

Modelling Verbal Morphology in Nen

Saliha Muradođlu^{ΩΦ} Nicholas Evans^{ΩΦ} Ekaterina Vylomova^μ

^ΩThe Australian National University (ANU) ^μThe University of Melbourne

^ΦARC Centre of Excellence for the Dynamics of Language (CoEDL)

saliha.muradgolu@anu.edu.au, nicholas.evans@anu.edu.au,

ekaterina.vylomova@unimelb.edu.au

Abstract

Nen verbal morphology is remarkably complex; a transitive verb can take up to 1,740 unique forms. The combined effect of having a large combinatoric space and a low-resource setting amplifies the need for NLP tools. Nen morphology utilises distributed exponence – a non-trivial means of mapping form to meaning. In this paper, we attempt to model Nen verbal morphology using state-of-the-art machine learning models for morphological reinflection. We explore and categorise the types of errors these systems generate. Our results show sensitivity to training data composition; different distributions of verb type yield different accuracies (patterning with E-complexity). We also demonstrate the types of patterns that can be inferred from the training data through the case study of syncretism.

1 Introduction

A long-standing research direction in NLP targets the development of robust language technology applicable across the wide variety of the world’s languages. Unfortunately, the vast majority of machine learning models are being developed for a small fraction of nearly 7,000 languages in the world, such as English, German, French, or Chinese. With introduction of highly multilingual corpora such as UniversalDependencies (Nivre et al., 2016) and UniMorph (Sylak-Glassman et al., 2015; Kirov et al., 2018) the situation started to change. For instance, SIGMORPHON organized a number of shared tasks on morphological reinflection starting from 10 languages in 2016 (Cotterell et al., 2016) and up to 90 languages in 2020 (Vylomova et al., 2020). In 2020, languages were sampled from various typologically diverse families: Indo-European, Oto-Manguean, Tungusic, Turkic, Niger-Congo, Bantu, and others. Still, just one language, namely, Murrinh-patha, an Australian Aboriginal language (Mansfield, 2019), represented the

whole linguistic variety of the Oceania region. In this paper, we aim at filling the gap by exploring Nen, a Papuan language spoken by approximately 400 people in Papua New Guinea. Nen is known for its rich verbal morphology, with a transitive verb inflecting for up to 1,740 feature combinations. Distributed exponence, the phenomenon which gives rise to this large paradigm size, provides insight into modelling complex mappings between surface forms and feature bundles.

We conduct a series of experiments on morphological reinflection task recently introduced under the umbrella of SIGMORPHON (Cotterell et al., 2016, 2018). We train several state-of-the-art machine learning models for verbal inflection in Nen and provide an extensive error analysis. We investigate the relationship between the distribution of verb type (inflection classes) in the data and performance. Finally, we show that the system learns properties of the data that are not explicitly given, but may be inferred.

The rest of the paper is organized as follows: In Section 2, we give a brief overview of related work. Section 3 provides an overview of Nen verbal morphology, Section 4, details our methodology, and Section 5 presents our results. Finally, Section 6 concludes the paper.

2 Related Work

Muradoglu et al. (2020) is the only reported work on the computational modelling of the Nen language. Similar to this study, the main focus is on modelling Nen verbal morphology, but using finite-state architecture instead. The accuracy achieved by the FST system is 80.3% obtained across the corpus, with approximately 10% of the accuracy attributable to the modelling of prefixing verbs (the regularity of copula verbs boosts the accuracy from 70.5%). The accuracies reported are not directly

comparable with those presented here due to the different data splits, and increased amount of data.

In our error analysis, we follow the error taxonomy proposed by Gorman et al. (2019) upon a detailed analysis of typical errors produced by morphologically inflection systems. A similar study was conducted for Tibetan (Di et al., 2019).

3 The Nen Language

Nen is a Papuan language of the Morehead-Marô (or Yam) family, located in the southern part of New Guinea (Evans, 2017). It is spoken in the village of Bimadbn in the Western Province of Papua New Guinea, by approximately 400 people, for which it is a primary language (Evans, 2015, 2020). Most inhabitants are multilingual, typically speaking several of the neighbouring languages.

The subject of this paper – verbs – are the most complicated word-class in Nen (Evans, 2015, 2019b). They are demarcated into three separate categories: prefixing, middle, and ambifixing verbs. The latter two are mostly regular in terms of morphophonological rules. In the remainder of this section, we elaborate on these characteristics, to give the reader enough background to follow the discussion in subsequent sections.

3.1 Verbal morphology

We begin our description from the maximal case – transitive *ambifixing* verbs. Examples of this verb type include *yis* ‘to plant’ and *waprs* ‘to do’. These verbs allow for full prefixing and suffixing possibilities. Evans (2016) provides the canonical paradigms for the undergoer prefixes, thematic and desinences. Suffix combinations are constructed by concatenating the corresponding thematic and the desinence. Between the undergoer prefix and verb stem is a directional prefix slot, available for all verb types. This slot is occupied by $\{-n-\}$ ¹ to convey a ‘towards’, $\{-ng-\}$ for ‘away’ or left empty to convey a directionally neutral semantic.

Middle verbs such as *owabs* ‘to speak’ or *an̄gs* ‘to return’, are also ambifixing, but the prefixal slot is restricted to $\{n-\}$ (α -series), $\{k-\}$ (β -series), $\{g-\}$ (γ -series). These prefixes are person and number invariant, and mark the verb as being a dynamic monovalent verb. The prefix set is divided through the use of arbitrarily labels: α , β , and γ . These

¹We follow linguistic convention with ‘{ }’ denoting morphemes, and examples are italicised.

dummy indices do not carry specific semantic values until they are unified with other TAM (Tense, Aspect, and Mood) markings on the verb (Evans, 2015).

Prefixing verbs have separate closed paradigms, tailored to the subtype. Prefixing verbs are mostly distinguished through semantics; positional verbs such as *kmangr* ‘to be lying down’, the verb ‘to own/have’ *awans*, the verb ‘to walk’ *tan* and the copula verb *m* with its directional variants (be hither (i.e. come) or be thither (go)).

Inflectional prefixes for these verbs, mostly resemble the process with ambifixing verbs, yet the suffixes are limited. Of the 50 or so prefixing verbs, the vast majority are positional (Evans, 2020). An additional distinguishing feature of prefixing verbs, is the lack of infinitives. Both ambifixing and middle verbs form infinitives through suffixing *-s* to the verb stem. For the purposes of this study, we have listed the prefixing verb lemmas as the verb stem.

Methodologically, it is more convenient to segment a word as a classical bijective mapping between form to meaning. However, the Nen verbal system distributes information in a more complicated way. The prefixes (undergoer and future imperative) and suffixes (thematic and desinence) are not independent values. Nen verbal morphology is characterised by *distributed exponence (DE)*; “morphosyntactic feature values can only be determined after unification of multiple structural positions” (Carroll, 2016).

There are two consequences for morphological parsing:

- a) Provisional unspecified values occur regularly, whether
 - (i) These involve partial specification that will be filled in later in the word-parse, such as the left-edge prefix $\{yaw-\}$ (1st person non-singular undergoer), which will only be made more precise in its number value (dual, or plural) when the thematic is encountered after the verb stem: thus *yaw-aka-t-an* ‘I see them² (more than two)’, where the ‘non-dual’ marker $\{-ta-\}$ eliminates the dual (them two) but *yaw-akae-w-n* ‘I see them

²Can also mean ‘I see you (more than two)’, resolved by combining with an appropriate free pronoun, *bm* ‘you (absolute)’, but for present purposes we ignore this further complication.

(two)', where the 'dual thematic' {-w-} eliminates the plural (them more than two) reading.

- (ii) These involve semantically-unspecified prefix series which only acquire meaning when they are combined with suffixes at the other end of the word: thus {yaw-}, in the above example, belongs to the α -series which, if it combines with the 'basic imperfective', will be given a (broadly) non-past reading, but when it combines with the 'past perfective' it will be given a past reading and when it combines with a 'projected imperative' it will be given a future meaning; a β -series form like {taw-}, by contrast, will have a 'yesterday past' interpretation when combining with the 'basic imperfective' suffixes but when combining with imperatives it will have a 'now/immediate command' meaning

- b) More problematically, prefixes that normally have one reading (such as the yaw-example just discussed, which normally marks second/third person non-singular objects) sometimes have to be given a different meaning (e.g. large plural intransitive subjects) if further parsing to the right encounters a 'middle' rather than a 'transitive dynamic' stem (Evans 2017, 2019).

In principle that this means left-to-right morphological parsing is sometimes non-monotonic (particularly in the case of (b)), so that semantic values, as parsing proceeds, need to be sometimes held as provisionally unspecified, sometimes as partially specified, and sometimes as specified but subject to later override.

3.2 Distributed Exponence

One of the primary motivations for choosing Nen as a case study is the phenomenon that gives rise to this combinatorial power: distributed exponence. Essentially distributed exponence is a morphological phenomenon that gives rise to some types of non-monotonicity.

In linguistics, the notion of extended exponence was first introduced by Mathews (1974) and is now commonly referred to as multiple exponence (ME). Mathews defined ME as a category that would have exponents in two or more distinct positions.

Distributed exponence is a kind of ME, which involves the use of more than one morphological segment to convey meaning. It requires all relevant morphs to yield a precise interpretation of the feature value in question (Carroll, 2016; Harris, 2017).

- (1) n-ng-owan-t-e
M: α -VEN-set.off-ND:IPF.NP-
IPF.NP.2|3SGA
'You/(s)he are/is setting off.'³

In the example above, no one marker marks the singular person. The information of the agent being singular is distributed across the thematic (dual/non-dual) and the desinence (single/dual/plural). If a non-dual thematic is present then the desinence cannot have dual features; the only options are singular or plural. Another morpheme present in this example is the prefix -ng- which marks the verb with the directional *thither*. The prefix n- marks this verb as a middle verb; it reduces the valency of the verb and yields information about the membership of the class α . Together with the prefix, thematic and desinence, the TAM feature can be obtained.

4 Methodology

4.1 Morphological inflection task

Morphological inflection is a task of predicting a target word form from a corresponding word lemma and a set of morphosyntactic features (specifying the target slot, e.g. its part of speech (POS), tense, number, gender). For instance, a system is provided with a lemma "to sing" and a set of tags "Verb; Past" and needs to generate "sang". Morphological *reinflection* is a variation of the task when a lemma form is replaced with some other form and (optionally) its tags. The task has been traditionally solved with finite-state transducers, either hand-engineered (Koskenniemi, 1983; Kaplan and Kay, 1994) or trainable models that rely on both expert knowledge and data (Mohri, 1997; Eisner, 2002). In 2016 SIGMORPHON started a series of shared tasks on morphological reinflection, and neural models demonstrated superior performance when compared to finite-state or rule-based approaches, especially in high-resource languages (Cotterell et al., 2016; Vylomova et al., 2020).

³Example adapted from (Evans, 2020)

4.2 Data

The data used in this study comes from a Nen verb corpus (approximately 6,000 verb samples representing 2,231 unique inflected forms) created by [Muradoğlu \(2017\)](#). This dataset is a distilled subset from the approximately 8-hour natural speech corpus for the Nen language. As such it entails a frequency sorted list of all the verb forms occurring.

The training data is a set of triples comprising a lemma, morphosyntactic features, and an inflected form (i.e. we will only focus on morphological inflection).

Sampling Following the methodology in [Cotterell et al. \(2018\)](#) we split the data into training, development, and test sets. Training splits were created by sampling without replacement for three set sizes: all (*ALL*), medium (*MR*), and low (*LR*).

In virtue of coming from a natural corpus, the list of verb forms we use is Zipfian. This study does not distinguish between the feature bundles and only considered surface (inflected) forms. To facilitate the nature of our study, we uniformly distribute frequency across each syncretic cell.

For the *ALL* training set we start by sampling the first 1,931 forms, in accordance with the Zipfian ranking across the corpus. In other words, we sample the 1,931 most frequent verb forms. We randomly shuffle the remaining 300 forms into a 200 form test, and 100 form development (dev) sets. The test and dev sets remain the same through this experiment. Zipfian sampling is considered more realistic in this case, as it mimics the stimulus a language learner encounters. The dev and test set are randomly shuffled since supervised methods usually generalise from frequently encountered words.

For the *LR* and *MR* settings we take the first 100 and 1,000 forms from the *ALL* training set, respectively. In addition, we create a high-resource (*HR*) set by supplementing the *ALL* set with synthetic forms, the final set contains 10,000 forms. In order to generate synthetic samples, we use data hallucination technique proposed in [Anastasopoulos and Neubig \(2019\)](#). Note that the low-resource (*LR*) training set is a subset of the medium-resource (*MR*), which is superseded by the *ALL* (and by extension the high-resource (*HR*) data set).

Finally, we contrast Zipfian sampling, when forms are sampled based on their frequency, to random sampling. Both sets (*LR* and *MR*) for the

random sampling are created in a similar manner to Zipfian sampling, except frequency is not considered. Note that due to initial data size constraints, the *ALL* (and, therefore, *HR*) data sets for *both* the Zipfian and random sampling are the same.⁴

4.3 Experiments

In the current study we conducted three experiments to address our research questions.

4.3.1 Experiment 1: Testing across various data sizes and sampling methods

Research Question: How does training size and sampling method affect the models' performance, and what kind of errors are likely across these conditions?

We evaluate modelling accuracies across four different training sizes, which is further contrasted across sampling type. Our experimental setup mirrors those of the SIGMORPHON reinflection tasks ([Cotterell et al., 2016, 2017, 2018](#); [Vylomova et al., 2020](#)): given an input lemma and a set of feature tags, models generate inflected forms. The final accuracy is computed as the percentage of matches between the gold and predicted forms.

4.3.2 Experiment 2: Testing compositionality of training data

Research Question: Does the composition of the training data affect the resultant accuracies, and, if so, how?

We test the effects of the verb type composition (i.e. how much of each verb type there is) in the training set. This study consists of seven (arising from all combinations of the three verb types) training data sets obtained through the sampling methods outlined above. We compare training sets of ambifixing verbs only, prefixing verbs only, middle verbs only, a two-way combination of each verb class: ambifixing and prefixing verbs, ambifixing and middle verbs, and prefixing and middle verbs and, finally an equal distribution of all three verb types, as listed in Table 4. Each set contains 386 forms (instances), stipulated by the amount of prefixing verbs available. The test and development set are 100 forms each, and is made up of 34 ambifixing, 33 middle and 33 prefixing verbs⁵

⁴Since the test and dev set are the same for both sampling methods, and are generated from the **remaining** 300 tokens (i.e. the least frequent items), it renders the random sampling of the *ALL* (and thus *HR*) the same.

⁵Uniform distribution is unlikely in natural language, in fact, [Muradoğlu \(2017\)](#) shows that the distribution is skewed

4.3.3 Experiment 3: Testing syncretism

Research Question: Do the models infer properties of the language which are not annotated in the data?

In Nen, the second and third person feature bundles often correspond to the same surface form across the available TAM categories (i.e. are syncretic). We test the likelihood of both models predicting the *unseen* second person singular for the past perfective TAM category as syncretic with the *seen* third-person singular variant. This is the one instance across the Nen verbal paradigm where this syncretism does not hold. In essence, we examine linguistic patterns that may be inferred from an annotated dataset.

The main focus here, is to categorise the type of prediction rather than the overall accuracy, as such training and development sets are identical to those generated for the *ALL* setting in the first experiment. The test set is comprised of 100 inflections of the past perfective second singular tags, most of these have been gathered from the Nen dictionary (Evans, 2019a).

4.4 Models

For our experiments, we will utilise two models that have shown superior performance in SIGMORPHON-CoNLL 2017 Shared Task on morphological reinflection in low- and medium-resource settings (Cotterell et al., 2017). Both of them are essentially neural sequence-to-sequence models implemented in Dynet (Neubig et al., 2017). In addition, we also compare the results with a simple non-neural baseline used in 2017–2018 tasks on morphological reinflection (Cotterell et al., 2017, 2018).

Hard Monotonic Attention (Aharoni and Goldberg, 2017) An external aligner (Sudoh et al., 2013) first produces transformation operations between an input (lemma) and a target (inflected form) character sequences. The alignment operations (steps) are then fed into a neural encoder–decoder model. The network, therefore, is trained to mimic the transformation steps, and at inference time it predicts the actions based on the input (lemma) sequence. Unlike soft attention models, this model attends to a single input state at each step and either writes a symbol to the output sequence

to favour a higher number of ambifixing verbs in terms of the number of inflected forms.

or advances its pointer to the next state. Hard attention models demonstrate superior performance in languages that employ suffixing morphology with stem changes.

Neural Transition-based (Makarov and Clematide, 2018) The model is essentially derived from Aharoni and Goldberg (2017) by enriching it with explicit insertion, deletion or, alternatively, copy mechanisms. The copy mechanism led to significant accuracy gains in low-resource settings. Following Rastogi et al. (2016), the model can be seen as a neural parameterization of a weighted finite-state machine.

Non-neural Baseline (Cotterell et al., 2017, 2018) The non-neural system first aligns lemma and inflected form strings using Levenstein distance (Levenstein, 1966) and then extracts prefix- and suffix-based transformation rules.

4.5 Settings

The hyperparameters of the models are set to the values reported in the corresponding papers as per Table 1.

Hyperparameters	A&G	M&C
Input dim	100	100
Hidden dim	100	100
Epochs	100	50
Layer	2	1

Table 1: Hyperparameters for both A&G (2017) and M&C (2018) models.

5 Results

Table 2 shows the accuracies achieved for each system for each training set size and sampling type from Experiment 1. For all setups the M&C model performed best with random sampling (where applicable). As expected the high-resource setting performs best overall. The random sampling yields slightly higher accuracies than the Zipfian counterpart, this is likely due to the fact that prefixing verbs, particularly the copula and its 40 distinct forms occupy a majority of the top 100 positions in the Zipfian distribution. Thus when random sampling is utilized the training set includes more examples of ambifixing verbs.

	A&G 2017		M&C 2018		Non-Neural baseline (NNB)	
	Random	Zipf	Random	Zipf	Random	Zipf
HR	0.610		0.650		0.015	
ALL	0.390		0.510		0.010	
MR	0.295	0.285	0.445	0.420	0.000	0.000
LR	0.020	0.005	0.080	0.030	0.010	0.010

Table 2: Data set, model and sampling accuracies. ALL is a total of 1,931 verbs, HR is 10,000, MR is 1,000 and LR is 100 samples for the training set.

	ALL			HR			MR			LR		
	A&G	M&C	NNB	A&G	M&C	NNB	A&G	M&C	NNB	A&G	M&C	NNB
Allomorphy	56	55	190	54	46	144	61	77	188	17	162	190
Free Variation	30	24	0	14	15	11	13	24	0	0	2	0
Target	8	8	8	8	8	8	8	8	8	8	8	8
Stem	28	11	0	2	1	5	61	7*	2	174†	22	0
Total	122	98	198	78	70	168	143	116	198	199	194	198

Table 3: Absolute number of errors on the test set (200 instances) made by each system trained in ALL, HR, MR and LR setting. *contains 5 looping errors, † 17 looping errors.

5.1 Error Analysis

We analysed the errors produced in prediction following the taxonomy laid out by Gorman et al. (2019); Di et al. (2019).

We have taken a hierarchical approach to our error classification; whereby if more than one error is present, the category higher up is reported. For example, if a predicted form exhibits both target and allomorphy errors (error types are described in the following subsections), then only the target error is reported. The motivation for this lies in the nature of the error; free variation is technically not even an error. By contrast, misapplication of a morphophonological rule does indeed yield an incorrect form. Additionally, we have marked Target errors higher up as the system cannot be expected to correctly predict a form if the gold standard is incorrect. The hierarchy is as follows: Target>Stem>Allomorphy>Free Variation. 3 Table 3 summarises the types of errors across the different training sizes for each model. Overall, for both systems allomorphy errors remain relatively unimproved between the ALL and HR setting, but show a leap of reduction from the LR to MR conditions. Free variation errors are more prevalent in the ALL setting. This is probably a consequence of seeing more of the golden data and thus observing more of the systematic variations. This also explains why these errors reduce in number for

the HR setting. The target errors are consistent across each experiment, as these are systematic issues with the gold data. Interestingly, stem errors reduce in the HR setting. This is despite the use of hallucinated data.

5.1.1 Allomorphy

This category consists of errors which are characterised by a misapplication of morphophonological rules, or feature category mappings. Frequent errors include the absence of vowel harmony or place assimilation rules, and incorrect mapping of feature bundles to surface forms. Most errors are of this category.

Vowel harmony. The Nen language exhibits vowel harmony. Consider the form *yn̄jite* generated by one of the models, in a canonical sense the inflection is correct, but the presence of the high front vowel *i* requires the general *e* to harmonize to become *yn̄jiti*.

Morphophonological Rules. When combining *r* final stems with *t* phonemes (which occurs in inflections via the non-dual thematics or certain desinences with \emptyset thematics), the resultant sound is *n* (Evans, 2016). The M&C systems predicts that the stem *tar* inflected for the non-prehodiernal, first person actor and third person undergoer as *ytaretan*. Presumably, the break down is *y-tar-e-ta-n*. Interestingly, it inserts an *e* between the *r* and *t*, rather than concatenates the stem with the {-ta-n}

	A&G	M&C	NNB
Ambifixing only	0.111	0.170	0.010
Middle only	0.121	0.210	0.111
Prefixing only	0.212	0.250	0.010
Ambi + Pre	0.111	0.190	0.010
Ambi + Mid	0.071	0.130	0.040
Mid + Pre	0.141	0.290	0.040
Ambi + Mid + Pre	0.061	0.200	0.040

Table 4: Data sets for each composition type, model and sampling accuracies. The training size for each is 386 forms (defined by the available prefixing verbs).

by providing comprehensive lists of the morphemes in a given language (such as [Bickel and Nichols \(2005\)](#); [Shosted \(2006\)](#)). Thus, the complexity of an inflectional system is measured by enumerating the number of inflectional categories and the range of available markers for their realisation (i.e. E-complexity). The bigger the number, the more complex the resulting system is.⁷ With this in mind, we would expect that, given the same training size for each verb type, the ambifixing would perform the worst,⁸ then the middle followed by the prefixing verbs. Our results, shown in Table 4, confirm this hypothesis.

More revealing than the overall accuracy for each set and model combination, is a decomposition of accuracy according to the verb class. Table 5 summarises the performance for each category according to verb class. Unsurprisingly, when the training set contains only one type of verb, it performs best for the type of verb seen in the training data.

From a linguist perspective, with principle parts from the middle verbs (mainly the suffixal system, recall that the middle verb takes a dummy prefix to reduce valency) and prefixing verbs (prefixal paradigm) we can construct the full paradigm available to ambifixing verbs. The results presented here show no such compositionality; instead, we see a simple correspondence to verb type observed.

As expected, we see the weak *leaking* or overlap between ambifixing and middle verbs, with very little transferability from prefixing to other verb types. It highlights the importance of tag choice;

⁷Although more recent works have explored the issues with E-complexity ([Ackerman and Malouf, 2013](#)), we use it here as a guiding principle and acknowledge that further work is required to make a more nuanced statement.

⁸The combinatorial space for a transitive verb is 1,740 cells ([Muradoglu et al., 2020](#))

middle verbs have a [M] tag for the undergoer prefix, to mark the dummy prefix. If this tag were absent, would we see more transferability between ambifixing and middle verbs? Linguistically, no information would be lost as the absence of this tag still allows for the middle verbs to be clustered together.

5.3 Syncretism test

[Experiment 3](#), entailed testing the systems with an unseen feature bundle and analysing the predicted forms, to gauge whether the models learnt syncretic behaviour.

As can be seen by the suffixal paradigm found in [Evans \(2016\)](#),⁹ where both numbers are available, almost all the TAM categories exhibit syncretism across the second and third-person singular actor. The past perfective slot is the only case with distinct forms for the second and third singular person numbers. We are testing the prediction of an exception. The second singular is formed with $\{-nd-\emptyset-\}$ and the third person singular with the $\{-nd-a\}$ suffix. We note the similarity between the second singular and dual forms, where the second dual is $\{-a-nd\}$. This becomes particularly pertinent when a vowel is inserted between consonants for ease of articulation but must also adhere to vowel harmony. In such cases, the second dual and second singular may appear the same.

Using the [Aharoni and Goldberg \(2017\)](#) architecture, the model incorrectly predicts 81 out of the 100 test forms as the third singular perfective category with the suffix $\{-nd-a\}$ instead of $\{-nd-\emptyset-\}$. Four forms predicted correctly (likely due to the similarity between the surface forms of the second person dual and singular tags) and the remaining fifteen distributed across second person dual and plural actor of the same TAM category, second/third singular for the imperfective non-prehodiernal TAM category, and several instances of nonce inflections such as $\{-ngt\}$ or $\{-ngw\}$.

Similarly, the [Makarov and Clematide \(2018\)](#) system overwhelmingly predicts the unseen second singular form to be syncretic with the third singular (90 out of the 100 forms are predicted as such). Of the remaining ten instances three are correct, four are incorrectly modelled as the imperfective imperative (yet given the prefixing series is α , the future imperative prefix is absent) and one of

⁹Table 23.14 (pg 563) and Table 23.16 (pg 565)

	Ambifixing		Middle		Prefixing	
	AG	MC	AG	MC	AG	MC
Ambifixing only	11	15	2	0	0	2
Middle only	2	1	12	19	0	1
Prefixing only	0	0	0	0	21	24
Ambi + Pre	1	1	1	0	10	18
Ambi + Mid	1	4	6	8	0	1
Mid + Pre	0	3	3	10	11	16
Ambi + Mid + Pre	0	6	4	8	3	6

Table 5: Absolute number of correct predictions for each setup.

each: second/third imperfective non-prehodiernal, second/third neutral preterite or second dual past perfective.

From these results, it is clear that such systems not only observe patterns that are directly stipulated through annotation but also others that may be inferred from the data. It is important to note this behaviour, particularly in cases such as the one presented here as the verb corpus only entails two instances of the second singular past perfective.

6 Conclusion

Diversity representation of languages in NLP is vital to test the generalisations of models. We present the first-ever neural network-based analysis of Nen, the first representation of the Yam language family and to the best of our knowledge, of a Papuan language. Nen provides an interesting case study as it exhibits non-monotonic morphological mapping: distributed exponence.

We compare state-of-the-art models for morphological inflection across various training sizes and two sampling methods: random and Zipfian. The results show no significant difference between sampling methods, and minor differences may be attributed to training set composition differences. In the Zipfian case, the prefixing verb types are over-represented as they are more frequent in natural speech. We provide extensive analysis of types of errors generated by each system and show that the most common error type is allomorphy errors; a misapplication of morphophonological rules, or feature category mappings. We introduce a new subcategory of error type: free variation, which is a consequence of the natural speech origins of the corpus.

We further explore composition effects by generating training sets with incremental distributions for the three verb classes noted. As expected, we

found that the models trained with one class had higher prediction accuracy for that class. Across homogeneous compositions, the prefixing verb class performed the best. This is likely due to a smaller E-complexity – or more simply – a smaller combination of feature tags for which the system must learn mappings. Finally, we explore the likelihood of learning syncretic behaviour and using this as a predictor for an unseen feature bundle – the second singular past perfective. Overwhelmingly, the system incorrectly predicts syncretism with over 80% for the A&G system and 90% for the M&C system. These results highlight that these systems can infer patterns from the data sets provided. Although in our case the prediction of syncretism mirrors that of a human learner, there may be underlying, unwanted properties learnt from the data given, which calls for careful preparation of data and observation of output.

References

- Farrell Ackerman and Robert Malouf. 2013. Morphological organization: The low conditional entropy conjecture. *Language*, pages 429–464.
- Roei Aharoni and Yoav Goldberg. 2017. [Morphological inflection generation with hard monotonic attention](#). In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 2004–2015, Vancouver, Canada.
- Antonios Anastasopoulos and Graham Neubig. 2019. Pushing the limits of low-resource morphological inflection. In *Proc. EMNLP*, Hong Kong.
- Balthasar Bickel and Johanna Nichols. 2005. Inflectional synthesis of the verb. *The world atlas of language structures*, pages 94–97.
- Matthew J. Carroll. 2016. *The Ngkolmpu Language*. Ph.D. thesis, The Australian National University.

- Ryan Cotterell, Christo Kirov, John Sylak-Glassman, Géraldine Walther, Ekaterina Vylomova, Arya D. McCarthy, Katharina Kann, Sabrina J. Mielke, Garrett Nicolai, Miikka Silfverberg, David Yarowsky, Jason Eisner, and Mans Hulden. 2018. [The CoNLL–SIGMORPHON 2018 shared task: Universal morphological reinflection](#). In *Proceedings of the CoNLL–SIGMORPHON 2018 Shared Task: Universal Morphological Reinflection*, pages 1–27, Brussels. Association for Computational Linguistics.
- Ryan Cotterell, Christo Kirov, John Sylak-Glassman, Géraldine Walther, Ekaterina Vylomova, Patrick Xia, Manaal Faruqui, Sandra Kübler, David Yarowsky, Jason Eisner, and Mans Hulden. 2017. [CoNLL–SIGMORPHON 2017 shared task: Universal morphological reinflection in 52 languages](#). In *Proceedings of the CoNLL SIGMORPHON 2017 Shared Task: Universal Morphological Reinflection*, pages 1–30, Vancouver. Association for Computational Linguistics.
- Ryan Cotterell, Christo Kirov, John Sylak-Glassman, David Yarowsky, Jason Eisner, and Mans Hulden. 2016. [The SIGMORPHON 2016 shared Task—Morphological reinflection](#). In *Proceedings of the 14th SIGMORPHON Workshop on Computational Research in Phonetics, Phonology, and Morphology*, pages 10–22, Berlin, Germany. Association for Computational Linguistics.
- Qianji Di, Ekaterina Vylomova, and Tim Baldwin. 2019. [Modelling Tibetan verbal morphology](#). In *Proceedings of the The 17th Annual Workshop of the Australasian Language Technology Association*, pages 35–40, Sydney, Australia. Australasian Language Technology Association.
- Jason Eisner. 2002. Parameter estimation for probabilistic finite-state transducers. In *Proceedings of the 40th Annual Meeting of the Association for Computational Linguistics*, pages 1–8.
- Nicholas Evans. 2015. Valency in Nen. In Andrej Malchukov, Martin Haspelmath, Bernard Comrie, and Iren Hartmann, editor, *Valency classes: A comparative handbook*, pages 1069–1116. Berlin: Mouton de Gruyter.
- Nicholas Evans. 2016. [Inflection in Nen](#). In Matthew Baerman, editor, *The Oxford Handbook of Inflection*, pages pages 543–575. Oxford University Press, USA.
- Nicholas Evans. 2017. Quantification in nen. In *Handbook of Quantifiers in Natural Language: Volume II*, pages 571–607. Springer.
- Nicholas Evans. 2019a. [Nen dictionary](#). *Dictionaria*, pages 1–5005.
- Nicholas Evans. 2019b. [Waiting for the Word: Distributed Deponency and the Semantic Interpretation of Number in the Nen Verb](#), pages 100–123. Edinburgh University Press.
- Nicholas Evans. 2020. Waiting for the word: distributed deponency and the semantic interpretation of number in the Nen verb. In Andrew Hippisley Matthew Baerman, Oliver Bond, editor, *Morphological perspectives*, pages 100–123. Edinburgh: Edinburgh University Press.
- Nicholas Evans and Julia Colleen Miller. 2016. Nen. *Journal of the International Phonetic Association*, 46(3):331–349.
- Kyle Gorman, Arya D. McCarthy, Ryan Cotterell, Ekaterina Vylomova, Miikka Silfverberg, and Magdalena Markowska. 2019. [Weird inflects but OK: Making sense of morphological generation errors](#). In *Proceedings of the 23rd Conference on Computational Natural Language Learning (CoNLL)*, pages 140–151, Hong Kong, China. Association for Computational Linguistics.
- Alice C Harris. 2017. *Multiple exponence*. Oxford University Press.
- Ronald M Kaplan and Martin Kay. 1994. Regular models of phonological rule systems. *Computational linguistics*, 20(3):331–378.
- Christo Kirov, Ryan Cotterell, John Sylak-Glassman, Géraldine Walther, Ekaterina Vylomova, Patrick Xia, Manaal Faruqui, Sebastian J Mielke, Arya D McCarthy, Sandra Kübler, et al. 2018. Unimorph 2.0: Universal morphology. In *Proceedings of the Eleventh International Conference on Language Resources and Evaluation (LREC 2018)*.
- Kimmo Koskenniemi. 1983. *Two-level morphology: A general computational model for word-form recognition and production*, volume 11. University of Helsinki, Department of General Linguistics Helsinki, Finland.
- Vladimir I Levenstein. 1966. Binary codes capable of correcting deletions, insertions and reversals. In *Soviet physics doklady*, volume 10, pages 707–710.
- Peter Makarov and Simon Clematide. 2018. [Neural transition-based string transduction for limited-resource setting in morphology](#). In *Proceedings of the 27th International Conference on Computational Linguistics*, pages 83–93, Santa Fe, New Mexico, USA. Association for Computational Linguistics.
- John Mansfield. 2019. *Murrinhpatha morphology and phonology*, volume 653. Walter de Gruyter GmbH & Co KG.
- Peter H Mathews. 1974. *Morphology: an introduction to the theory of word-structure*. Cambridge, England: Cambridge University Press.
- Mehryar Mohri. 1997. Finite-state transducers in language and speech processing. *Computational linguistics*, 23(2):269–311.

- Saliha Muradoglu, Nicholas Evans, and Hanna Suominen. 2020. *To compress or not to compress? a finite-state approach to Nen verbal morphology*. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics: Student Research Workshop*, pages 207–213, Online. Association for Computational Linguistics.
- Saliha Muradoğlu. 2017. *When is enough enough ? A corpus-based study of verb inflection in a morphologically rich language (Nen)*. Masters thesis, The Australian National University.
- Graham Neubig, Chris Dyer, Yoav Goldberg, Austin Matthews, Waleed Ammar, Antonios Anastasopoulos, Miguel Ballesteros, David Chiang, Daniel Clothiaux, Trevor Cohn, et al. 2017. *Dynet: The dynamic neural network toolkit*. *arXiv preprint arXiv:1701.03980*.
- Joakim Nivre, Marie-Catherine De Marneffe, Filip Ginter, Yoav Goldberg, Jan Hajic, Christopher D Manning, Ryan McDonald, Slav Petrov, Sampo Pyysalo, Natalia Silveira, et al. 2016. *Universal dependencies v1: A multilingual treebank collection*. In *Proceedings of the Tenth International Conference on Language Resources and Evaluation (LREC'16)*, pages 1659–1666.
- Pushpendre Rastogi, Ryan Cotterell, and Jason Eisner. 2016. *Weighting 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, San Diego, California. Association for Computational Linguistics.
- Ryan K Shosted. 2006. *Correlating complexity: A typological approach*. *Linguistic Typology*, 10(1):1–40.
- Katsuhito Sudoh, Shinsuke Mori, and Masaaki Nagata. 2013. *Noise-aware character alignment for bootstrapping statistical machine transliteration from bilingual corpora*. In *Proceedings of the 2013 Conference on Empirical Methods in Natural Language Processing*, pages 204–209.
- John Sylak-Glassman, Christo Kirov, David Yarowsky, and Roger Que. 2015. *A language-independent feature schema for inflectional morphology*. In *Proceedings of the 53rd Annual Meeting of the Association for Computational Linguistics and the 7th International Joint Conference on Natural Language Processing (Volume 2: Short Papers)*, pages 674–680, Beijing, China.
- Ekaterina Vylomova, Jennifer White, Elizabeth Salesky, Sabrina J. Mielke, Shijie Wu, Edoardo Maria Ponti, Rowan Hall Maudslay, Ran Zmigrod, Josef Valvoda, Svetlana Toldova, Francis Tyers, Elena Klyachko, Ilya Yegorov, Natalia Krizhanovsky, Paula Czarnowska, Irene Nikkarinen, Andrew Krizhanovsky, Tiago Pimentel, Lucas Torroba Hennigen, Christo Kirov, Garrett Nicolai, Adina Williams, Antonios Anastasopoulos, Hilaria Cruz, Eleanor Chodroff, Ryan Cotterell, Miikka Silfverberg, and Mans Hulden. 2020. *SIGMORPHON 2020 shared task 0: Typologically diverse morphological inflection*. In *Proceedings of the 17th SIGMORPHON Workshop on Computational Research in Phonetics, Phonology, and Morphology*, pages 1–39, Online. Association for Computational Linguistics.