

Morphological Reinflection with Multiple Arguments: An Extended Annotation schema and a Georgian Case Study

David Guriel, Omer Goldman, Reut Tsarfaty

Bar-Ilan University

{davidgu1312, omer.goldman}@gmail.com, reut.tsarfaty@biu.ac.il

Abstract

In recent years, a flurry of morphological datasets had emerged, most notably UniMorph, a multi-lingual repository of inflection tables. However, the flat structure of the current morphological annotation schema makes the treatment of some languages quirky, if not impossible, specifically in cases of polypersonal agreement, where verbs agree with multiple arguments using true affixes. In this paper we propose to address this phenomenon, by expanding the UniMorph annotation schema to hierarchical feature structure that naturally accommodates complex argument marking. We apply this extended schema to one such language, Georgian, and provide a human-verified, accurate and balanced morphological dataset for Georgian verbs. The dataset has 4 times more tables and 6 times more verb forms compared to the existing UniMorph dataset, covering all possible variants of argument marking, demonstrating the adequacy of our proposed scheme. Experiments with a standard reinflection model show that generalization is easy when the data is split at the form level, but extremely hard when splitting along lemma lines. Expanding the other languages in UniMorph to this schema is expected to improve both the coverage, consistency and interpretability of this benchmark.

1 Introduction

In recent years, morphological (re)inflection tasks have gained a lot of attention in NLP.¹ Subsequently, several multi-lingual morphological datasets have emerged to allow for the supervised training of morphological models, most notably UniMorph (McCarthy et al., 2020), that organizes words into inflectional tables, annotating each inflected word-form with its respective feature-set.

While western languages are widely represented in UniMorph, many *morphologically rich languages* (Tsarfaty et al., 2010, 2020) exhibit rich

¹Cf. the series of SIGMORPHON shared tasks: <https://sigmorphon.github.io/sharedtasks/>

and diverse inflection patterns that make them less compatible with the flat feature-sets in the UniMorph schema. Concretely, in some cases it is completely impossible to annotate parts of the inflectional paradigm with a flat bundle, as is the case with *case stacking*, and in other cases, such as *polypersonal agreement*, the annotation solutions provided are unnatural, non-transparent, and are barely used in practice. As a result, languages exhibiting such phenomena are under-represented in UniMorph, and when they are, the inflection tables for these languages are often incomplete.

In this paper we propose a general solution for annotating such structures, thus extending the UniMorph annotation schema to fully cover a wider range of morphologically-complex argument-marking phenomena. Following Anderson (1992), we propose a so-called *layered* annotation of features, where the inflectional features take the form of a *hierarchical* structure, in the spirit of formal linguistic frameworks as that of Johnson (1988); Pollard and Sag (1994); Shieber (2003); Bresnan et al. (2015). We organize the features of multiple arguments in a hierarchical structure, rather than the current flat structure that accommodates only subject concords. This schema shift allows for an adequate annotation of *polypersonal agreement* and of *possessed nominals*, where a word has multiple number and gender features, as well as forms with *case stacking*, where a word has multiple cases.

We apply the suggested solution to Georgian, an agglutinative language with a convoluted verbal system, that indicates both subjects and objects with true affixes (rather than clitics that are omissible from the inflection tables). We create a new human-verified dataset for Georgian, that covers most of the grammatical phenomena in Georgian verbs, and includes 118 lemmas, adding up to about 21k verb forms, compared with the 47 lemmas and 3.3k verb forms, some of which are erroneous, currently available in the Georgian UniMorph.

გაგინებთ (<i>gagišvebt</i>)					
გა-	გ-	ი-	შვ	-ებ	-თ
ga-	g-	i-	šv	-eb	t
FUT	O2SG	TRANS	LET GO	THEME	S1PL

We will let you(sg.) go

Table 1: A typical Georgian verb. Note the 2 argument markers, one object (tagged with O) and one subject (S).

We use the new dataset to train a standard morphological reinflection model (Silfverberg and Hulden, 2018) and show that training on the Georgian inflections currently available in UniMorph is not sufficient for generalizing to the more inclusive set of inflections that are allowed by the new scheme. We conclude that our annotation approach provides a more complete representation of linguistic behaviors, and that our proposed Georgian dataset provides a much better depiction of the morphological phenomena that exist in the data and the computational challenge reflected therein.

We therefore call to apply layered annotation to all currently existing morphological data in UniMorph, to more consistently and transparently capture the linguistic reality and morphological complexity reflected in the world languages.

2 The Problem: Multiple Arguments

Models of morphological reinflection are trained to generate forms within a lemma L , given another form and the features of source $_i$ and target $_j$ forms:

$$(\langle feat_i^L, form_i^L \rangle, \langle feat_j^L, _ \rangle) \mapsto form_j^L$$

For example, for the Russian lemma ЛЕТЕТЬ: reinflecting from (PRS;1;SG,лечу) to (IMP;2;SG,лети) will be represented as:

$$(\langle PRS;1;SG,лечу \rangle, \langle IMP;2;SG, _ \rangle) \mapsto лети$$

Standardly, the data for training morphological models (e.g., Wu et al., 2020; Makarov and Clematide, 2018) is taken from UniMorph (McCarthy et al., 2020), a multilingual morphological dataset in which words are grouped by lemma into inflection tables, each word is tagged with an unordered set of morphological features. The features list is shared across languages. The inflection tables are meant to be exhaustive, i.e., covering all possible forms of a lemma, regardless of usability.

Although the features were designed to apply cross-lingually, some blind-spots exist. Most relevant to our work is the assumption that every feature set includes at most one pronominal feature bundle (i.e., person-gender-number).

However, this assumption does not apply to verbs with object concords, as exhibited in Georgian (see Table 1), Inuit and many Bantu languages *inter alia*, nor does it apply to possessed nouns that mark the features of both the possessor and the possessee. Examples (1a)–(1d) illustrate this:

- (1) a. Georgian: *gagišvebt* ‘We will let you go’ (SUBJ-1PL, OBJ-2SG)
- b. Turkish: *kedisinin* ‘you are his cat’ (NOUN-SG, SUBJ-2SG, POSS-3SG)
- c. Swahili: *ninakupenda* ‘I love you’ (SUBJ-1SG, OBJ-2SG)
- d. Hebrew: *emdata* ‘her position’ (NOUN-SG, POSS-3SG-FEM)

The solution proposed in UniMorph to annotating these phenomena is via concatenating several properties into a single string, lacking any internal structure; e.g., ARGAC2S indicates a form with a 2nd person singular accusative argument (Sylak-Glassman, 2016). However, there are at least two shortcomings to this solution. First, it is not sufficiently transparent. ARGAC2S is an opaque string, that does not decompose into the known features licensed by the UniMorph features list (i.e., ACC, 2, SG). Secondly, and possibly due to this lack of transparency, this annotation hack is hardly ever used in practice. Hence, from all examples in (1), only the Hebrew form is included in UniMorph, and tagged as N;SG;FEM;PSS3S with multiple possessor features merged into the flat string PSS3S.

The crux of the matter is that in the current annotation schema, complex features assigned to additional arguments are treated as a single non-decomposable feature, that lack any internal structure, unlike the features of the main (so-called ‘internal’) argument, that are individually spelled out. We argue that the lack of transparency and usability are due to the misrepresentation of the inherently *hierarchical* and *compositional* structure of the features in such forms. We suggest to explicitly annotate these forms with features that are all explicitly composed of the same primitive features.

All in all, the lack of a sufficiently expressive annotation standard leads to a data distribution that is skewed, unrealistically simple, and, when language-specific annotation solutions are painfully needed, they suffer from inconsistencies and ad-hoc decisions. For these reasons, we set out to extend the UniMorph annotation schema to accommodate all such cases and to enable a proper coverage of languages, such as Georgian and many others.

3 The Proposed Schema

We propose to extend the UniMorph annotation schema to cover multiple pronominal feature-bundles in the same word-form, via a *layering* approach, originally proposed for morphological systems by Anderson (1992). Anderson suggests to arrange the *morphosyntactic representation* (MSR) of words in a hierarchy (dubbed *layers*) of features, in the sense that every element of the unordered set of features can be composed of another unordered set of features. That is, a general feature annotation looks as in (2a). A specific transitive verb annotation could be as depicted in (2b):

- (2) a. $[f_1, f_2, \dots, [F_i : f_{i_1}, f_{i_2}, \dots, [F_j : f_{j_1} \dots]]]$
 b. $[V, Tense,$
 $[nom : Per, Num, Gen],$
 $[acc : Per, Num, Gen]]$

This hierarchical feature structure is reminiscent of *unification grammars* or *attribute-value grammars* (Shieber, 2003; Johnson, 1988) that are extensively used in syntactic theories such as GPSG, HPSG, and resemble the *f-structures* in LFG (Gazdar et al., 1989; Pollard and Sag, 1994; Bresnan et al., 2015).

Here we employ these structures to organize the features of morphologically-marked arguments hierarchically, so an argument is characterized by a feature composite of all features pertaining to that argument. That is, each argument’s feature-bundle is specifically marked with the argument it belongs to, and is decomposed into the primitive features licensed by the UniMorph scheme. It also homogeneously annotates the different kinds of arguments, in contrast with the current schema where the subject features are assigned to the verb directly. Thus, the English form *thinks* previously annotated as $V;PRS;3;SG$ will be annotated as $V;PRS;NOM(3;SG)$. In languages that mark multiple arguments, different kinds of arguments can be marked with their feature-bundles without conflicts. The proposed schema thus facilitates the annotation of the poorly-treated or untreated phenomena as illustrated in (1). These are, respectively:

- (3) a. Georgian: *gagišvebt* ‘We will let you go’
 $V;FUT;NOM(1;PL);ACC(2;SG)$
 b. Turkish: *kedisin* ‘you are his cat’
 $N;SG;NOM(2;SG);POSS(3;SG)$
 c. Swahili: *ninakupenda* ‘I love you’
 $V;PRS;NOM(1;SG);ACC(2;SG)$
 d. Hebrew: *emdata* ‘her position’
 $N;SG;POSS(3;SG;FEM)$

Table 2 compares the annotation of these examples in the current UniMorph schema compared with our proposed annotation schema.² The hierarchical structures, beyond being more transparent, opens the door further for future study on compositional generalization in morphology.

The resemblance of our proposed schema to ideas in other fields of theoretical linguistics, most prominently to the *f-structure* in LFG (Bresnan et al., 2015) and to the nested *Attribute-Value matrices* in HPSG (Pollard and Sag, 1994), points to a natural interface with further syntactic and semantic annotations downstream.

4 A Case Study from Georgian

Linguistic Background Georgian is an agglutinative language with a verbal system that makes a vast use of affixes to convey a wide array of meanings, both inflectional and derivational (see Table 1). The Georgian verbal paradigm is divided into 5 classes known as: transitive, intransitive, medial, indirect and stative (Hewitt, 1995). The verbs are inflected to reflect 12 Tense-Aspect-Mood (TAM) combinations (traditionally known as *screeves*) sorted into 4 series: present and future, aorist, perfective, and the imperative. Each series has its own morpho-syntactic characteristics, most notably split-ergativity is manifested in the aorist.

The characteristic most essential to this work is that Georgian verbs always agree on person and number with the direct and indirect objects, on top of the subject-verb agreement. The Georgian data in UniMorph follows the convention of including objects *only* in third person singular — thus failing to provide a comprehensive coverage of the word-forms that can be attested in the language.

Additional issues with the current morphological data in UniMorph for Georgian verbs are: sparsity, as it includes only 47 inflection tables; lack of diversity, as all table are from the transitive class; and lack of accuracy, as the data was produced automatically without verification by native speakers.

Data Annotation A key contribution of this work is the creation of a new dataset for Georgian that follows the layered annotation schema and addresses the other shortcomings just described. We selected a list of 118 verb lemmata from all differ-

²Although not explicitly shown here, annotation of case stacking is also possible with our approach, while non-hierarchical annotations do not account for such cases. For example, Korean 교사에게이 can be tagged as $N;SG;NOM(DAT)$.

	Flat structure	Hierarchical Structure
Georgian: <i>gagišvebt</i> Trans: ‘We will let you go’ Args: SUBJ-1PL, OBJ-2SG	<pre> V / \ FUT ARGNO1P ARGAC2S </pre>	<pre> V / \ FUT NOM ACC / \ 1 PL 2 SG </pre>
Turkish: <i>kedisinin</i> Trans: ‘you are his cat’ Args: NOUN-SG, SUBJ-2SG, POSS-3SG	<pre> N / \ SG ARGNO2S PSS3S </pre>	<pre> N / \ SG NOM POSS / \ 2 SG 3 SG </pre>
Swahili: <i>ninakupenda</i> Trans: ‘I love you’ Args: SUBJ-1SG, OBJ-2SG	<pre> V / \ PRS ARGNO1S ARGAC2S </pre>	<pre> V / \ PRS NOM ACC / \ 1 SG 2 SG </pre>
Hebrew: <i>emdata</i> Trans: ‘her position’ Args: NOUN-SG, POSS-3SG-FEM	<pre> N / \ SG PSS3S FEM </pre>	<pre> N / \ SG POSS / \ 3 SG FEM </pre>

Table 2: Examples for word-forms with multiple argument agreements. On the left we present the flat structure currently employed in UniMorph. All examples save Hebrew are not included in the UniMorph inflection tables, presumably due to their lack of transparency. On the right we present our proposed hierarchical structure, which is more transparent, and also amenable for compositional generalization.

ent classes.³ Every verb was manually annotated with its stem, its thematic affix and principal parts, to automatically generate the full inflection tables.

This automatic generation of Georgian verbs is prone to some errors, for instance, in accounting for idiosyncratic phonologically-conditioned stem changes. Hence, we ran our data through 3 native Georgian speakers to assert its correctness, or fix when needed. In cases where speakers were unsure we used a Georgian morphological analyzer (Doborjginidze and Lobzhanidze, 2012) for consultation. In cases of disagreement, we used a majority vote among the speakers. On average, at least one speaker was uncertain in about 5% of the forms, but a disagreement that necessitated a majority vote occurred only on about 0.7% of the cases.

Table 3 summarizes the statistics over our annotated data. In total, we produced 21,054 verb forms, of 118 lemmata. The data is quite evenly balanced across the classes, with more verbs drawn from the more frequent transitive class. For comparison, the current UniMorph data has fewer lemmas, 3,300 forms, and includes only verbs that are transitive.⁴

³We based the list of verbs on those whose inflection tables appear on Hewitt (1995) and added some commonly-used verbs suggested by native speakers.

⁴All our data is publicly available at <https://github.com/Onlp/GeorgianMorphology>.

	Trans.	Intrans.	Med.	Indi.	Stat.
#Infl. Tables	40	21	29	16	12
#Verb Forms	12506	2560	3132	2626	230

Table 3: Distribution of the Georgian verbs over classes.

5 Experiments

To assess the usability of our dataset, we trained a standard reinflection model, the character-level LSTM of Silfverberg and Hulden (2018), on our data.⁵ We sampled from our data 2 datasets for training morphological reinflection models, containing train, validation and test sets in sizes $8k$, $1k$ and $1k$ examples, respectively. Following Goldman et al. (2021), one dataset employed an easier form-split, i.e., no forms appear in both train and test,⁶ and the other with the more challenging lemma-split, where lemmas from train, dev and test are disjoint. To assess the generalization capacity we varied the sources of both the train and test sets.⁷ We report 2 evaluation metrics: *accuracy* over exact matches, and *average edit distance* from gold.

⁵For hyper-parameters tuning see Appendix C.

⁶This is the splitting method used in SIGMORPHON’s shared tasks on reinflection (e.g., Cotterell et al., 2018).

⁷We harmonized the train and test features vocabulary, so that the old data bears the new scheme. So the only difference between Original and New is in which forms are included.

Train Set	Test set	Form Split		Lemma Split	
		Acc	Avg ED	Acc	Avg ED
New	New	94.9%	0.15	1.3%	4.66
New	Original	84.7%	0.3	0.3%	4.39
Original	New	35.2%	1.36	0.0%	6.22
Original	Original	99.3%	0.01	0.0%	6.13

Table 4: Accuracy (**Acc**, higher is better) and Average Edit Distance (**Avg ED**, lower is better) for morphological reinflection on different train-test combinations.

Results and Analysis Table 4 presents the model’s performance for all train-test combinations. It shows that the model’s performance on the new data (top line combination) is largely on par comparing to its performance over training and testing on UniMorph’s original data (bottom combination). However, the model generalizes poorly from the original partial data to the forms in our test set which reflect the entire Georgian inflectional system. Generalization from our data to UniMorph’s set is a lot better. The results also show that the splitting method is crucial for success of the model, as it inflects easily to unseen forms, but much harder when inflecting forms in a previously unseen lemma.⁸ These results corroborate the results of Goldman et al. (2021) regarding the difficulty of lemma-split data. Although the accuracy over the lemma split data is negligible, the average edit distance in that case points again to the conclusion that generalization from UniMorph to our data is harder than the other way around.

Error Analysis To provide insights into the challenge of reinflecting morphologically complex forms, we manually sampled the erroneous output of the model trained and tested over our lemma-split data, to draw insights on the points of failure. In many cases the model succeeded in copying and modifying the verb stem, but failed to output the other morphemes correctly. Sometimes the errors were due to inflection to an incorrect TAM combination of the same lexeme, and sometimes the inflection was done to the correct TAM but to a different derivationally-related lemma (e.g. change of voice in addition to the change of TAM). We conclude that the fact that our datasets include lemmas from diverse classes that may have derivational relations makes the inflection task significantly harder. Interestingly, the model managed to predict the correct subject and object affixes most of the time.

⁸For learning curves on the splits see Appendix A.

6 Conclusion

This paper proposes a transition of the UniMorph annotation standard to a layered hierarchical annotation of features. This revised schema caters for complex marking phenomena including multiple pronominal agreement. We apply it to Georgian, and construct a corresponding new dataset that is large, balanced, complete with respect to grammatical phenomena in the Georgian verb system and verified by native-speakers. Our experiments with a standard reinflection model on the old and new Georgian datasets shows that the old UniMorph dataset does not generalize well to the new test-set, due to its partial coverage. This work is intended to encourage the community to extend the annotation of different languages to include phenomena such as polypersonal agreement and others that can be dealt with using a hierarchical annotation, ultimately leading to more complete and consistent benchmarks for studying non-trivial and less-explored areas of computational morphology.

Acknowledgements

The first author would like to thank the native Georgian speakers: Simon Guriel, Silvia Guriel-Agiashvili and Nona Atanelov for their invaluable help in the data annotation process. This research was funded by the European Research Council under the European Union’s Horizon 2020 research and innovation programme (grant agreement No. 677352) and by a research grant from the Ministry of Science and Technology (MOST) of the Israeli Government, for which we are grateful.

References

- Stephen R Anderson. 1992. *A-morphous morphology*. 62. Cambridge University Press.
- Joan Bresnan, Ash Asudeh, Ida Toivonen, and Stephen Wechsler. 2015. *Lexical-functional syntax*. John Wiley & Sons.
- 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.

- Nino Doborjginidze and Irina Lobzhanidze. 2012. Corpus of the georgian language - morphological analyzer. <http://corpora.iliauni.edu.ge/?q=search-words>.
- Nino Doborjginidze and Irina Lobzhanidze. 2016. *Corpus of the Georgian Language*. In *Proceedings of the XVII EURALEX International Congress*, pages 328–335, Tbilisi, Georgia. Ivane Javakhishvili Tbilisi University Press.
- Gerald Gazdar, Ewan Klein, Geoffrey Pullum, and Ivan Sag. 1989. *Generalized phrase structure grammar*. *Philosophical Review*, 98(4):556–566.
- Omer Goldman, David Guriel, and Reut Tsarfaty. 2021. (un)solving morphological inflection: Lemma overlap artificially inflates models’ performance.
- B. G. Hewitt. 1995. *Georgian a structural reference grammar / B.G. Hewitt*. London Oriental and African language library, v. 2. John Benjamins Pub., Amsterdam.
- Mark Johnson. 1988. *Attribute-value logic and the theory of grammar*. Center for the Study of Language and Information.
- 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.
- Arya D. McCarthy, Christo Kirov, Matteo Grella, Amrit Nidhi, Patrick Xia, Kyle Gorman, Ekaterina Vylomova, Sabrina J. Mielke, Garrett Nicolai, Miikka Silfverberg, Timofey Arkhangelskiy, Nataly Krizhanovsky, Andrew Krizhanovsky, Elena Klyachko, Alexey Sorokin, John Mansfield, Valts Ernštreits, Yuval Pinter, Cassandra L. Jacobs, Ryan Cotterell, Mans Hulden, and David Yarowsky. 2020. *UniMorph 3.0: Universal Morphology*. In *Proceedings of the 12th Language Resources and Evaluation Conference*, pages 3922–3931, Marseille, France. European Language Resources Association.
- Carl Pollard and Ivan A Sag. 1994. *Head-driven phrase structure grammar*. University of Chicago Press.
- Stuart M Shieber. 2003. *An introduction to unification-based approaches to grammar*. Microtome Publishing.
- Miikka Silfverberg and Mans Hulden. 2018. *An encoder-decoder approach to the paradigm cell filling problem*. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 2883–2889, Brussels, Belgium. Association for Computational Linguistics.
- John Sylak-Glassman. 2016. The composition and use of the universal morphological feature schema (uni-morph schema). *Johns Hopkins University*.
- Reut Tsarfaty, Dan Bareket, Stav Klein, and Amit Seker. 2020. *From SPMRL to NMRL: What did we learn (and unlearn) in a decade of parsing morphologically-rich languages (MRLs)?* In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 7396–7408, Online. Association for Computational Linguistics.
- Reut Tsarfaty, Djamé Seddah, Yoav Goldberg, Sandra Kuebler, Yannick Versley, Marie Candito, Jennifer Foster, Ines Rehbein, and Lamia Tounsi. 2010. *Statistical parsing of morphologically rich languages (SPMRL) what, how and whither*. In *Proceedings of the NAACL HLT 2010 First Workshop on Statistical Parsing of Morphologically-Rich Languages*, pages 1–12, Los Angeles, CA, USA. Association for Computational Linguistics.
- Shijie Wu, Ryan Cotterell, and Mans Hulden. 2020. *Applying the transformer to character-level transduction*.

A Learning Curves

Fig. 1 exemplifies the sufficiency of our dataset for training an inflection model on form-split data as doubling the data amount from 4,000 to 8,000 yields relatively minor improvement. It also shows that for the lemma-split data, the model completely fails. It starts improve marginally with more than 2,000 examples, although its performance remains far from satisfactory. This leaves room for exploration of bootstrapping and augmentation methods or more sophisticated modeling to improve results.




Figure 1: Inflection accuracy over *form-split* and *lemma-split* test sets as a function of train set size.

B Tech-Spec

All algorithms described in the paper were executed on a single machine equipped with one NVIDIA TITAN Xp GPU, 16 Intel i7-6900K(3.20GHz) CPUs and 126GB RAM. Since the LSTM algorithm was implemented on DyNet, there was no need of the GPU, and all the calculations were done using only the CPU.

C Hyper Parameters

1. Embedding size = 100
2. Hidden state size = 100
3. Attention size = 100
4. Number of LSTM layers = 1

During training, we experimented with several values for the hyper-parameters detailed above. However, for all the combinations we tried, the results barely changed both at the form-split setting and the lemma-split setting.