

Evaluating Neural Morphological Taggers for Sanskrit

Ashim Gupta¹, Amrith Krishna², Pawan Goyal³, Oliver Hellwig⁴

¹School of Computing, University of Utah

²ITU Copenhagen, ³Dept. of Computer Science and Engineering, IIT Kharagpur

⁴University of Zürich, IVS

ashim@cs.utah.edu, amrk@itu.dk,

pawang@cse.iitkgp.ac.in, oliver.hellwig@uzh.ch

Abstract

Neural sequence labelling approaches have achieved state of the art results in morphological tagging. We evaluate the efficacy of four standard sequence labelling models on Sanskrit, a morphologically rich, fusional Indian language. As its label space can theoretically contain more than 40,000 labels, systems that explicitly model the internal structure of a label are more suited for the task, because of their ability to generalise to labels not seen during training. We find that although some neural models perform better than others, one of the common causes for error for all of these models is mispredictions due to syncretism.¹

1 Introduction

Sanskrit is a fusional Indo-European language with rich morphology, both at the inflectional and derivational level. The language relies heavily on morphological markers to determine the syntactic, and to some extent the semantic roles, of words in a sentence. There exist limited and partly incompatible solutions (Hellwig, 2016; Goyal and Huet, 2016; Krishna et al., 2018) for morphological tagging of Sanskrit that heavily rely on lexicon driven shallow parsers and other linguistic knowledge. However recently, neural sequential labelling models have achieved competitive results in morphological tagging for multiple languages (Cotterell and Heigold, 2017; Tkachenko and Sirts, 2018; Malaviya et al., 2018). We therefore explore the efficacy of such models in performing morphological tagging for Sanskrit without access to extensive linguistic information.

Most recent research treats morphological tagging as a structured prediction problem where the morphological class of a word is either treated as

a monolithic label or as a composite label with multiple features (Müller et al., 2013; Cotterell and Heigold, 2017). Schmid and Laws (2008); Hakkani-Tür et al. (2002) model the morphological tags as a sequence of individual morphological features. Recently, Tkachenko and Sirts (2018) proposed to generate this sequence of morphological features using a neural encoder-decoder architecture. Hellwig (2016) shows a significant improvement in performance for morphological tagging in Sanskrit by using a monolithic tagset with recurrent neural network based tagging model. In systems using monolithic labels, multiple feature values pertaining to a word are combined to form a single label (Müller et al., 2013; Heigold et al., 2017), which leads to data sparsity for morphologically rich languages such as Czech, Turkish and Sanskrit. The sparsity issue can be addressed by using composite labels which model the internal structure of a class as a set of individual features (Tkachenko and Sirts, 2018; Zalmout and Habash, 2017). Malaviya et al. (2018) use a neural factorial CRF to model inter-dependence between individual categories of the composite morphological label.

However, as the decision for monolithic vs. composite labels is one of the central design choices when tagging morphologically rich languages, we use Sanskrit as a test case for a systematic evaluation for this choice. For this evaluation, we consider several neural architectures with different modelling principles. For the monolithic tag model, the neural architecture is based on a bi-directional LSTM with a linear CRF layer stacked on top of it (Lample et al., 2016; Huang et al., 2015). For composite labels, we explored a neural generation model that generates a sequence of morphological features for each word in the input sequence. In order to explicitly capture the inter-dependencies between the morphological features, we use a model based on a factorial Conditional

¹Code and data available at <https://github.com/ashim95/sanskrit-morphological-taggers>

Feature	Values	Noun	Fin. verb	Participle	Compound	Others
Tense	18		✓	✓		
Case	8	✓		✓		
Number	3	✓	✓	✓		
Gender	3	✓		✓		
Person	3		✓			
Last char.	31	✓		✓	✓	
Other	5					✓
Total	71	2232	162	40176	31	5

Table 1: Grammatical features and their distribution over inflectional classes. ‘Others’ include forms such as infinitives, absolutes, Indeclinables, etc.

Random Field (CRF) (Malaviya et al., 2018). Additionally, independent classifiers trained under a multi-task setting with sharing of parameters are also explored (Inoue et al., 2017; Søgaard and Goldberg, 2016). Our experiments specifically focus on the following problems and questions:

- **Syncretism:** We will show that syncretism, i.e., inflected forms of a lemma that share the same surface form for multiple morphological tags, is the major source of mispredictions. We evaluate if and how models with monolithic and composite labels deal with this phenomenon.
- **Performance on unseen tags:** For models with composite labels, it should be possible to predict morphological classes which were not seen in the training data. Our experiments show that the performance of the systems remains more or less the same irrespective of the neural architecture.

This raises an important point: Models that perform marginally better in terms of evaluation metrics are supposedly not superior, since we see similar performances on special test sets targeting particular statistical phenomenon (unseen tags) and linguistic phenomenon (syncretism).

2 Problem Formulation and Models

Tagset: Sanskrit, similar to Hungarian (Zsibrita et al., 2013) and Turkish (Çetinoğlu and Kuhn, 2013), relies on suffixes for marking inflectional information. As Sanskrit has a rich inflectional system, the size of the tagset plays a relevant role. Hellwig (2016) uses a tagset with 86 possible labels that merges some grammatical features based

on linguistic considerations. Krishna et al. (2018) use an extended tagset of 270 features, by adding the feature tense, but only to finite verbs. As the systems tested in this paper do not use external linguistic information that could restrict the range of applicable features, we choose a tagset consisting of the features shown in Table 1, which is in principle motivated by the traditional grammatical analysis of Sanskrit (Apte, 1885).² As the declensional type of a noun in Sanskrit is determined by the last character in the non-inflected stem of the word, we add the last character of the stem as a morphological feature in our predictions.

Notation: Given a sequence of tokens $\mathbf{x} = x_1, x_2, \dots, x_T$, we aim at predicting a sequence of labels $\mathbf{y} = y_1, y_2, \dots, y_T$, one for each token. Each label y_i is a composite label $y_i = \{y_{i0}, y_{i1}, \dots, y_{it}\}$ and consists of a collection of grammatical features for x_i .

Encoder: All neural sequence labelling models tested in this paper (see below) use an encoder that generates the input representations of words as follows. Given a sequence of tokens as input, for each token $x_i \in \mathbf{x}$, its vector representation is obtained by concatenating its word embedding with a sub-word character embedding obtained from a bi-directional LSTM similar to Lample et al. (2016). These word representations are passed through a word level Bi-LSTM to obtain hidden state h_i for each token in the sequence.

Monolithic Sequence Model (MonSeq): This is a standard neural sequence labelling model with a neural-CRF tagger (Lample et al., 2016; Huang et al., 2015). A linear chain CRF is used as the output layer for the monolithic labels used.

Neural Factorial CRF (FCRF): The model proposed by Malaviya et al. (2018) is an end to end neural sequence labeling model with a factorial CRF (Sutton et al., 2007). The model is shown in Figure 1a. In order to model the inter-dependence between different morphological types, a pairwise potential between cotemporal variables and a transitional potential between variables of the same type of tags is used. As exact inference is computationally expensive, loopy belief propagation is

²Exhaustive list of inflectional tags used by three prevalent tagging schemes in Sanskrit computational linguistics: <https://sanskritlibrary.org/helpmorphids.html>

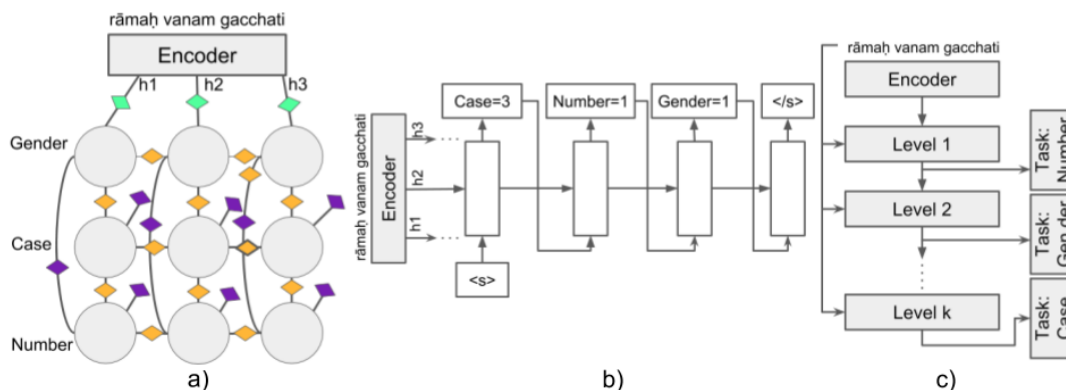


Figure 1: Models for composite tag prediction, demonstrated for the toy sentence *rāmaḥ vanam gacchati* (English: ‘Rāma goes to the forest’). a) FCRF b) Seq c) MTL-Hierarchy. For illustrative purposes, we consider a tagset consisting of only three categories.

used to compute approximate factor and variable marginals.

Sequence Generation Model (Seq): We use a sequence generation model (Tkachenko and Sirts, 2018) that predicts composite labels as sequences of feature values for every token in the sequence. For token x_i , the LSTM hidden state h_i is fed to an LSTM based decoder, which generates a sequence of feature values conditioned on the context vector h_i and the previous feature value. As shown in Figure 1b, the decoding is initiated with a special marker passed as input and terminates when the decoder predicts an end marker.

Multi task learning (MTL): Here, we consider each grammatical feature as shown in Table 1 as a separate task. We experiment with two settings for multi-task learning. In **MTL-Shared** the encoder parameters are shared across all the tasks, and supervision for all the tasks is performed at the same layer with each task having its own independent output CRF layer. In **MTL-Hierarchy**, as proposed by Søgaard and Goldberg (2016), we establish a hierarchy between the grammatical categories. A hierarchical inductive bias is introduced by supervising low-level tasks at the bottom layers and higher-level tasks at the higher layers (Sanh et al., 2019). Concretely, for a task k , only the parameters at the output CRF layer and those at the shallower layers are updated. The model is shown in Figure 1c. In order for the higher layers to have access to the inputs, we use shortcut connections as proposed in Hashimoto et al. (2017).

3 Experiments

Data: We use a training set of 50,000 and a test set of 11,000 sentences from the Digital Corpus

of Sanskrit (DCS, Hellwig (2010–2020)). To prepare the training set, we sample sentences such that there exist at least 100 instances for each of the 71 features in Table 1. The training data still covers only 2,757 out of possible 42,606 labels, indicating that the true dimension of the target space is much lower than could be expected from Table 1. The test data we use contain only about 0.5 % of the tokens with labels not present in the training set. Additionally we use a separate **unseen** test set, where every sentence contains at least one word with a monolithic label not present in the training data.

3.1 Evaluation

We report the performance using average token-level accuracy and F1-scores (see Malaviya et al., 2018; Cotterell and Heigold, 2017; Buys and Botha, 2016). The average token-level accuracy is reported on the exact match of a morphological tag for a token, i.e., if it predicts all the morphological features correctly. The F1 measure is computed on a tag-by-tag basis, i.e. macro and micro averaged at the grammatical category level, which provides partial credit to partially correct tag sets.

3.2 Results

Table 2 shows the results for the five models studied in this paper. We find that three of the four models using composite labels obtain overall comparable results, with hierarchical Multi-task model obtaining the highest Macro and Micro F1-Scores, token accuracy as well as outperforming other models for category specific evaluation for 4 out of 5 categories. This highlights that there is some gains to be had by inducing a hierarchical bias among these morphological categories. As can be observed from

System	Token Accuracy	Macro F1* Score	Micro F1 Score	T	C	N	G	L
MonSeq	75.15	85.16	85.14	79.64	86.39	91.05	80.37	83.86
MTL-Shared	81.72	90.82	90.84	85.98	88.31	95.61	90.84	90.00
MTL-Hierarchy	86.74	94.11	94.11	93.77	92.17	97.21	93.30	94.01
FCRF	85.37	93.01	93.01	92.35	89.08	96.38	92.44	94.03
Seq	85.95	92.72	92.72	91.10	89.66	96.52	92.01	93.12

Table 2: Performance of different systems for the morphological tagging task. All the reported values for Macro-F1 are statistically significant ($p < 0.05$), based on pairwise t-tests between the systems. The features (refer Table 1) are marked using their first letter.

	Syncretism	Unseen
System	Macro F1	Macro F1
Seq	70.44	55.98
MTL-Shared	70.26	55.02
MTL-Hierarchy	70.62	55.21
FCRF	68.80	55.62
MonSeq	52.55	–

Table 3: Macro-F1 score for the tokens which exhibit syncretism (top 25 pairs, based on reported mispredictions) and the unseen labels during training.

the results, all the composite models clearly outperform MonSeq in terms of Macro- and Micro F1-Score, indicating their better performance for rare morphological classes. Among the composite models, MTL-Shared clearly underperforms, which is probably due to the fact that most of its parameters are shared by all the tasks and no task specific adaptation was possible. We also perform pairwise t-tests and find that the gains reported are statistically significant ($p < 0.05$).

One of our key findings is that syncretism is a major source of error for all these systems. For the composite models, about 20 % to 25 % of all the mispredictions in nouns arise due to syncretism. As expected it is worse for MonSeq, where close to 37 % mispredictions are due to this linguistic phenomenon. For a more detailed analysis, we check the top 25 label-pairs of mis-predictions for each system. In Table 3, we report macro F1-Scores for the tokens which exhibit syncretism from this filtered set. The reported results are far below the overall F1-Macro scores, as shown in Table 2.

Composite-label models are also able to make partially correct predictions for more unusual forms. The 3rd person plural perfect form, *sasarjire* (English: ‘they have created’), for example, is analysed as 3rd sg. perfect by Seq. This decision

Hierarchy	Number	Tense
N-G-C-T-L	97.21	93.77
T-C-N-G-L	96.49	90.12
N-G-L-T-C	97.24	92.85

Table 4: Results for Number and Tense based on different configurations in MTL-Hierarchy. The shallow to deep levels are marked by first character of the features.

should be influenced by the last letter -e, which can indicate the 3rd singular of the perfect, while the correct, and relatively rare, affix is *-ire*. MonSeq predicts a locative singular of a non-existing noun **sasarjira* in this case. Again, this decision is probably based on the last letter -e, which in most cases derives the locative singular.

Table 3 (right half) shows how the models perform for tags unseen during training (on **unseen** test set). We consider 11 of such case, number and gender combinations, 14 tense, person and number combinations and two of the tenses additionally for the participles. Among four different composite models, the macro F1-Score for ‘Seq’ model is similar to what is observed in [Tkachenko and Sirts \(2018\)](#) for a fusional language like Czech. Moreover, the behaviour for all the composite models remains more or less same for unseen labels.

Next, we explore if there is a natural hierarchy for supervision of morphological categories in MTL-hierarchy model. For this, we train the system in different permutations of feature hierarchies as shown in Table 4. We observe that the feature *number* benefits from supervision at a shallower level, whereas *tense* always benefits from supervision at a deeper level. The trends for other features were not as conclusive, but these results show there might be an inherent hierarchy among some of the morphological features.

[Krishna \(2019\)](#), an extended version of the en-

ergy based model proposed in Krishna et al. (2018), report the state of the art results for morphological parsing in Sanskrit. They report a token level accuracy, macro averaged over sentences, of 95.33 % on a test set of 9,576 sentences. The FCRF and SeqGen models, when tested on their test dataset, achieve an average sentence level token accuracy of 80.26 % and 81.79 % respectively. Here, it needs to be noted that the morphological tagger used in Krishna (2019) relies on a lexicon driven shallow parser to obtain a smaller search space of candidates. However, this makes the model a closed vocabulary model as it would fail for cases where the words are not recognised by the lexicon-driven parser, as there will be no analyses for such words. On the contrary, none of the models presented in this work are constrained by the vocabulary of any lexicon.

4 Conclusion and Future Work

In this work, we evaluated various neural models for morphological tagging of Sanskrit, concentrating on models that are capable of using composite labels. We find that all the composite label models outperform MonSeq by significant margins. These models, with an exception to MTL-Shared, achieve overall competitive results when enough training data is available. A major problem for all the sequence labelling models studied in this paper is syncretism of morphological categories, which should constitute the main focus of future research.

References

- Vaman Shivaram Apte. 1885. *The student's guide to Sanskrit composition: a treatise on Sanskrit syntax for the use of schools and colleges*. Chowkhamba Sanskrit Series Office.
- Jan Buys and Jan A. Botha. 2016. [Cross-lingual morphological tagging for low-resource languages](#). In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1954–1964. Association for Computational Linguistics.
- Özlem Çetinoğlu and Jonas Kuhn. 2013. [Towards joint morphological analysis and dependency parsing of Turkish](#). In *Proceedings of the Second International Conference on Dependency Linguistics (DepLing 2013)*, pages 23–32. Charles University in Prague, Matfyzpress, Prague, Czech Republic.
- Ryan Cotterell and Georg Heigold. 2017. [Cross-lingual character-level neural morphological tagging](#). In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing*, pages 748–759. Association for Computational Linguistics.
- Pawan Goyal and Gérard Huet. 2016. Design and analysis of a lean interface for sanskrit corpus annotation. *Journal of Language Modelling*, 4(2):145–182.
- Dilek Z Hakkani-Tür, Kemal Oflazer, and Gökhan Tür. 2002. Statistical morphological disambiguation for agglutinative languages. *Computers and the Humanities*, 36(4):381–410.
- Kazuma Hashimoto, caiming xiong, Yoshimasa Tsu-ruoka, and Richard Socher. 2017. [A joint many-task model: Growing a neural network for multiple NLP tasks](#). In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing*, pages 1923–1933. Association for Computational Linguistics.
- Georg Heigold, Guenter Neumann, and Josef van Genabith. 2017. An extensive empirical evaluation of character-based morphological tagging for 14 languages. In *Proceedings of the 15th Conference of the European Chapter of the Association for Computational Linguistics: Volume 1, Long Papers*, volume 1, pages 505–513.
- Oliver Hellwig. 2010–2020. [DCS - The Digital Corpus of Sanskrit](#).
- Oliver Hellwig. 2016. Improving the morphological analysis of classical Sanskrit. In *Proceedings of the 6th Workshop on South and Southeast Asian Natural Language Processing (WSSANLP2016)*, pages 142–151.
- Zhiheng Huang, Wei Xu, and Kai Yu. 2015. Bidirectional LSTM-CRF models for sequence tagging. *arXiv preprint arXiv:1508.01991*.
- Go Inoue, Hiroyuki Shindo, and Yuji Matsumoto. 2017. Joint prediction of morphosyntactic categories for fine-grained arabic part-of-speech tagging exploiting tag dictionary information. In *Proceedings of the 21st Conference on Computational Natural Language Learning (CoNLL 2017)*, pages 421–431.
- Amrith Krishna. 2019. *Addressing Language Specific Characteristics for Data-Driven Modelling of Lexical, Syntactic and Prosodic Tasks in Sanskrit*. Ph.D. thesis, IIT Kharagpur, Kharagpur, West Bengal, India.
- Amrith Krishna, Bishal Santra, Sasi Prasanth Bandaru, Gaurav Sahu, Vishnu Dutt Sharma, Pavankumar Satuluri, and Pawan Goyal. 2018. [Free as in free word order: An energy based model for word segmentation and morphological tagging in Sanskrit](#). In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 2550–2561. Association for Computational Linguistics.

- Guillaume Lample, Miguel Ballesteros, Sandeep Subramanian, Kazuya Kawakami, and Chris Dyer. 2016. Neural architectures for named entity recognition. In *Proceedings of NAACL-HLT*, pages 260–270.
- Chaitanya Malaviya, Matthew R. Gormley, and Graham Neubig. 2018. [Neural factor graph models for cross-lingual morphological tagging](#). In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 2653–2663. Association for Computational Linguistics.
- Thomas Müller, Helmut Schmid, and Hinrich Schütze. 2013. Efficient higher-order CRFs for morphological tagging. In *Proceedings of the 2013 Conference on Empirical Methods in Natural Language Processing*, pages 322–332.
- Victor Sanh, Thomas Wolf, and Sebastian Ruder. 2019. A hierarchical multi-task approach for learning embeddings from semantic tasks.
- Helmut Schmid and Florian Laws. 2008. [Estimation of conditional probabilities with decision trees and an application to fine-grained pos tagging](#). In *Proceedings of the 22Nd International Conference on Computational Linguistics - Volume 1, COLING '08*, pages 777–784, Stroudsburg, PA, USA. Association for Computational Linguistics.
- Anders Søgaard and Yoav Goldberg. 2016. [Deep multi-task learning with low level tasks supervised at lower layers](#). In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, pages 231–235. Association for Computational Linguistics.
- Charles Sutton, Andrew McCallum, and Khashayar Rohanimanesh. 2007. Dynamic conditional random fields: Factorized probabilistic models for labeling and segmenting sequence data. *Journal of Machine Learning Research*, 8(Mar):693–723.
- Alexander Tkachenko and Kairit Sirts. 2018. [Modeling composite labels for neural morphological tagging](#). In *Proceedings of the 22nd Conference on Computational Natural Language Learning*, pages 368–379. Association for Computational Linguistics.
- Nasser Zalmout and Nizar Habash. 2017. [Don't throw those morphological analyzers away just yet: Neural morphological disambiguation for Arabic](#). In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing*, pages 704–713. Association for Computational Linguistics.
- János Zsibrita, Veronika Vincze, and Richárd Farkas. 2013. [magyarlanc: A tool for morphological and dependency parsing of hungarian](#). In *Proceedings of the International Conference Recent Advances in Natural Language Processing RANLP 2013*, pages 763–771. INCOMA Ltd. Shoumen, BULGARIA.