát, Lukás Burget, Jan Černocký, and Sanjeev Khudanpur. 2010. Recurrent neural network based language model. In *Proceedings of Interspeech*, pages 1045–1048.
- Mehryar Mohri. 1997. Finite-state transducers in language and speech processing. *Computational linguistics*, 23(2):269–311.
- M. Puakea Nogelmeier. 2010. *Mai Pa‘a i ka Leo: Historical Voice in Hawaiian Primary Materials: Looking Forward and Listening Back*. Bishop Museum Press, Honolulu.
- Robert Östling and Jörg Tiedemann. 2017. Neural machine translation for low-resource languages. *Computing Research Repository*, arXiv:1708.05729. Version 1.
- Oiwi Parker Jones. 2010. *A computational phonology and morphology of Hawaiian*. Ph.D. thesis, University of Oxford.
- Oiwi Parker Jones. 2018. Illustrations of the IPA: Hawaiian. *Journal of the International Phonetic Association*, 48:103–115.
- Mary Kawena Pukui and Samuel H. Elbert. 1957. *Hawaiian-English Dictionary*. University of Hawai‘i Press, Honolulu.
- Mary Kawena Pukui, Holo Ho‘opai, Oiwi Parker Jones, and Keao NeSmith. 2006. No ka mahi‘ai ‘ana, māhele 6. *Ka Ho‘oilina*, 5(1):2–23.
- Pushpendre Rastogi, Ryan Cotterell, and Jason Eisner. 2016. Weighted finite-state transductions with neural context. In *Proceedings of the 2016 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 623–633.
- Kevin Scannell. 2014. Statistical models for text normalization and machine translation. In *Proceedings of the First Celtic Language Technology Workshop*, pages 33–40.
- Albert J. Schütz. 1994. *The Voices of Eden: A history of Hawaiian language studies*. University of Hawai‘i Press, Honolulu.
- Brendan Shillingford and Oiwi Parker Jones. 2018. Recovering missing characters in old Hawaiian writing. *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 4929–4934.
- Richard Sproat and Navdeep Jaitly. 2016. RNN approaches to text normalization: A challenge. *Computing Research Repository*, arXiv:1611.00068. Version 2.
- Ilya Sutskever, Oriol Vinyals, and Quoc V Le. 2014. Sequence to sequence learning with neural networks. In *Advances in neural information processing systems*, pages 3104–3112.
- Sam L. No‘eau Warner. 2001. The movement to revitalize Hawaiian language and culture. In Leanne Hinton and Kenneth Hale, editors, *The Green Book of Language Revitalization in Practice*, pages –144. Academic Press, San Diego, CA.
- William H. Wilson. 1976. Standardized Hawaiian orthography. Manuscript, University of Hawai‘i.
- William H. Wilson and Kauanoē Kamanā. 2001. “Mai loko mai o ka ‘i‘ini: Proceeding from a dream”: The ‘Aha Pūnana Leo connection in Hawaiian language revitalization. In Leanne Hinton and Kenneth Hale, editors, *The Green Book of Language Revitalization in Practice*, pages 147–176. Academic Press, San Diego, CA.

A Appendix

Sample predictions from newspaper data

Each block of three lines shows the input, followed by the prediction, followed by the ground-truth. Characters omitted by the model as compared to the ground-truth are denoted by *blue and italic*, whereas characters that are erroneously inserted or substituted (or should have been substituted) for another character are denoted by red underline. When the incorrect character is a space, the space is replaced with ‘_’.

First 10 sentences in Newspaper 1:

Input 1: Weheia ka Malapua Alii a Kanuia na Uluwehi no ia Wao.

Prediction 1: Wehe ‘ia ka Māla pua Ali‘i a Kanu_‘ia na Uluwehi nō ia Wao.

Ground-truth 1: Wehe ‘ia ka Māla Pua Ali‘i a Kanu ‘ia nā Uluwehi no ia Wao.

Input 2: E like no hoi me ka mea i hoike akea ia ae no ka manawa a me ka la e weheia ai a e kanuia ai hoi o na pua a me na mea ulu e ae ma kahi nona ka inoa kilakila maluna ae, pela no i hoea io mai ai i ka Poaha iho la, hora 9 A. M. a mahope mai. Mamua ae o ia manawa, ua lehulehu na poe i pii aku me na mea kanu, maluna o na kaa a malalo no hoi.

Prediction 2: E like nō ho‘i me ka mea i hō‘ike ākea ‘ia a‘e no ka manawa a me ka lā e wehe ‘ia ai a e kanu ‘ia ai ho‘i _o nā pua a me nā mea ulu ‘ē a‘e ma kahi nona ka inoa kilakila ma luna a‘e, pēlā nō i hō‘ea ‘i‘o mai ai i ka Pō‘ahā iho la, hora 9 A. M. a ma hope mai. Ma mua a‘e o ia manawa, ua lehulehu nā po‘e i pi‘i aku me nā mea kanu, ma luna o nā ka‘a a ma lalo nō ho‘i.

Ground-truth 2: E like nō ho‘i me ka mea i hō‘ike ākea ‘ia a‘e no ka manawa a me ka lā e wehe ‘ia ai a e kanu ‘ia ai ho‘i o nā pua a me nā mea ulu ‘ē a‘e ma kahi nona ka inoa kilakila ma luna a‘e, pēlā nō i hō‘ea ‘i‘o mai ai i ka Pō‘ahā iholā, hola 9 A. M. a ma hope mai. Ma mua a‘e o ia manawa, ua lehulehu nā po‘e i pi‘i aku me nā mea kanu, ma luna o nā ka‘a a ma lalo nō ho‘i.

Input 3: Hoomakaia ke Kanu Ana.

Prediction 3: Ho‘omaka ‘ia ke Kanu ‘Ana.

Ground-truth 3: Ho‘omaka ‘ia ke Kanu ‘Ana.

Input 4: Ua hoea ae no ilaila ka Puali Puhiohe Lahui, a i ka aneane ana ae i ka manawa, a i ole ia, ua hala no paha he hapalua hora mahope iho o ka hora 9, ua uhene mai la lakou i ke mele Liliuokalani, a o ka wa no ia o Kamalii Kawanānakoā, ma ka aoao o ke Aliiāimoku, i kanu iho ai i kekahi kumu lehua o Mokaulele iwaenakonu, i hoopuniia ae me na ohawai a me kekahi mau mea kanu Hawaii e ae iloko o kekahi ponaha poepoe, a makai iho hoi o ia wahi i kanu ai o Kamalii Kalanianaole i kekahi kumu lehua ahihi ma ka aoao o ke Alii ka Moiwahine Kanemake. Pau keia mau hana ae la, ua noa i na mea a pau, a ua hele no hoi ia wahi a eeu i na oiwi palupalu o kakou, e kanu ana i kela a me keia mea, a he mau oiwi oolea no hoi kekahi malaila e kokua ana. He mea makehewa paha ke helupapa aku i na pua i kanuia. Hookahi a makou mea i mahalo, oia no kekahi wahi i kanu mua e ia no makai iki mai o ka puka komo, me ka inoa o ia kihapai e kau ae la maluna, a he ku maoli no i ka nani.

Prediction 4: Ua hō‘ea a‘e nō i laila ka Pū‘ali Puihi ‘ohe Lāhui, a i ka ‘ane‘ane ‘ana a‘e i ka manawa, a i ‘ole ia, ua hala nō paha he hapalua hora ma hope iho o ka hora 9, ua ‘uhene mai la lakou i ke mele Lili‘uokalani, a ‘o ka wā nō ia o Kāmali‘i Kawānakoā, ma ka ‘ao‘ao o ke Ali‘i‘ai moku, i kanu iho ai i kekahi kumu lehua o Mōkaulele i waenakonu, i ho‘opuni ‘ia a‘e me nā ‘ohā wai a me kekahi mau mea kanu Hawai‘i ‘ē a‘e i loko o kekahi pōnaha poepoe, a ma kai iho ho‘i o ia wahi i kanu ai ‘o Kāmali‘i Kalaniana‘ole i kekahi kumu lehua ‘āhihi ma ka ‘ao‘ao o ke Ali‘i‘i, ka Mō‘i_wahine Kāne make. Pau kēia mau hana a‘e_lā, ua noa i nā mea a pau, a ua hele nō ho‘i ia wahi a ‘e‘eu i nā ‘ōiwi palupalu o kākou, e kanu ana i kēlā a me kēia mea, a he mau ‘ōiwi ‘o‘ole‘a nō ho‘i kekahi ma laila e kōkua ana. He mea makehewa paha ke helu papa aku i nā pua i

kanu 'ia. Ho'okahi a mākou mea i mahalo, 'o ia nō kekahi wahi i kanu mua 'ē 'ia nō ma kai iki mai o ka puka komo, me ka inoa o ia kīhāpai e kau a'e 'lā ma luna, a he kū maoli nō i ka nani.

Ground-truth 4: Ua hō'ea a'e nō i laila ka pū'ali puhi 'ohe lāhui, a i ka 'ane'ane 'ana a'e i ka manawa, a i 'ole ia, ua hala nō paha he hapalua hola ma hope iho o ka hola 9, ua 'uhene maila lākou i ke mele Lili'uokalani, a 'o ka wā nō ia o kamāli'i Kawānanakoa, ma ka 'ao'ao o ke ali'i 'ai moku, i kanu iho ai i kekahi kumu lehua o Mokaulele i waenakonu, i ho'opuni 'ia a'e me nā 'ōhāwai a me kekahi mau mea kanu Hawai'i 'ē a'e i loko o kekahi pōnaha poepoe, a ma kai iho ho'i o ia wahi i kanu ai 'o kamāli'i Kalaniana'ole i kekahi kumu lehua 'āhihi ma ka 'ao'ao o ke ali'i, ka mō'i wahine kāne make. Pau kēia mau hana a'ela, ua noa i nā mea a pau, a ua hele nō ho'i ia wahi a 'e'eu i nā 'ōiwi palupalu o kākou, e kanu ana i kēlā a me kēia mea, a he mau 'ōiwi 'o'ole'a nō ho'i kekahi ma laila e kōkua ana. He mea makehewa paha ke helu papa aku i nā pua i kanu 'ia. Ho'okahi a mākou mea i mahalo, 'o ia nō kekahi wahi i kanu mua e ia nō ma kai iki mai o ka puka komo, me ka inoa o ia kīhāpai e kau a'ela ma luna, a he kū maoli nō i ka nani.

Input 5: Ka Lanai Pea Ahaaina.

Prediction 5: Ka Lāna'i Pe'a 'Aha'aina.

Ground-truth 5: Ka Lānai Pe'a 'Aha'aina.

Input 6: Mauka iki aku hoi o keia mala, ma ka aoao maluna o ka hale noho o J. Mana, ua kukuluia ae la he lanai pea, a ilaila kahi i hoomakaukauia ai o kekahi papaaina. E houhene mau ana no hoi ka puahi puhiohe i kela a me keia wa a hiki i ka wa ai, i ke kau pono o ka la i ka piko a mahope iho paha, me ka hoomau aku no i ka hoolealea ana a pau ka papaaina mua a ke Alii, o na Kamaliikane no ma kona a me ko kona kaikoeko wahi. Ekolu papaaina i hoonohoia ai a ua ai na poe a pau i hiki aku a lawa pono, me ka hoounaia o kahi mea-ai no kekahi poe i maopopo i hiki ole aku.

Prediction 6: Ma uka iki aku ho'i o kēia māla, ma ka 'ao'ao ma luna o ka hale noho o J. Mānā, ua kūkulu 'ia a'e 'la he lānai pe'a, a i laila kahi i ho'omākaukau 'ia ai 'o kekahi papa_ 'aina. E ho'uhene mau ana nō ho'i ka pū'ali puhi 'ohe i kēlā a me kēia wā a hiki i ka wā 'ai, i ke kau pono o ka lā i ka piko a ma hope iho paha, me ka ho'omau aku nō i ka ho'ole'ale'a 'ana a pau ka papa 'aina mua a ke Ali'i, 'o nā Kamāli'i_ kāne nō ma kōna a me kō kōna kaiko'eke wahi. 'Ekolu papa 'aina i ho'onoho 'ia ai a ua 'ai nā po'e a pau i hiki aku a lawa pono, me ka ho'ouna 'ia o kahi mea_ 'ai no kekahi po'e i maopopo i hiki 'ole aku.

Ground-truth 6: Ma uka iki aku hoi o kēia māla, ma ka 'ao'ao ma luna o ka hale noho o J. Mana, ua kūkulu 'ia a'ela he lānai pe'a, a i laila kahi i ho'omākaukau 'ia ai o kekahi papa 'aina. E ho'uhene mau ana nō ho'i ka pū'ali puhi 'ohe i kēlā a me kēia wā a hiki i ka wā 'ai, i ke kau pono o ka lā i ka piko a ma hope iho paha, me ka ho'omau aku nō i ka ho'ole'ale'a 'ana a pau ka papa 'aina mua a ke ali'i, o nā kamāli'i kāne nō ma kōna a me kō kōna kaiko'eke wahi. 'Ekolu papa 'aina i ho'onoho 'ia ai a ua 'ai nā po'e a pau i hiki aku a lawa pono, me ka ho'ouna 'ia o kahi mea'ai no kekahi po'e i maopopo i hiki 'ole aku.

Input 7: O na keiki puhiohe no hoi kekahi i kanu mau wahi mea kanu, a he mau mea ulu Hawaii wale no hoi ka lakou o ke ano papa kahuna. Mamuli o ka oluolu a me ka lokomaikai o kahi o lakou, ua loa mai ia makou ka papa hoike a me na wehewehe ana o ka lakou.

Prediction 7: 'O nā keiki puhi 'ohe nō ho'i kekahi i ka nū mau wahi mea kanu, a he mau mea ulu Hawai'i wale nō ho'i kā lākou o ke 'ano papa kahuna. Ma muli o ka 'olu'olu a me ka lokomaika'i o kekahi o lākou, ua loa'a mai iā mākou ka papa hō'ike a me nā wehewehe 'ana o kā lākou.

Ground-truth 7: 'O nā keiki puhi 'ohe nō ho'i kekahi i kanu mau wahi mea kanu, a he mau mea ulu Hawai'i wale nō ho'i ka lākou o ke 'ano papa kahuna. Ma muli o ka 'olu'olu a me ka lokomaika'i o kekahi o lākou, ua loa'a mai iā mākou ka papa hō'ike a me nā wehewehe 'ana o kā lākou.

Input 8: Mau Mea Kanu Hookalakupua

Prediction 8: Mau Mea Kanu Ho'okalakupua.

Ground-truth 8: Mau Mea Kanu Ho'okalakupua.

Input 9: o ke ano hookahuna oiaio maoli no. Eia iho no ia papa hoike:

Prediction 9: 'o ke 'ano ho'okahuna 'oia'i'o maoli nō. Eia iho nō ia papa hō'ike:

Ground-truth 9: 'O ke 'ano ho'okahuna 'oia'i'o maoli nō. Eia iho nō ia papa hō'ike:

Input 10: hala polapola. Ko puni e Kalani o ka lei e, leiia hoi o Halaomapuana, onaona i ka ihu, huihui ke hanu iho.

Prediction 10: hala polapola. Kō puni ē Ka_lani o ka lei ē, lei 'ia ho'i 'o Hālaomapuana, onaona i ka ihu, hu'ihu'i ke hanu iho.

Ground-truth 10: hala polapola. Kō puni e Kalani o ka lei ē, lei 'ia ho'i 'o Halaomāpuana, onaona i ka ihu, hu'ihu'i ke hanu iho.

First 10 sentences in Newspaper 2:

Input 1: O Kahahana, he alii kapu ia no o Oahu.

Prediction 1: 'O Kahahana, he ali'i kapu ia nō o O'ahu.

Ground-truth 1: 'O Kahahana, he ali'i kapu ia no O'ahu.

Input 2: O Kaionuilalahai ka makuahine, ka moopuna a Kalaniomaiheuila, ke kaikamahine a Kalanikahimakaiialii, a laua o Kualu ke kaikuahine, a mua hoi o Kaulahea ka Moi o Maui.

Prediction 2: 'O Ka'ionuilalahai ka makuahine, ka mo'opuna a Kalani'o ma'i he uila, ke kaikamahine a Kalani_kahi_maka'i ali'i, 'ā lāua 'o Kū'alu ke kaikuahine, a mua ho'i o Ka'ula_hea, ka Mō'i o Maui.

Ground-truth 2: 'O Ka'ionuilalahai ka makuahine, ka mo'opuna a Kalani'ōmaiheuila, ke kaikamahine a Kalanikahimaka'iali'i, a lāua 'o Kū'alu ke kaikuahine, a mua ho'i o Ka'ulahea, ka Mō'i o Maui.

Input 3: A o ka makuakane, oia hoi o Elani, no ka ohana a Kupanihi, a o Keōpūolani.

Prediction 3: A 'o ka makua kāne, 'o ia ho'i 'o 'Elani, no ka 'ohana a Kupanihi, a 'o Keōpūolani.

Ground-truth 3: A 'o ka makua kāne, 'o ia ho'i 'o 'Ēlani, no ka 'ohana a Kūpānihi, a 'o Keōpūolani.

Input 4: I ko Kahahana manawa kamalii, ua kii ia mai e Kahekili e lawe i Maui i keiki nana.

Prediction 4: I ko Kahahana manawa kamali'i, ua ki'i 'ia mai e Kahekili e lawe i Maui i keiki nāna.

Ground-truth 4: I ko Kahahana manawa kamali'i, ua ki'i 'ia mai e Kahekili e lawe i Maui i keiki nāna.

Input 5: Ua hanai kapu ia oia i Maui.

Prediction 5: Ua hānai kapu 'ia 'o ia i Maui.

Ground-truth 5: Ua hānai kapu 'ia 'o ia i Maui.

Input 6: A i kona lilo ana ae i kanaka makua, ua lilo oia i kanaka maikai, a ua nani hoi kona helehelena, a ua piipii maikai kona lauoho; a ua kapaia oia i kekahi wa, he piipii hahai moa.

Prediction 6: A i kona lilo 'ana a'e i kanaka mākuua, ua lilo 'o ia i kanaka maika'i, a ua nani ho'i kona helehelena, a ua pi'ipi'i maika'i kona lauoho; a ua kapa 'ia 'o ia i kekahi wā, he pi'ipi'i hahai moa.

Ground-truth 6: A i kona lilo ‘ana a’e i kanaka makua, ua lilo ‘o ia i kanaka maika‘i, a ua nani ho‘i kona helehelena, a ua pi‘ipi‘i maika‘i kona lauoho; a ua kapa ‘ia ‘o ia i kekahi wā, he pi‘ipi‘i hahai moa.

Input 7: I kona lilo ana i kanaka makua, hooipoipo aku la oia me na wahine kaukualii, nolaila, ua ali ia ke kapu, a ua kapa hou ia ka inoa o Walia, a ua hoopauia ka inoa Ahi, Wela, Hahana.

Prediction 7: I kona lilo ‘ana i kanaka mākua, ho‘oipoipo aku la ‘o ia me nā wāhine kaukuali‘i, no laila, ua ‘ali ‘ia ke kapu, a ua kapa hou ‘ia ka inoa ‘o Wali‘a, a ua ho‘opau ‘ia ka ‘ino a ‘Ahi, Wela, Hahana.

Ground-truth 7: I kona lilo ‘ana i kanaka makua, ho‘oipoipo akula ‘o ia me nā wāhine kaukuali‘i, no laila, ua ‘ali ‘ia ke kapu, a ua kapa hou ‘ia ka inoa ‘o Walia, a ua ho‘opau ‘ia ka inoa Ahi, Wela, Hahana.

Input 8: Lawe ae la o Kahahana i wahine nana, o Kekuapo i ka inoa.

Prediction 8: Lawe a‘e la ‘o Kahahana i wahine nāna, ‘o Kekuapo‘i ka inoa.

Ground-truth 8: Lawe a‘ela ‘o Kahahana i wahine nāna, ‘o Kekuapo‘i ka inoa.

Input 9: Ua olelo ia no hoi kela wahine, aohe ona lua iloko o ke aupuni Hawaii nei, a ua kaulana kona inoa mai Hawaii a Kauai, a ua lilo hoi i kaa, a ua kiekie hanohano kona kino; aole no hoi i ike ia kahi kina mai luna o ke poo a hiki i na kapuai wawae; he mau maka manu nunu kona i like me ko ka mohoea; a ua like hoi na helehelena o na hiohiona papalina me ka opuu rose i mohala maikai i ke kakahiaka; a ua nui na loli o kona mau hiohiona, i ke kakahiaka, i ke awakea, ahiahi a me ka po; o ka puo kelakela, oiiai oia iloko o ia manawa, ua kupono ke ali ana o ke kapu.

Prediction 9: Ua ‘olelo ‘ia nō ho‘i kēlā wahine, ‘a‘ohe ona lua i loko o ke aupuni Hawai‘i nei, a ua kaulana kona inoa mai Hawai‘i a Kaua‘i, a ua lilo ho‘i i ka‘ao, a ua ki‘eki‘e hanohano kona kino; ‘a‘ole nō ho‘i i ike ‘ia kahi ‘kīna mai luna o ke po‘o a hiki i nā kapua‘i wāwae; he mau maka manu nūnū kona i like me ko ka mohoe ‘ā; a ua like ho‘i nā helehelena o nā hi‘ohi‘ona pāpālina me ka ‘ōpu‘u rose i mōhala maika‘i i ke kakahiaka; a ua nui nā loli o kona mau hi‘ohi‘ona, i ke kakahiaka, i ke awakea, ahiahi a me ka pō; ‘o ka pu‘ō kelakela, ‘oiiai ‘o ia i loko o ia manawa, ua kūpono ke ‘ali ‘ana o ke kapu.

Ground-truth 9: Ua ‘olelo ‘ia nō ho‘i kēlā wahine, ‘a‘ohe ona lua i loko o ke aupuni Hawai‘i nei, a ua kaulana kona inoa mai Hawai‘i a Kaua‘i, a ua lilo ho‘i i ka‘ao, a ua ki‘eki‘e hanohano kona kino; ‘a‘ole nō ho‘i i ike ‘ia kahi kīnā mai luna o ke po‘o a hiki i nā kapua‘i wāwae; he mau maka manu nūnū kona i like me ko ka mohoea; a ua like ho‘i nā helehelena o nā hi‘ohi‘ona pāpālina me ka ‘ōpu‘u loke i mōhala maika‘i i ke kakahiaka; a ua nui nā loli o kona mau hi‘ohi‘ona, i ke kakahiaka, i ke awakea, ahiahi a me ka pō; ‘o ka pu‘ō kelakela, ‘oiiai ‘o ia i loko o ia manawa, ua kūpono ke ‘ali ‘ana o ke kapu.

Input 10: Ua uluhua na‘lii, na kahuna a me na makaainana o ke aupuni o Oahu i ko lakou Moi ia Kumahana, i ke keiki a Peleiōholani.

Prediction 10: Ua uluhua nā ‘‘ali‘i, nā kāhuna a me nā maka‘āinana o ke aupuni o O‘ahu i ko lākou Mō‘ī, iā Kūmahana, i ke keiki a Peleiōhōlani.

Ground-truth 10: Ua uluhua nā ali‘i, nā kāhuna a me nā maka‘āinana o ke aupuni o O‘ahu i ko lākou Mō‘ī, iā Kūmahana, i ke keiki a Peleiōhōlani.