

Corpus-based measures discriminate inflection and derivation cross-linguistically

Coleman Haley Edoardo M. Ponti Sharon Goldwater

School of Informatics

University of Edinburgh, United Kingdom

{coleman.haley, eponti, sgwater}@ed.ac.uk

In the field of morphology, a distinction is commonly drawn between *derivation*, processes that form “new” words, and *inflection*, processes that merely create new “forms” of words (Booij, 2007). While the theoretical nature of this distinction is a subject of ongoing debate, it is widely employed throughout theoretical linguistics, 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, which now includes derivations from 30 languages, is extracted from the Wiktionary open online dictionary¹, which organises constructions into inflections and derivations based on typical traditions for a given language, rather than features or properties of constructions.

Despite this relative consistency at the level of annotation, there is considerable disagreement amongst linguists about the fundamental properties that might underlie or explain these traditional categorizations—such as the degree of syntactic or semantic change. As an example, Plank (1994) covers no fewer than 28 tests for these categories. Upon applying them to just 6 English morphological constructions, Plank (1994) finds significant contradictions between the criteria. Such 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).

One major issue in evaluating these theoretical claims is the lack of large-scale, cross-linguistic evidence based on quantitative measures (rather than subjective diagnostic tests). Work in theoretical linguistics has established that the intuitions underlying diagnostic tests can be problematic in specific cases, yet it remains unclear to what extent they successfully capture the traditional concepts of inflection and derivation on a large scale—perhaps measures based on these subjective tests can indeed be used to classify the vast majority of morphological relationships across languages in a way that is consistent with traditional distinctions. If so, a large-scale empirical study could also provide evidence regarding the gradient of the inflection–derivation distinction.

We take inspiration from both linguistic theory and prior computational work on single languages to 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—two based on orthographic form and two based on distributional semantics characteristics—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:

¹<https://en.wiktionary.org>

- $\|\Delta_{form}\|$, the average edit distance between the base and constructed forms of a construction,
- $\|\Delta_{distribution}\|$, the Euclidean distance between the FastText (Bojanowski et al., 2017) distributional embeddings of the base and constructed forms,
- $\text{var}(\Delta_{form})$, the average edit distance between the edit sequences between corresponding base and constructed forms within a construction,
- $\text{var}(\Delta_{distribution})$, the total variance of the difference vectors between base and constructed form in the FastText embedding space.

If, across a variety of languages, belonging to different grammatical traditions, language families, and morphological typologies, the UniMorph annotations can be predicted with high accuracy based on our four measures, this would provide evidence 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.

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 (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%). Together, these results provide large-scale cross-linguistic evidence that, despite the apparent difficulty in designing diagnostic tests to definitively identify inflectional versus derivational relations, the comparative concepts of inflection and derivation are nevertheless associated with distinct and measurable formal and distributional signatures that behave relatively consistently across a variety of languages.

We also identify in a quantitative way *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

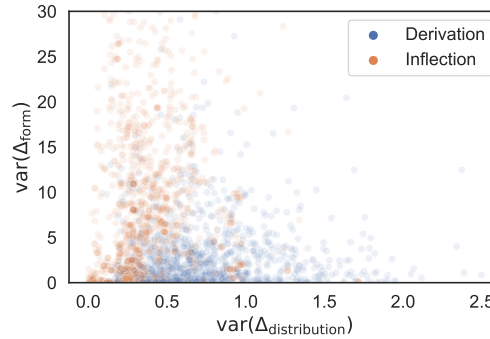


Figure 1: Our two most predictive measures for inflectional and derivational constructions in UniMorph. While these measures can be used to correctly classify 84% of UniMorph constructions, the inflection–derivation distinction appears gradient rather than categorical with respect to them.

(1996) characterisation of inherent inflection as non-canonical.

Finally, we note that while there is a high degree of consistency in the use of the terms inflection and derivation in terms of our measures, we still find many constructions near the model’s decision boundary between the two categories, indicating a gradient, rather than categorical, distinction. This gradient region is relatively small but does not suggest inflection and derivation as categories *naturally clustering* in our measures.

Acknowledgements

We would like to thank Paul Smolensky, Itamar Kastner, Frank Mollica, and Kate McCurdy for feedback on early versions of this work. This work was supported in part by the UKRI Centre for Doctoral Training in Natural Language Processing, funded by the UKRI (grant EP/S022481/1) at the University of Edinburgh, School of Informatics and School of Philosophy, Psychology, & Language Science.

References

- K. Batsuren, O. Goldman, S. Khalifa, N. Habash, W. Kieraś, G. Bella, B. Leonard, G. Nicolai, K. Gorman, Y. G. Ate, M. Ryskina, S. Mielke, E. Budianskaya, C. El-Khaissi, T. Pimentel, M. Gasser, W. A. Lane, M. Raj, M. Coler, J. R. M. Samame, D. S. Camaiteri, E. Z. Rojas, D. López Francis, A. Oncevay, J. López Bautista, G. C. S. Villegas, L. T. Hennigen, A. Ek, D. Guriel, P. Dirix, J.-P. Bernardy, A. Scherbakov, A. Bayyr-ool, A. Anastasopoulos, R. Zariquiey, K. Sheifer, S. Ganieva, H. Cruz, R.

- Karahóga, S. Markantonatou, G. Pavlidis, M. Plugarov, E. Klyachko, A. Salehi, C. Angulo, J. Baxi, A. Krizhanovsky, N. Krizhanovskaya, E. Salesky, C. Vania, S. Ivanova, J. White, R. H. Maudslay, J. Valvoda, R. Zmigrod, P. Czarnowska, I. Nikkarinen, A. Salchak, B. Bhatt, C. Straughn, Z. Liu, J. N. Washington, Y. Pinter, D. Ataman, M. Wolinski, T. Suhardijanto, A. Yablonskaya, N. Stoehr, H. Dolatian, Z. Nuriah, S. Ratan, F. M. Tyers, E. M. Ponti, G. Aiton, A. Arora, R. J. Hatcher, R. Kumar, J. Young, D. Rodionova, A. Yemelina, T. Andrushko, I. Marchenko, P. Mashkovtseva, A. Serova, E. Prud'hommeaux, M. Nepomniashchaya, F. Giunchiglia, E. Chodroff, M. Hulden, M. Silfverberg, A. D. McCarthy, D. Yarowsky, R. Cotterell, R. Tsarfaty, and E. 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.
- P. Bojanowski, E. Grave, A. Joulin, and T. Mikolov. 2017. [Enriching word vectors with subword information](#). *Transactions of the Association for Computational Linguistics*, 5:135–146.
- G. Booij. 1996. Inherent versus contextual inflection and the split morphology hypothesis. In *Yearbook of Morphology 1995*, pages 1–16. Springer.
- G. Booij. 2007. [Inflection](#). In *The Grammar of Words: An Introduction to Linguistic Morphology*. Oxford University Press.
- W. U. Dressler. 1989. Prototypical differences between inflection and derivation. *STUF-Language Typology and Universals*, 42(1):3–10.
- M. Haspelmath. in press. Inflection and derivation as traditional comparative concepts. *Linguistics*.
- F. Plank. 1994. Inflection and derivation. In *The Encyclopedia of Language and Linguistics*, pages 1671–1679. Elsevier Science and Technology, Amsterdam.