Language-agnostic measures discriminate inflection and derivation

Coleman Haley **Edoardo M. Ponti** **Sharon Goldwater**
Institute for Language, Cognition and Computaton, University of Edinburgh, United Kingdom
{coleman.haley,eponti,sgwater}@ed.ac.uk

In the field of morphology, a distinction is commonly drawn between *derivations*, processes that form "new" words, and *inflections*, processes that merely create new "forms" of words (Dressler, 1989). While the theoretical nature of this distinction is a subject of ongoing debate, it is widely employed throughout linguistic theory, computational and corpus linguistics, and even psycholinguistics.

Dictionaries and grammars roughly agree on which morphological relationships are inflectional and which are derivational within a language. There is even a degree of cross-linguistic consistency in the constructions which are typically/traditionally considered inflections—e.g., tense marking on verbs is widely considered to be inflectional. This cross-linguistic consistency is highlighted by the development of UniMorph (Batsuren et al., 2022), a resource which annotates inflections across 182 languages using a unified feature scheme. This is despite the fact that UniMorph data is extracted from the Wiktionary open online dictionary[1], which organises constructions into inflections and derivations based on typical traditions for a given language. This is in line with Haspelmath's (in press) view of these terms as *traditional comparative concepts*, being based on the ways in which Western dictionaries and grammar books are traditionally structured.

While linguists have proposed many tests or prototypical properties of these categories, such as derivations producing larger semantic changes or occurring closer to the root of the word, difficulties in producing a cross-linguistically consistent definition have led many researchers to conclude that the inflection–derivation distinction is gradient rather than categorical (e.g., Dressler, 1989) or even to take position that the distinction carries no theoretical weight (Haspelmath, in press). In particular, Haspelmath (in press) argues that many such properties of inflection and derivation are not proven to apply in a consistent way across languages.

One major issue in evaluating these theoretical claims is the lack of large-scale, cross-linguistic evidence based on quantitative measures (rather than subjective tests). While several studies have also computationally operationalised linguistic intuitions about the inflection–derivation distinction, they have been limited in terms of the languages studied, focusing on French (Bonami and Paperno, 2018; Copot et al., in press) and Czech (Rosa and Žabokrtský, 2019). We here expand the set of measures and languages studied to evaluate whether traditional concepts of inflection and derivation relate to their claimed properties cross-linguistically.

We develop a set of four quantitative measures of morphological constructions, including measures of *both* the magnitude and the variability of the changes introduced by each construction. Crucially, our measures can be computed directly from a linguistic corpus, allowing us to consistently operationalise them across many languages and morphological constructions. That is, given a particular morphological construction (such as "the nominative plural in German") and examples of word pairs that illustrate that construction (e.g., '*Frau*, *Frauen*', '*Kind*, *Kinder*'), we compute four corpus-based measures which quantify the idea that derivations produce *larger* and *more variable* changes to words compared to inflections. We then ask whether, for a given construction, knowing just these measures is sufficient to predict its inflectional versus derivational status in UniMorph.

In particular we consider for each construction:

- $\|\Delta_{form}\|$, the average edit distance between the base and constructed forms,
- $\|\Delta_{distribution}\|$, the Euclidean distance between the distributional embeddings of the base and constructed forms,
- $\text{var}(\Delta_{form})$, the average edit distance between the edit sequences between base and

[1] https://en.wiktionary.org

Features	Logistic	MLP
Majority class (Inflection)	0.57	–
$\|\Delta_{distribution}\|$	0.67	0.68
$\|\Delta_{form}\|$	0.59	0.60
var($\Delta_{distribution}$)	0.76	0.76
var(Δ_{form})	0.71	0.71
Form/distribution magnitude*	0.66	0.67
Form/distribution variability*	0.84	0.84
Form magnitude/variability*	0.70	0.75
Distribution magnitude/variability*	0.77	0.77
All measures*	**0.86**	**0.90**

Table 1: Accuracy in reconstructing Unimorph's inflection–derivation distinction by various supervised classifiers.

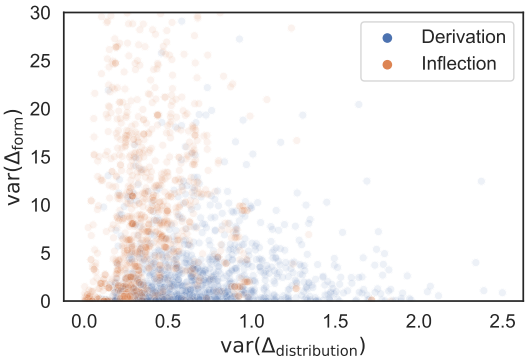

Figure 1: Our two most predictive measures for inflectional and derivational constructions in UniMorph. While these measures can be used to correctly classify 84% of UniMorph constructions, they display a clearly gradient mapping onto the categories.

constructed forms within a construction,

- var($\Delta_{distribution}$), the total variance of the difference vectors between base and constructed form in the distributional embedding space.

If, across languages belonging to different language families and morphological typologies, the UniMorph annotations can be predicted with high accuracy based on our measures, this would indicate that traditional concepts of inflection and derivation *do* correspond to intuitions about the different *types* of changes inflection and derivation induce.

To explore this, we train a logistic regression classifier and a multilayer perceptron (MLP). Since we are interested in the cross-linguistic consistency of our four predictors, the models are not given access to the input language or any of its typological features. In experiments on 26 languages[2] (including five from non-Indo-European families) and 2,772 constructions, we find that both models are able to predict with high accuracy whether a held-out construction is listed as inflection or derivation in UniMorph (86% and 90%, respectively, for the two models, compared to a majority-class baseline of 57%). We additionally find that our distributional measures alone are more predictive than our formal ones, and our variability measures alone are more predictive than our magnitude ones; still, combining all four features yields the best results.

We also identify *how prototypical* various categories of inflections are in terms of our measures. We determine that inherent inflectional meanings are particularly likely to be classified as derivation by our model, in line with Booij's (1996) characterisation of inherent inflection as non-canonical.

We provide initial evidence about non-Indo-European languages, obtaining 82% accuracy compared to 91% for Indo-European languages. While still indicating generalisation, this suggests that the application of the inflection–derivation distinction to non-Indo-European languages may be less consistent as suggested by Haspelmath (in press). For example, Turkish is a highly agglutinative language with, in traditional descriptions, an exceptionally rich inflectional system—reflected by an extremely large number of inflectional constructions and relatively small number of derivations in our dataset. Our classifier over-uses the label derivation for this language, suggesting a degree of mis-alignment with the way linguists typically operationalise inflection and derivation in this language.

Nevertheless, together these results provide large-scale cross-linguistic evidence that, despite the apparent difficulty in designing diagnostic tests for inflection and derivation, these concepts are nevertheless associated with distinct and measurable formal and distributional signatures that behave consistently across a variety of languages. Further analysis of our results does not, however, support the view of these concepts as clearly discrete categories. While our measures largely discriminate inflection and derivation, we still find many constructions near the model's decision boundary between the two categories, indicating a gradient, rather than categorical, distinction (Figure 1).

[2] cat, ces, dan, deu, eng, ell, fin, fra, gle, hun, hye, ita, kaz, lat, lav, mon, nob, nld, pol, por, ron, rus, spa, swe, tur, ukr

References

Khuyagbaatar Batsuren, Omer Goldman, Salam Khalifa, Nizar Habash, Witold Kieraś, Gábor Bella, Brian Leonard, Garrett Nicolai, Kyle Gorman, Yustinus Ghanggo Ate, Maria Ryskina, Sabrina Mielke, Elena Budianskaya, Charbel El-Khaissi, Tiago Pimentel, Michael Gasser, William Abbott Lane, Mohit Raj, Matt Coler, Jaime Rafael Montoya Samame, Delio Siticonatzi Camaiteri, Esaú Zumaeta Rojas, Didier López Francis, Arturo Oncevay, Juan López Bautista, Gema Celeste Silva Villegas, Lucas Torroba Hennigen, Adam Ek, David Guriel, Peter Dirix, Jean-Philippe Bernardy, Andrey Scherbakov, Aziyana Bayyr-ool, Antonios Anastasopoulos, Roberto Zariquiey, Karina Sheifer, Sofya Ganieva, Hilaria Cruz, Ritván Karahóǧa, Stella Markantonatou, George Pavlidis, Matvey Plugaryov, Elena Klyachko, Ali Salehi, Candy Angulo, Jatayu Baxi, Andrew Krizhanovsky, Natalia Krizhanovskaya, Elizabeth Salesky, Clara Vania, Sardana Ivanova, Jennifer White, Rowan Hall Maudslay, Josef Valvoda, Ran Zmigrod, Paula Czarnowska, Irene Nikkarinen, Aelita Salchak, Brijesh Bhatt, Christopher Straughn, Zoey Liu, Jonathan North Washington, Yuval Pinter, Duygu Ataman, Marcin Wolinski, Totok Suhardijanto, Anna Yablonskaya, Niklas Stoehr, Hossep Dolatian, Zahroh Nuriah, Shyam Ratan, Francis M. Tyers, Edoardo M. Ponti, Grant Aiton, Aryaman Arora, Richard J. Hatcher, Ritesh Kumar, Jeremiah Young, Daria Rodionova, Anastasia Yemelina, Taras Andrushko, Igor Marchenko, Polina Mashkovtseva, Alexandra Serova, Emily Prud'hommeaux, Maria Nepomniashchaya, Fausto Giunchiglia, Eleanor Chodroff, Mans Hulden, Miikka Silfverberg, Arya D. McCarthy, David Yarowsky, Ryan Cotterell, Reut Tsarfaty, and Ekaterina Vylomova. 2022. UniMorph 4.0: Universal Morphology. In *Proceedings of the Thirteenth Language Resources and Evaluation Conference*, pages 840–855, Marseille, France. European Language Resources Association.

Olivier Bonami and Denis Paperno. 2018. Inflection vs. derivation in a distributional vector space. *Lingue e linguaggio*, 17(2):173–196.

Geert Booij. 1996. Inherent versus contextual inflection and the split morphology hypothesis. In *Yearbook of Morphology 1995*, pages 1–16. Springer.

Maria Copot, Timothee Mickus, and Olivier Bonami. in press. Idiosyncratic frequency as a measure of derivation vs. inflection. *Journal of Language Modelling*.

Wolfgang U Dressler. 1989. Prototypical differences between inflection and derivation. *STUF-Language Typology and Universals*, 42(1):3–10.

Martin Haspelmath. in press. Inflection and derivation as traditional comparative concepts. *Linguistics*.

Rudolf Rosa and Zdeněk Žabokrtský. 2019. Attempting to separate inflection and derivation using vector space representations. In *Proceedings of the Second International Workshop on Resources and Tools for Derivational Morphology*, pages 61–70, Prague, Czechia. Charles University, Faculty of Mathematics and Physics, Institute of Formal and Applied Linguistics.