

Morphological Processing of Low-Resource Languages: Where We Are and What’s Next

Adam Wiemerslage^{‡*} Miikka Silfverberg^{‡*} Changbing Yang[‡] Arya D. McCarthy[‡]
Garrett Nicolai[‡] Eliana Colunga[‡] Katharina Kann[‡]

[‡]University of Colorado Boulder [‡]University of British Columbia

[‡]Johns Hopkins University

Abstract

Automatic morphological processing can aid downstream natural language processing applications, especially for low-resource languages, and assist language documentation efforts for endangered languages. Having long been multilingual, the field of computational morphology is increasingly moving towards approaches suitable for languages with minimal or no annotated resources. First, we survey recent developments in computational morphology with a focus on low-resource languages. Second, we argue that the field is ready to tackle the logical next challenge: understanding a language’s morphology from raw text alone. We perform an empirical study on a truly unsupervised version of the paradigm completion task and show that, while existing state-of-the-art models bridged by two newly proposed models we devise perform reasonably, there is still much room for improvement. The stakes are high: solving this task will increase the language coverage of morphological resources by a number of magnitudes.

1 Introduction

Automatic morphological processing tools have the potential to drastically speed up language documentation (Moeller et al., 2020) and thereby help combat the language endangerment crisis (Austin and Sallabank, 2011). Explicit morphological information also benefits myriad NLP tasks, such as parsing (Hohensee and Bender, 2012; Seeker and Çetinoğlu, 2015), language modeling (Blevins and Zettlemoyer, 2019; Park et al., 2021; Hofmann et al., 2021), and machine translation (Dyer et al., 2008; Tamchyna et al., 2017).

For low-resource languages, valuable morphological resources are typically small or non-existent. Of late, the field of computational morphology has increased its efforts to extend the coverage of multilingual morphological resources (Kirov et al., 2016,

2018; McCarthy et al., 2020a; Metheniti and Neumann, 2020). Simultaneously, there has been a revival of minimally supervised and unsupervised models for morphological tasks, such as segmentation (Eskander et al., 2019), inflection (Kann et al., 2017b), and lemmatization (Bergmanis and Goldwater, 2019). Given the speed of recent developments, it is important to reflect on where we are as a field and what future challenges lie ahead.

To this end, we survey recent computational morphology: we review existing multilingual resources (§2) and tasks and systems (§3), with a focus on low-resource languages. Given recent developments in unsupervised segmentation, low-resource morphological inflection, and unsupervised morphological paradigm completion (Jin et al., 2020; Erdmann et al., 2020)—which we argue is not fully unsupervised—we believe the community is poised for the next logical step: inferring a language’s morphology purely from raw text.

In §4, we formalize a new task: *truly unsupervised morphological paradigm completion* (*tUMPC*). We then introduce a pipeline with two novel components (§4.3): one model for aligning paradigm slots across lexemes and another for predicting the slots of observed forms. With these, we assess several state-of-the-art models and the influence of different types of unlabeled corpora within the framework of *tUMPC*. While existing methods leave room for improvement, they perform reasonably enough to support our argument that inferring a language’s morphology from raw text is within reach and worthy of community efforts.

To summarize, we present the following contributions: (i) a survey of tasks and systems in computational morphology with a focus on low-resource languages; (ii) models for the tasks of paradigm slot alignment and slot prediction, (iii) a formalization of the task of truly unsupervised morphological paradigm completion and (iv) an evaluation of state-of-the-art approaches and differ-

*The first two authors contributed equally.

ent corpora within the framework of this task. Our code and data are publicly available.¹

2 Morphological Resources

Manually created resources are necessary for developing and evaluating NLP systems. They also serve as a basis for research questions in a multilingual context (Pimentel et al., 2019; Wu et al., 2019).² Below, we review the two largest active multilingual resources for morphology and a number of language-specific resources.

Background and Notation The canonical form of a word is called its *lemma*, and the set of all surface forms of a lemma is referred to as that lemma’s *paradigm*. As is common, we formally write the paradigm of a lemma ℓ as:

$$\pi(\ell) = \langle f(\ell, \vec{t}_\gamma) \rangle_{\gamma \in \Gamma(\ell)}, \quad (1)$$

with $f : \Sigma^* \times \mathcal{T} \rightarrow \Sigma^*$ defining a mapping from a tuple consisting of the lemma and a vector $\vec{t}_\gamma \in \mathcal{T}$ of morphological features to the corresponding inflected form. Σ is an alphabet of discrete symbols: the characters used in the language of lemma ℓ . $\Gamma(\ell)$ is the set of slots in ℓ ’s paradigm.

UniMorph The UniMorph project (Sylak-Glassman et al., 2015a,b; Kirov et al., 2016) is a database of triples organized into paradigms, where each triple represents a word as its lemma ℓ , morpho-syntactic description \vec{t}_γ , and surface form $f(\ell, \vec{t}_\gamma)$. An English example triple is:

mutate mutates V;3;SG;PRS

This structure provides training data for inflection generation, lemmatization, or paradigm completion. The most recent version of UniMorph (McCarthy et al., 2020a) includes 118 languages and 14.8 million triples, with more languages under development. As it is semi-automatically created, issues have been noted—particularly, it is a convenience sample across languages (Gorman et al., 2019; Malouf et al., 2020). Still, related efforts validate themselves using UniMorph, including Metheniti and Neumann (2020)—another Wiktionary-derived resource for morphology. Wik-inflection captures segmentation information (§3.2) from Wiktionary templates, though the authors note some limits in the morphological tags that are extracted to accompany these.

¹<https://github.com/Adamits/tUMPC>

²This approach has been criticized by Malouf et al. (2020) due to incompleteness and quality of existing resources.

Universal Dependencies Whereas UniMorph contains type-level annotations, the Universal Dependencies project (UD) is a resource of token-level annotations. As of writing, the latest release (v2.8; Zeman et al., 2021) spans 114 languages, typically semi-automatically extracted from existing corpora, sometimes with less comprehensive annotations (Malaviya et al., 2018). The structure is useful for morphological tagging (§3.1) at the sentence level (Goldman and Tsarfaty, 2021), and several languages have parallel text, enabling evaluation of projection-based approaches for morphology induction, parsing, and other tasks (Yarowsky et al., 2001; Rasooli and Collins, 2017).

Mapping between UniMorph and Universal Dependencies The UD2 morphological annotations borrow several features from UniMorph.³ Consequently, there is great harmony between the two schemas. A deterministic mapping (McCarthy et al., 2018) has shown the synergy; for instance, Bergmanis and Goldwater (2019) augment a contextual tagger with UniMorph inflection tables.

Language-Specific Resources Throughout the years, many language-specific morphological resources have been created. These include corpora and treebanks like the morphologically annotated corpus for Emirati Arabic by Khalifa et al. (2018). Resources also come in the form of morphological databases, such as CELEX for Dutch, English and German (Baayen et al., 1996), or morphological analyzers, such as the Paraguayan Guaraní analyzer presented by Zueva et al. (2020).

Creation of morphological resources is an ongoing effort which in recent years has increasingly focused on low-resource languages. Several conferences and workshops like LREC (Calzolari et al., 2020), SIGMORPHON (Nicolai et al., 2021), ComputEL (Arppe et al., 2021), AmericasNLP (Mager et al., 2021), PYLO (Klavans, 2018) and FSMNLP (Maletti and Constant, 2011) have presented and continue to present language-specific tools and datasets for computational morphology.

3 Where We Are: Tasks and Systems

3.1 Morphological Tagging

Morphological tagging is a sequence-labeling task similar to part-of-speech (POS) tagging. As a token-level task, it considers words in context.

³<http://universaldependencies.org/v2/features.html#comparison-with-unimorph>

Given a sentence, it consists of assigning to each word $f(\ell, \vec{t}_\gamma)$ a morphosyntactic description (MSD), i.e., a tag representing the morphological features \vec{t}_γ it expresses. For instance, in the sentence *The virus mutates*, the word *mutates* would be assigned the tag `v;3;SG;PRS`. Morphological tagging was featured in the SIGMORPHON 2019 shared task (McCarthy et al., 2019).

Systems A leading non-neural morphological tagger is MARMOT (Mueller et al., 2013), a higher-order conditional random field (CRF; Lafferty et al., 2001) tagger. Of late, LSTM (Hochreiter and Schmidhuber, 1997) and Transformer (Vaswani et al., 2017) models have been used for tagging (Heigold et al., 2016, 2017; Nguyen et al., 2021).

For low-resource languages, both projection-based approaches (Buys and Botha, 2016) and cross-lingual transfer approaches via multitask training (Cotterell and Heigold, 2017) have been developed. 16 systems were submitted to the SIGMORPHON 2019 shared task⁴ (McCarthy et al., 2019), which featured 66 languages. The winning team (Kondratyuk, 2019) built a tagger based on multilingual BERT (Devlin et al., 2019), thus employing cross-lingual transfer; for other systems, we refer the reader to the shared task overview. The largest multilingual morphological tagging effort to date is that by Nicolai et al. (2020) who build morphological analyzers for 1108 languages using projection from a high-resource to a low-resource language via the aligned text in the JHU Bible Corpus (McCarthy et al., 2020b).

3.2 Morphological Segmentation

The goal of morphological segmentation (Goldsmith, 2010) is to split words into their smallest meaning-bearing units: morphemes. We discuss both surface and canonical segmentation here.

3.2.1 Surface Segmentation

During surface segmentation, a word is split into morphemes in a way such that the concatenation of all parts exactly results in the original word. An example (with “*” marking boundaries) is:

mutates → mutate * s

Surface segmentation was the focus of the Morpho Challenge from 2005 to 2010 (Kurimo et al., 2010).

⁴The task is concerned with joint lemmatization and tagging, but systems can be used for separate tagging as well.

The competition featured datasets in Finnish, Turkish, German, English, and Arabic. Additionally, segmentation was a track (alongside morphological analysis and generation) of LowResourceEval-2019 (Klyachko et al., 2020), a shared task which featured four low-resource languages from Russia. The shared task overview lists morphological resources for other Russian languages.

Systems Many approaches to this task are unsupervised. Harris (1970) identifies morpheme boundaries in English based on the frequency of characters at the end of a word. LINGUISTICA (Goldsmith, 2001) finds sets of stems and suffixes that represent the minimum description length of the data. MORFESSOR (Creutz and Lagus, 2002) introduces a family of probabilistic models for identifying morphemes, which have seen wide use, including variations of the original model (Virpioja et al., 2009; Smit et al., 2014). Lignos et al. (2009) learn rewrite rules that can explain many types in the corpus. Poon et al. (2009) apply a CRF to unsupervised segmentation by learning parameters with contrastive estimation (Smith and Eisner, 2005). Incorporating semantic similarity between related words that form “chains” has also been shown to be effective (Narasimhan et al., 2015). Monson et al. (2007) propose a segmentation algorithm that exposes the properties of partial morphological paradigms in order to learn segments. Xu et al. (2018) iteratively refine segments according to their distribution across paradigms. They filter unreliable paradigms with statistically reliable ones, and induce segments with the proposed partial paradigms. Both systems can only model suffix concatenation. Xu et al. (2020) follow a similar strategy, but incorporate language typology, expanding beyond suffixes, and outperform Xu et al. (2018). MorphAGram (Eskander et al., 2020) is a publicly available tool for unsupervised segmentation based on adaptor grammars (Johnson et al., 2007).

Supervised (Creutz and Lagus, 2005; Ruokolainen et al., 2013; Cotterell et al., 2015) and semi-supervised systems (Ruokolainen et al., 2014) also exist. Non-neural systems are often based on CRFs. Ruokolainen et al. (2013) focus explicitly on low-resource settings and perform experiments on Arabic, English, Hebrew, Finnish, and Turkish with training set sizes as small as 100 instances.

Neural models have also been applied to surface segmentation: Wang et al. (2016) obtain strong re-

sults with window LSTM neural networks in the high-resource setting, [Seker and Tsarfaty \(2020\)](#) introduce a pointer network ([Vinyals et al., 2015](#)) for segmentation and tagging, and [Micher \(2017\)](#) propose a segmental RNN ([Kong et al., 2015](#)) for segmentation and tagging of Inuktitut. [Kann et al. \(2018b\)](#) explore LSTM-based sequence-to-sequence (seq2seq) models for segmentation in combination with data augmentation, multitask and multilingual training; they evaluate on datasets they introduce for four low-resource Mexican languages. [Eskander et al. \(2019\)](#) apply an *unsupervised* approach based on adaptor grammars to the same languages; it outperforms supervised methods in some cases. [Sorokin \(2019\)](#) show that CNNs outperform RNN-based models on that data as well as on North Sámi ([Grönroos et al., 2019](#)).

Additional contributions have been made by [Yarowsky and Wicentowski \(2000\)](#), [Schone and Jurafsky \(2001\)](#), and [Clark \(2001\)](#). Linguistically informed approaches show demonstrable value compared to approaches like BPE; see [Church \(2020\)](#) and [Hofmann et al. \(2021\)](#). Still, not all morphological phenomena are suited for a segmentation-based analysis, as in fusional morphology that sometimes leaves ambiguity as to where a morpheme boundary lies; indeed in some cases there is no consensus among linguists as to the proper segmentation of a word. Therefore, (especially surface) segmentation is not necessarily meaningful for all languages.

3.2.2 Canonical Segmentation

Canonical segmentation is more complex: its aim is to jointly split a word into morphemes and to undo the orthographic changes which have occurred during word formation. As a result, each word is segmented into its *canonical* morphemes. While often not being modeled this way in practice, the task can be seen as the following two-step process:

manic → maniaic → mania * ic

Systems The state-of-the-art pre-neural system is the CRF-based model by [Cotterell et al. \(2016c\)](#), which is jointly trained on segmentation and restoration of orthographic changes. The unsupervised system of [Bergmanis and Goldwater \(2017\)](#) builds upon MorphoChains ([Narasimhan et al., 2015](#)). Neural models are typically based on seq2seq architectures: [Kann et al. \(2016\)](#) use a seq2seq GRU and a feature-based reranker. Like [Cotterell et al. \(2016c\)](#), they evaluate on German, English, and Indonesian. [Ruzsics and Samardžić](#)

(2017) use a similar system, but add a language model over canonical segments and do not require external resources. In addition to German, English, and Indonesian, they evaluate on Chintang, a truly low-resource language spoken in Nepal. [Wang et al. \(2019\)](#) use a character-level seq2seq model for (surface and) canonical segmentation in Mongolian. [Mager et al. \(2020\)](#) show the benefit of copy mechanisms and introduce datasets for two low-resource Mexican languages. [Moeng et al. \(2021\)](#) show that Transformers outperform RNNs for canonical segmentation in four Nguni languages.

3.3 Lemmatization, Inflection, Reinflection

Inflection and reinflection have recently gained popularity in computational morphology by being featured in yearly SIGMORPHON shared tasks ([Cotterell et al., 2016b](#)). They are concerned with generating inflected forms $f(\ell, \vec{t}_\gamma)$ of a lemma ℓ ; the target inflected form can be specified in different ways, depending on the exact task formulation. While the terms *inflection* and *reinflection* are sometimes used synonymously in the literature, *inflection* refers to generating a word form from a given *lemma*, while *reinflection* refers to generation from an *arbitrary* given form in the paradigm. Lemmatization is a special case of reinflection: instead of generating an indicated inflected form, a lemma is produced. As the target form is implicitly determined by the task definition, lemmatization generally does not require tags to indicate which form to generate.

3.3.1 Type-level Versions

Most commonly, lemmatization, inflection and reinflection are type-level tasks. The input consists of an input form together with the target MSD (which can be omitted for lemmatization). The output is the corresponding inflected form, for instance:

mutated V;3;SG;PRS → mutates

The version of reinflection featured in the SIGMORPHON 2016 shared task also provides the MSD of the *source form*, but performance improvements are usually minor ([Cotterell et al., 2016a](#)).

Systems Pre-neural systems for the task include those by [Durrett and DeNero \(2013\)](#) and [Nicolai et al. \(2015\)](#). These systems align lemmas and inflections before extracting character-level transductions for training CRF-inspired models. [Faruqui](#)

et al. (2016) propose the first neural model for morphological inflection, an RNN seq2seq model, but fail to outperform prior approaches on some of the datasets they evaluate on. The breakthrough for neural models was the SIGMORPHON 2016 shared task (Cotterell et al., 2016a), with about one third of the systems being neural: the winning system (Kann and Schütze, 2016a,b) used multi-task training by encoding MSDs together with the character sequence of the source word. This approach has now become the standard for the task, and while a multilingual version of the model by Kann and Schütze (2016a) was submitted to the SIGMORPHON 2021 shared task (Pimentel et al., 2021; Szolnok et al., 2021), the same multitask approach has since been used with other seq2seq models such as Transformers (Wu et al., 2021). Ensembles have been shown to improve performance for inflection (Kann and Schütze, 2016a) and have been systematically studied for the task by Kylliäinen and Silfverberg (2019).

The SIGMORPHON shared tasks on morphological inflection have focused increasingly on low-resource settings. Seq2seq models with hard monotonic attention (Aharoni and Goldberg, 2017), a copy mechanism (Sharma et al., 2018; Singer and Kann, 2020), or both (Makarov et al., 2017; Makarov and Clematide, 2018a,b) obtain great results for training sets as small as 100 examples. Cross-lingual transfer via multitask training was proposed by Kann et al. (2017b) for GRU seq2seq models and has later been used with other architectures, e.g., in the SIGMORPHON 2019 shared task on cross-lingual transfer (McCarthy et al., 2019).

Another approach suitable for low-resource languages is data augmentation. For morphological inflection, this was suggested by several contemporaneous works (Kann and Schütze, 2017; Bergmanis et al., 2017; Silfverberg et al., 2017). In the following years, other augmentation strategies have been developed (Anastasopoulos and Neubig, 2019). The success of data augmentation is mixed, as it is largely dependent on the architecture (*Does it have to learn how to copy or is there a copy mechanism?*) as well as on the quality of the original data, which influences the quality of artificially generated examples.

3.3.2 Token-level Versions

The token-level version of the task is often referred to as lemmatization or inflection *in context*. Here the information about which form to generate is

explicitly given via a sentence context in which the target word should be embedded, e.g.:

mutate – The virus [MASK]. → mutates

A drawback of this formulation is that typically many different inflected forms are possible within the same context: in the given example, *mutates* is the gold solution, but *mutated* would be equally grammatical. To overcome this, multiple gold solutions can be provided (Cotterell et al., 2018). It might be impossible to unambiguously define the target form for some languages if the speaker’s intention is unknown.

Systems Lemmatization in context is arguably easier than inflection or reinflection, as the target form for generation is implicitly defined. Neural models for inflection are seq2seq architectures: Bergmanis and Goldwater (2018) propose Lematus, a character-level LSTM, which they later extend to the low-resource setting by training on labeled data in combination with raw text (Bergmanis and Goldwater, 2019). They explore data settings as small as 1k types each from UD and UniMorph. Zalmout and Habash (2020) use a similar architecture to Lematus but add subword features. Malaviya et al. (2019) present a joint model for tagging and lemmatization and show that joint training benefits low-resource languages. They evaluate on 20 languages, using data from UD. The best lemmatizer in the SIGMORPHON 2019 shared task (McCarthy et al., 2019), UDPipe (Straka et al., 2019), is based on BERT (Devlin et al., 2019).

Inflection in context can be tackled by neural seq2seq models too. Models typically either see a context window around the target word (Makarov and Clematide, 2018c; Kann et al., 2018a; Ács, 2018) and then are optionally trained via multitask training (Kementchedjhieva et al., 2018) or predict the MSD of the form to generate as a first step (Liu et al., 2018). Kementchedjhieva et al. (2018) show that a multilingual model can aid low-resource languages via cross-lingual transfer.

3.4 Paradigm Completion

The paradigm cell filling problem (Ackerman et al., 2009) – also called supervised paradigm completion (Cotterell et al., 2017a) – is yet another inflection task, but differs from the above ones in that the inflected forms for *all* slots $\Gamma(\ell)$ of lemma ℓ ’s paradigm need to be generated and that the input can consist of *one or more* forms.

Systems Many approaches for the paradigm cell filling problem are effectively systems for morphological reinflection and generate all forms of a paradigm individually and from a single input form, e.g., [Silfverberg et al. \(2017\)](#); [Silfverberg and Hulden \(2018\)](#); [Moeller et al. \(2020\)](#). [Kann et al. \(2017a\)](#) propose a model for multi-source inflection, showing that multiple available forms per paradigm can be beneficial for generation, but do not evaluate on paradigm completion. Two notable exceptions which design approaches explicitly for the paradigm cell filling problem are [Cotterell et al. \(2017b\)](#) and [Kann and Schütze \(2018\)](#). [Cotterell et al. \(2017b\)](#) rely on the notion of principal parts ([Finkel and Stump, 2007](#)) to jointly generate all forms in the paradigm. [Kann and Schütze \(2018\)](#) use a transductive training approach, fine-tuning on a paradigm’s *input* forms before generating missing target forms. The latter shows good performance for training sets with as few as 10 paradigms.

3.5 Paradigm Clustering

Paradigm clustering can be seen as a first step towards the unsupervised analysis of a language’s morphology and is typically part of pipelines for unsupervised paradigm completion (§3.6). The goal of paradigm clustering is to group all types in a corpus into (partial) morphological paradigms. For example, the input *The, virus, mutates, after, it, has, mutated* should result in the paradigm cluster (*mutates, mutated*) and 5 singleton clusters. Systems for the task can be evaluated using best-match F1 (BMF1; [Wiemerslage et al., 2021](#)).

Systems Perhaps the seminal work in distributionally-based paradigm clustering is the work of [Yarowsky and Wicentowski \(2000\)](#). Their work predates embedding-based approaches while leveraging both distributional features of context and relative frequency, along with early statistical models of inflection-to-lemma string transduction. For instance, the work succeeds in identifying that the past tense of ‘sing’ is not ‘singed’ but ‘sang’, based on both the distributional signatures of music vs. fire terms in context, as well as the distribution of observed tense *frequency ratios*, where the regular *sing:singed* pairing can also be rejected given its frequency ratio is several standard deviations off of expectation, while the irregular *sing:sang* pairing occurs at nearly exactly the ratio expected. While contextual information has been incorporated in follow-up works ([Schone](#)

and [Jurafsky, 2001](#)) and in recent approaches by means of word embeddings, we do not see much follow-on use of the frequency ratio features, which remain ripe for disambiguation of paradigm members.

Segmentation approaches like [Goldsmith \(2001\)](#), developed to segment words into stems and affixes, can also be used to induce paradigm clusters. [Chan \(2006\)](#) formalizes the notion of a probabilistic paradigm — modeling conditional probabilities of suffixes given paradigms and paradigms given stems. However, they that a segmentation is given, and only model regular morphology for unambiguous words, or those with a known POS. Some segmentation algorithms induce paradigms as a byproduct, as in [Monson et al. \(2007\)](#), [Xu et al. \(2018\)](#) and [Xu et al. \(2020\)](#). These can also be employed as paradigm clustering systems.

Several systems have been proposed for the SIG-MORPHON 2021 shared task ([Wiemerslage et al., 2021](#)). The best performing system ([McCurdy et al., 2021](#)) segments input types with MorphA-Gram ([Eskander et al., 2020](#)), then groups the resulting stems into paradigm clusters. [Yang et al. \(2021\)](#) learn frequent transformation rules and cluster types together that result from rule application.

3.6 Unsupervised Paradigm Completion

Due to the recent progress on supervised morphological tasks, unsupervised paradigm completion (UMPC; or the *paradigm discovery problem* ([Elsner et al., 2019](#))) has recently (re)emerged as a promising way to automatically extend morphological resources such as UniMorph to more low-resource languages. Similar to the supervised version of the task, the goal is to generate the inflected forms corresponding to *all* slots $\Gamma(\ell)$ of lemma ℓ ’s paradigm. However, no morphological annotations are given during training. Two independent works propose similar unsupervised paradigm completion setups. In [Jin et al. \(2020\)](#), the basis of the SIG-MORPHON 2020 shared task ([Kann et al., 2020](#)), the input consists of 1) a corpus in a low-resource language and 2) a list of lemmas from one POS in that language. In [Erdmann et al. \(2020\)](#), the inputs are 1) a corpus and 2) a list of word forms belonging to a single POS. For both, the expected output is the paradigms for the words in the provided list.

As systems are trained without supervision, they cannot output human-readable MSDs and, instead, assign uninterpretable slot identifiers to generated

forms. Thus, evaluation against gold standard data from UniMorph is non-trivial. Jin et al. (2020) propose to evaluate systems via best-match accuracy (BMAcc): the best accuracy among all mappings from pseudo tags to paradigm slots.

Systems State-of-the-art systems for paradigm completion follow a pipeline approach similar to that by Jin et al. (2020): 1) based on the given input forms, they detect transformations which happen during inflection (and sometimes new lemmas), 2) the paradigm structure is detected based on the transformations, and 3) an inflection model is trained to generate missing surface forms. Jin et al. (2020) employ the inflection model by Makarov and Clematide (2018a), while Mager and Kann (2020) use the LSTM pointer-generator model from Sharma et al. (2018), and Singer and Kann (2020) implement a Transformer-based pointer-generator model. The performance across languages is mixed (Kann et al., 2020).

Is the Task Truly Unsupervised? Existing versions of the unsupervised paradigm completion task make small concessions to supervision requirements by providing lists of lemmas or surface forms from a single POS. This simplifies a difficult task, but also makes it less realistic. From the point of view of data availability, this method is not language-agnostic, as many languages do not have the required documentation: many of the world’s languages have fuzzy POS definitions, and no annotated POS corpora. From a language learning perspective, existing methods are closer to L2 than to L1 learning.

Under this framing, UMPC requires only discovering the set of inflection slots for a single paradigm, of a single POS that must be known a priori. The presence of a word list also allows systems to anchor to a privileged form and simplifies paradigm clustering to a retrieval task.

4 What’s Next: Truly Unsupervised Paradigm Completion

4.1 Motivation

We introduce a version of UMPC that more strictly removes human intervention. By removing the input lexicon and evaluating more than one POS, we minimize any prior human involvement with the data and better evaluate a system’s ability to generalize. This means that our only input is a raw text corpus, and it introduces two challenges. 1) We

must model the entire training corpus, rather than a filtered set of words. 2) We must predict which slots to generate at test time. We design test sets to evaluate these problems, ensuring they include paradigms from at least two POS, and prompt for input forms in context, half of which are unseen in the training corpora, so systems can infer the input word POS. We refer to this version of the task as *truly* unsupervised paradigm completion (*tUMPC*).

4.2 Data and Languages

Languages We select three development languages (English, Finnish, and Swedish) and four test languages (German, Greek, Icelandic, and Russian). We select our test languages to maximize orthographic and typological diversity, given three constraints: (1) a large number of available paradigms in UniMorph, (2) two or more POS in UniMorph, and (3) no known issues with the UniMorph data such as large numbers of missing forms. (We exclude all paradigms containing multiword forms.) We note that this yields a set of test languages that are all Indo-European, though it spans three different orthographies.

Raw Text Corpora We experiment on two corpora: the JHU Bible Corpus (McCarthy et al., 2020b) and a child-directed corpus we create by digitizing children’s books. While many studies in computational morphology focus on transcripts of child-directed speech from databases like CHILDES (MacWhinney, 2014), child-directed books are part of parent’s child-directed talk, and are thus an important source of language for many children (Montag et al., 2015). We translate the child-directed corpus into all of our languages from English using the Google Translate API following Dou and Neubig (2021). We tokenize with spaCy.⁵ Details are given in Table A.1.

Test Data Our test data consists of words in context from two different corpora – Wikipedia (Ginter et al., 2017) and JW300 (Agić and Vulić, 2019) –, plus their gold paradigms from UniMorph. A detailed description of the preparation of the test data can be found in Appendix C.2.

4.3 Models

To use existing state-of-the-art approaches and to evaluate them within the framework of *tUMPC*,

⁵<https://spacy.io>

we tackle the task with a pipeline approach, conducting 4 steps: 1) paradigm clustering, 2) slot alignment, 3) slot prediction, and 4) inflection generation. State-of-the-art models exist for Steps 1 and 4, and we propose systems for Steps 2 and 3 here, together with descriptions of those subtasks. Hyperparameters for all models are in [Appendix B](#).

Paradigm Clustering The first step for *tUMPC* is clustering words into paradigms. We compare 3 paradigm clustering algorithms: [McCurdy et al. \(2021, McC\)](#), [Xu et al. \(2018, Xu\)](#), and the baseline from [Wiemerslage et al. \(2021, SIG\)](#). We modify SIG so it does not predict clusters which are subsets of other clusters, which improves precision. For reference, we provide those systems’ paradigm clustering results in [Table A.2](#). In some clustering systems, each type appears in only one paradigm, which confounds our task for types that can instantiate more than one POS, and thus more than one inflectional paradigm, depending on the context.

Slot Alignment Slot alignment is concerned with identifying which words across paradigms express the same inflectional information.

The system we propose for the task first removes all singleton paradigm clusters from the input, as they contain no inflection pairs to learn from, and converts all remaining clusters into abstract paradigms $c_i \in C$ ([Hulden et al., 2014](#)) by computing the longest common substring (LCS) for each cluster. For example, the LCS of the (true) paradigm of *walk* is *walk*, and the abstract paradigm is *X0, X0+ed, X0+ing, X0+s*. We filter abstract forms that appear less than $\beta = 50$ times.

Next, we assign a POS tag to each cluster. With a set of latent tags Z , we define a Bayesian model:

$$P(k, c_i) = P(k) \prod_{f_j \in c_i} P(f_j | k) \quad (2)$$

$$P(c_i) = \sum_{k \in Z} P(c_i, k) \quad (3)$$

We then maximize the likelihood of the paradigm clusters $c_i \in C$ with an expectation maximization algorithm ([Demester et al., 1977](#)). The POS assignment for each c_i is thus $\arg\max_k (P(k, c_i))$, and $|Z|$ is a hyperparameter which we set to 3.

We now have sets C^k . We assign a slot to each form in an abstract paradigm, considering one C^k at a time. To this end, we compute a fastText ([Bojanowski et al., 2017](#)) embedding for each type in

the corpus and compute the embedding for an abstract form a as the average fastText embedding of all types whose abstract form is a . We define the similarity of two abstract forms a and a' as

$$\text{sim}(a, a') = \cos(a, a') \times (1 - J(a, a')), \quad (4)$$

where $\cos(a, a')$ is the cosine similarity, J is the Jaccard similarity

$$J(a, a') = \frac{|C^a \cap C^{a'}|}{|C^a \cup C^{a'}|}, \quad (5)$$

and C^a is the set of abstract paradigms containing a . Finally, we apply agglomerative clustering over the abstract forms with (4) as our similarity metric and a distance threshold of 0.15.

Slot Prediction Given a test form $f(\ell, \vec{t}_\gamma)$, the goal of slot prediction is to predict the source slot \vec{t}_γ and target slots $\Gamma(\ell)$. We treat this as a simplified POS tagging task and use a character-level Transformer seq2seq model to predict a word’s POS tag and source slot. The model is trained on the results of the slot alignment step. For every word from the raw-text corpus that was assigned a slot, we sample up to 5 unique contexts. A given target word is input with its left and right neighbors; context words that occur fewer than $\alpha = 50$ times in the training data are replaced with OOV. The outputs are the POS tag and the source slot generated by slot alignment. We train our model in FAIRSEQ ([Ott et al., 2019](#)); hyperparameters are in [Appendix B](#).

At test time, the model predicts $f(\ell, \vec{t}_\gamma)$ and the (pseudo) POS tag. Because the slot alignment step associates each POS tag with a unique set of slots, we can perform a simple lookup to find the slots that $f(\ell, \vec{t}_\gamma)$ inflects for.

Morphological Inflection To generate missing forms, we train state-of-the-art inflection models on the results of the slot alignment step and generate surface forms according to the slot prediction. We experiment with the following three models: [Makarov and Clematide \(2018a, M&C\)](#), [Wu et al. \(2021, Wu\)](#), and [Kann and Schütze \(2016b, K&S\)](#).

4.4 Non-neural Baseline

We compare against a rule-based system (baseline) that heuristically predicts the same set of slots for all words, and inflects by applying edit trees to input words. A detailed description is in [Appendix D](#), together with a comparison between baseline and our proposed POS-based system




Figure 1: BMacc for each paradigm clustering system for the POS-based slot aligner; averaged over inflectors.




Figure 2: BMacc for each inflector for the POS-based slot aligner; averaged over paradigm clusters.

for slot alignment and slot prediction. As the POS-based system clearly outperforms `baseline`, we focus the remainder of this paper on the former.

4.5 Results and Discussion

We present results from all experiments in terms of BMacc (Jin et al., 2020). Overall, t UMPC is difficult, though the variance in results over different components of our pipeline implies that there is a great deal of room for the community to innovate. We see the lowest scores for our Greek and Icelandic corpora. These have far fewer tokens than German and Russian, plus higher type–token ratios, which likely makes the task more challenging.

Impact of the Clustering System Figure 1 shows that the choice of paradigm clustering strategy strongly affects our pipeline’s downstream performance. McC, the best performing clustering system on the paradigm clustering task, frequently outperforms the other two strategies. The exception to this is Russian, where Xu gives the best results—by a large margin when learning from the child-directed training corpora.

Impact of the Inflection System From Figure 2 it is obvious that the choice of inflection model does *not* have a large effect on downstream results. All three systems we compare are known to be extremely competitive on the supervised inflection task, so it is reasonable to assume that they fit the generated training data relatively similarly. Future

work can assess how inflection generation can best account for the noisy nature of the data in this task, akin to Michel and Neubig (2018).

Impact of the Corpus The consistency of our results suggests that the child-directed corpus leads to slightly better downstream performance, except in German. Notably, the German Bible contains far more tokens and far fewer types than the corresponding child-directed corpus (Table A.1), which may significantly simplify the learning task.

5 Conclusion

Thanks to strong systems for inflection, segmentation, and paradigm completion, computational morphology is ripe to contribute to the large number of the world’s languages with very few digital resources. We explore this through the novel task t UMPC—which presents several challenges. We believe that truly unsupervised morphology is an important direction, and it can have a large impact on language technology for thousands of languages. With the goal of preserving endangered languages, we note that more than half the world’s languages have no writing system (Harmon, 1995). A frontier for this task would process speech as a strategy for language documentation in unwritten languages.

Acknowledgments

We thank David Yarowsky, members of NALA, and the anonymous reviewers for helpful feedback.

References

- Farrell Ackerman, James P Blevins, and Robert Malouf. 2009. Parts and wholes: Implicative patterns in inflectional paradigms. *Analogy in grammar: Form and acquisition*, pages 54–82.
- Judit Ács. 2018. **BME-HAS system for CoNLL–SIGMORPHON 2018 shared task: Universal morphological reinflection**. In *Proceedings of the CoNLL–SIGMORPHON 2018 Shared Task: Universal Morphological Reinflection*, pages 121–126, Brussels. Association for Computational Linguistics.
- Željko Agić and Ivan Vulić. 2019. **JW300: A wide-coverage parallel corpus for low-resource languages**. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 3204–3210, Florence, Italy. Association for Computational Linguistics.
- 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. Association for Computational Linguistics.
- Antonios Anastasopoulos and Graham Neubig. 2019. **Pushing the limits of low-resource morphological inflection**. In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 984–996, Hong Kong, China. Association for Computational Linguistics.
- Antti Arppe, Jeff Good, Atticus Harrigan, Mans Hulden, Jordan Lachler, Sarah Moeller, Alexis Palmer, Miikka Silfverberg, and Lane Schwartz, editors. 2021. *Proceedings of the 4th Workshop on the Use of Computational Methods in the Study of Endangered Languages Volume 1 (Papers)*. Association for Computational Linguistics, Online.
- Giuseppe Attardi. 2015. Wikiextractor. <https://github.com/attardi/wikiextractor>.
- P. Austin and J. Sallabank, editors. 2011. *The Cambridge Handbook of Endangered Languages*. Cambridge Handbooks in Language and Linguistics. Cambridge University Press.
- R. Harald Baayen, Richard Piepenbrock, and Leon Gulikers. 1996. The CELEX lexical database (CD-ROM).
- Dzmitry Bahdanau, Kyunghyun Cho, and Yoshua Bengio. 2015. Neural machine translation by jointly learning to align and translate. *CoRR*, abs/1409.0473.
- Toms Bergmanis and Sharon Goldwater. 2017. **From segmentation to analyses: a probabilistic model for unsupervised morphology induction**. In *Proceedings of the 15th Conference of the European Chapter of the Association for Computational Linguistics: Volume 1, Long Papers*, pages 337–346, Valencia, Spain. Association for Computational Linguistics.
- Toms Bergmanis and Sharon Goldwater. 2018. **Context sensitive neural lemmatization with Lematus**. In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long Papers)*, pages 1391–1400, New Orleans, Louisiana. Association for Computational Linguistics.
- Toms Bergmanis and Sharon Goldwater. 2019. **Data augmentation for context-sensitive neural lemmatization using inflection tables and raw text**. In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 4119–4128, Minneapolis, Minnesota. Association for Computational Linguistics.
- Toms Bergmanis, Katharina Kann, Hinrich Schütze, and Sharon Goldwater. 2017. **Training data augmentation for low-resource morphological inflection**. In *Proceedings of the CoNLL SIGMORPHON 2017 Shared Task: Universal Morphological Reinflection*, pages 31–39, Vancouver. Association for Computational Linguistics.
- Terra Blevins and Luke Zettlemoyer. 2019. **Better character language modeling through morphology**. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 1606–1613, Florence, Italy. Association for Computational Linguistics.
- Piotr Bojanowski, Edouard Grave, Armand Joulin, and Tomas Mikolov. 2017. Enriching word vectors with subword information. *Transactions of the Association for Computational Linguistics*, 5:135–146.
- 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, Berlin, Germany. Association for Computational Linguistics.
- Nicoletta Calzolari, Frédéric Béchet, Philippe Blache, Khalid Choukri, Christopher Cieri, Thierry Declerck, Sara Goggi, Hitoshi Isahara, Bente Maegaard, Joseph Mariani, et al., editors. 2020. *Proceedings of the 12th Language Resources and Evaluation Conference*. European Language Resources Association, Marseille, France.
- Erwin Chan. 2006. **Learning probabilistic paradigms for morphology in a latent class model**. In *Proceedings of the Eighth Meeting of the ACL Special Interest Group on Computational Phonology and Morphology at HLT-NAACL 2006*, pages 69–78, New

- York City, USA. Association for Computational Linguistics.
- Grzegorz Chrupała. 2008. *Towards a machine-learning architecture for lexical functional grammar parsing*. Ph.D. thesis, Dublin City University.
- Kenneth Ward Church. 2020. [Emerging trends: Subwords, seriously?](#) *Nat. Lang. Eng.*, 26(3):375–382.
- Alexander Clark. 2001. Learning morphology with pair hidden Markov models. In *ACL (Companion Volume)*.
- 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, Copenhagen, Denmark. Association for Computational Linguistics.
- Ryan Cotterell, Christo Kirov, John Sylak-Glassman, Géraldine Walther, Ekaterina Vylomova, Arya D. McCarthy, Katharina Kann, Sabrina J. Mielke, Garrett Nicolai, Miikka Silfverberg, et al. 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. 2017a. [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. 2016a. [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.
- Ryan Cotterell, Christo Kirov, John Sylak-Glassman, David Yarowsky, Jason Eisner, and Mans Hulden. 2016b. [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.
- Ryan Cotterell, Thomas Müller, Alexander Fraser, and Hinrich Schütze. 2015. [Labeled morphological segmentation with semi-Markov models](#). In *Proceedings of the Nineteenth Conference on Computational Natural Language Learning*, pages 164–174, Beijing, China. Association for Computational Linguistics.
- Ryan Cotterell, John Sylak-Glassman, and Christo Kirov. 2017b. [Neural graphical models over strings for principal parts morphological paradigm completion](#). In *Proceedings of the 15th Conference of the European Chapter of the Association for Computational Linguistics: Volume 2, Short Papers*, pages 759–765, Valencia, Spain. Association for Computational Linguistics.
- Ryan Cotterell, Tim Vieira, and Hinrich Schütze. 2016c. [A joint model of orthography and morphological segmentation](#). In *Proceedings of the 2016 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 664–669, San Diego, California. Association for Computational Linguistics.
- Mathias Creutz and Krista Lagus. 2002. [Unsupervised discovery of morphemes](#). In *Proceedings of the ACL-02 Workshop on Morphological and Phonological Learning*, pages 21–30. Association for Computational Linguistics.
- Mathias Creutz and Krista Lagus. 2005. Inducing the morphological lexicon of a natural language from unannotated text. In *Proceedings of the International and Interdisciplinary Conference on Adaptive Knowledge Representation and Reasoning (AKRR’05)*, 106–113, pages 51–59.
- Arthur P. Dempster, Nan M. Laird, and Donald B. Rubin. 1977. Maximum likelihood from incomplete data via the EM algorithm. *Journal of the Royal Statistical Society: Series B (Methodological)*, 39(1):1–22.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. [BERT: Pre-training of deep bidirectional transformers for language understanding](#). In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 4171–4186, Minneapolis, Minnesota. Association for Computational Linguistics.
- Zi-Yi Dou and Graham Neubig. 2021. [Word alignment by fine-tuning embeddings on parallel corpora](#). In *Proceedings of the 16th Conference of the European Chapter of the Association for Computational Linguistics: Main Volume*, pages 2112–2128, Online. Association for Computational Linguistics.
- Greg Durrett and John DeNero. 2013. [Supervised learning of complete morphological paradigms](#). In *Proceedings of the 2013 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 1185–1195, Atlanta, Georgia. Association for Computational Linguistics.

- Christopher Dyer, Smaranda Muresan, and Philip Resnik. 2008. [Generalizing word lattice translation](#). In *Proceedings of ACL-08: HLT*, pages 1012–1020, Columbus, Ohio. Association for Computational Linguistics.
- Micha Elsner, Andrea D Sims, Alexander Erdmann, Antonio Hernandez, Evan Jaffe, Lifeng Jin, Martha Booker Johnson, Shuan Karim, David L King, Luana Lamberti Nunes, et al. 2019. Modeling morphological learning, typology, and change: What can the neural sequence-to-sequence framework contribute? *Journal of Language Modelling*, 7(1):53–98.
- Alexander Erdmann, Micha Elsner, Shijie Wu, Ryan Cotterell, and Nizar Habash. 2020. [The paradigm discovery problem](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 7778–7790, Online. Association for Computational Linguistics.
- Ramy Eskander, Francesca Callejas, Elizabeth Nichols, Judith Klavans, and Smaranda Muresan. 2020. [MorphAGram, evaluation and framework for unsupervised morphological segmentation](#). In *Proceedings of the 12th Language Resources and Evaluation Conference*, pages 7112–7122, Marseille, France. European Language Resources Association.
- Ramy Eskander, Judith Klavans, and Smaranda Muresan. 2019. [Unsupervised morphological segmentation for low-resource polysynthetic languages](#). In *Proceedings of the 16th Workshop on Computational Research in Phonetics, Phonology, and Morphology*, pages 189–195, Florence, Italy. Association for Computational Linguistics.
- Manaal Faruqui, Yulia Tsvetkov, Graham Neubig, and Chris Dyer. 2016. [Morphological inflection generation using character sequence to sequence learning](#). In *Proceedings of the 2016 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 634–643, San Diego, California. Association for Computational Linguistics.
- Raphael Finkel and Gregory Stump. 2007. Principal parts and morphological typology. *Morphology*, 17(1):39–75.
- Filip Ginter, Jan Hajič, Juhani Luotolahti, Milan Straka, and Daniel Zeman. 2017. [CoNLL 2017 shared task - automatically annotated raw texts and word embeddings](#). LINDAT/CLARIAH-CZ digital library at the Institute of Formal and Applied Linguistics (ÚFAL), Faculty of Mathematics and Physics, Charles University.
- Omer Goldman and Reut Tsarfaty. 2021. [Well-defined morphology is sentence-level morphology](#). In *Proceedings of the 1st Workshop on Multilingual Representation Learning*, pages 248–250, Punta Cana, Dominican Republic. Association for Computational Linguistics.
- John Goldsmith. 2001. [Unsupervised learning of the morphology of a natural language](#). *Computational Linguistics*, 27(2):153–198.
- John A Goldsmith. 2010. Segmentation and morphology. *The handbook of computational linguistics and natural language processing*, 57:364.
- 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.
- Stig-Arne Grönroos, Sámi Virpioja, and Mikko Kurimo. 2019. [North Sámi morphological segmentation with low-resource semi-supervised sequence labeling](#). In *Proceedings of the Fifth International Workshop on Computational Linguistics for Uralic Languages*, pages 15–26, Tartu, Estonia. Association for Computational Linguistics.
- Jan Hajič and Dan Zeman, editors. 2017. *Proceedings of the CoNLL 2017 Shared Task: Multilingual Parsing from Raw Text to Universal Dependencies*. Association for Computational Linguistics, Vancouver, Canada.
- David Harmon. 1995. The status of the world’s languages as reported in “Ethnologue”. *Southwest Journal of Linguistics*, 14:1–28.
- Zellig S Harris. 1970. Morpheme boundaries within words: Report on a computer test. In *Papers in Structural and Transformational Linguistics*, pages 68–77. Springer.
- Georg Heigold, Guenter Neumann, and Josef van Genabith. 2016. Neural morphological tagging from characters for morphologically rich languages. *arXiv preprint arXiv:1606.06640*.
- 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*, pages 505–513, Valencia, Spain. Association for Computational Linguistics.
- Sepp Hochreiter and Jürgen Schmidhuber. 1997. Long short-term memory. *Neural computation*, 9(8):1735–1780.
- Valentin Hofmann, Janet Pierrehumbert, and Hinrich Schütze. 2021. [Superbizarre is not superb: Derivational morphology improves BERT’s interpretation of complex words](#). In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pages 3594–3608, Online. Association for Computational Linguistics.

- Matt Hohensee and Emily M. Bender. 2012. [Getting more from morphology in multilingual dependency parsing](#). In *Proceedings of the 2012 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 315–326, Montréal, Canada. Association for Computational Linguistics.
- Mans Hulden, Markus Forsberg, and Malin Ahlberg. 2014. [Semi-supervised learning of morphological paradigms and lexicons](#). In *Proceedings of the 14th Conference of the European Chapter of the Association for Computational Linguistics*, pages 569–578, Gothenburg, Sweden. Association for Computational Linguistics.
- Huiming Jin, Liwei Cai, Yihui Peng, Chen Xia, Arya McCarthy, and Katharina Kann. 2020. [Unsupervised morphological paradigm completion](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 6696–6707, Online. Association for Computational Linguistics.
- Mark Johnson, Thomas L Griffiths, Sharon Goldwater, et al. 2007. Adaptor grammars: A framework for specifying compositional nonparametric bayesian models. *Advances in neural information processing systems*, 19:641.
- Katharina Kann, Ryan Cotterell, and Hinrich Schütze. 2016. [Neural morphological analysis: Encoding-decoding canonical segments](#). In *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing*, pages 961–967, Austin, Texas. Association for Computational Linguistics.
- Katharina Kann, Ryan Cotterell, and Hinrich Schütze. 2017a. [Neural multi-source morphological reinflection](#). In *Proceedings of the 15th Conference of the European Chapter of the Association for Computational Linguistics: Volume 1, Long Papers*, pages 514–524, Valencia, Spain. Association for Computational Linguistics.
- Katharina Kann, Ryan Cotterell, and Hinrich Schütze. 2017b. [One-shot neural cross-lingual transfer for paradigm completion](#). In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1993–2003, Vancouver, Canada. Association for Computational Linguistics.
- Katharina Kann, Stanislas Lauly, and Kyunghyun Cho. 2018a. [The NYU system for the CoNLL–SIGMORPHON 2018 shared task on universal morphological reinflection](#). In *Proceedings of the CoNLL–SIGMORPHON 2018 Shared Task: Universal Morphological Reinflection*, pages 58–63, Brussels. Association for Computational Linguistics.
- Katharina Kann, Jesus Manuel Mager Hois, Ivan Vladimir Meza-Ruiz, and Hinrich Schütze. 2018b. [Fortification of neural morphological segmentation models for polysynthetic minimal-resource languages](#). In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long Papers)*, pages 47–57, New Orleans, Louisiana. Association for Computational Linguistics.
- Katharina Kann, Arya D. McCarthy, Garrett Nicolai, and Mans Hulden. 2020. [The SIGMORPHON 2020 shared task on unsupervised morphological paradigm completion](#). In *Proceedings of the 17th SIGMORPHON Workshop on Computational Research in Phonetics, Phonology, and Morphology*, pages 51–62, Online. Association for Computational Linguistics.
- Katharina Kann and Hinrich Schütze. 2016a. [MED: The LMU system for the SIGMORPHON 2016 shared task on morphological reinflection](#). In *Proceedings of the 14th SIGMORPHON Workshop on Computational Research in Phonetics, Phonology, and Morphology*, pages 62–70, Berlin, Germany. Association for Computational Linguistics.
- Katharina Kann and Hinrich Schütze. 2016b. [Single-model encoder-decoder with explicit morphological representation for reinflection](#). In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, pages 555–560, Berlin, Germany. Association for Computational Linguistics.
- Katharina Kann and Hinrich Schütze. 2017. [Unlabeled data for morphological generation with character-based sequence-to-sequence models](#). In *Proceedings of the First Workshop on Subword and Character Level Models in NLP*, pages 76–81, Copenhagen, Denmark. Association for Computational Linguistics.
- Katharina Kann and Hinrich Schütze. 2018. [Neural transductive learning and beyond: Morphological generation in the minimal-resource setting](#). In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 3254–3264, Brussels, Belgium. Association for Computational Linguistics.
- Yova Kementchedjheva, Johannes Bjerva, and Isabelle Augenstein. 2018. [Copenhagen at CoNLL–SIGMORPHON 2018: Multilingual inflection in context with explicit morphosyntactic decoding](#). In *Proceedings of the CoNLL–SIGMORPHON 2018 Shared Task: Universal Morphological Reinflection*, pages 93–98, Brussels. Association for Computational Linguistics.
- Salam Khalifa, Nizar Habash, Fadhl Eryani, Ossama Obeid, Dana Abdulrahim, and Meera Al Kaabi. 2018. [A morphologically annotated corpus of Emirati Arabic](#). In *Proceedings of the Eleventh International Conference on Language Resources and Evaluation (LREC 2018)*, Miyazaki, Japan. European Language Resources Association (ELRA).

- Diederik P Kingma and Jimmy Ba. 2014. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*.
- Christo Kirov, Ryan Cotterell, John Sylak-Glassman, Géraldine Walther, Ekaterina Vylomova, Patrick Xia, Manaal Faruqui, Sabrina J. Mielke, Arya 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)*, Miyazaki, Japan. European Language Resources Association (ELRA).
- Christo Kirov, John Sylak-Glassman, Roger Que, and David Yarowsky. 2016. **Very-large scale parsing and normalization of Wiktionary morphological paradigms**. In *Proceedings of the Tenth International Conference on Language Resources and Evaluation (LREC'16)*, pages 3121–3126, Portorož, Slovenia. European Language Resources Association (ELRA).
- Judith L. Klavans, editor. 2018. *Proceedings of the Workshop on Computational Modeling of Polysynthetic Languages*. Association for Computational Linguistics, Santa Fe, New Mexico, USA.
- Elena Klyachko, Alexey Sorokin, Natalia Krizhanovskaya, Andrew Krizhanovsky, and Galina Ryazanskaya. 2020. **LowResourceEval-2019: a shared task on morphological analysis for low-resource languages**. *arXiv preprint arXiv:2001.11285*.
- Dan Kondratyuk. 2019. **Cross-lingual lemmatization and morphology tagging with two-stage multilingual BERT fine-tuning**. In *Proceedings of the 16th Workshop on Computational Research in Phonetics, Phonology, and Morphology*, pages 12–18, Florence, Italy. Association for Computational Linguistics.
- Lingpeng Kong, Chris Dyer, and Noah A Smith. 2015. Segmental recurrent neural networks. *arXiv preprint arXiv:1511.06018*.
- Mikko Kurimo, Sami Virpioja, Ville Turunen, and Krista Lagus. 2010. **Morpho challenge 2005-2010: Evaluations and results**. In *Proceedings of the 11th Meeting of the ACL Special Interest Group on Computational Morphology and Phonology*, pages 87–95, Uppsala, Sweden. Association for Computational Linguistics.
- Ilmari Kylliäinen and Miikka Silfverberg. 2019. **Ensembles of neural morphological inflection models**. In *Proceedings of the 22nd Nordic Conference on Computational Linguistics*, pages 304–309, Turku, Finland. Linköping University Electronic Press.
- John Lafferty, Andrew McCallum, and Fernando CN Pereira. 2001. Conditional random fields: Probabilistic models for segmenting and labeling sequence data. In *ICML*.
- Constantine Lignos, Erwin Chan, Mitchell P Marcus, and Charles Yang. 2009. A rule-based unsupervised morphology learning framework. In *CLEF (Working Notes)*.
- Ling Liu, Ilamvazhuthy Subbiah, Adam Wiemerslage, Jonathan Lilley, and Sarah Moeller. 2018. **Morphological reinflection in context: CU boulder’s submission to CoNLL–SIGMORPHON 2018 shared task**. In *Proceedings of the CoNLL–SIGMORPHON 2018 Shared Task: Universal Morphological Reinflection*, pages 86–92, Brussels. Association for Computational Linguistics.
- Brian MacWhinney. 2014. *The CHILDES project: Tools for analyzing talk*. Psychology Press.
- Manuel Mager, Özlem Çetinoglu, and Katharina Kann. 2020. **Tackling the low-resource challenge for canonical segmentation**. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 5237–5250, Online. Association for Computational Linguistics.
- Manuel Mager and Katharina Kann. 2020. **The IMS–CUBoulder system for the SIGMORPHON 2020 shared task on unsupervised morphological paradigm completion**. In *Proceedings of the 17th SIGMORPHON Workshop on Computational Research in Phonetics, Phonology, and Morphology*, pages 99–105, Online. Association for Computational Linguistics.
- Manuel Mager, Arturo Oncevay, Abteen Ebrahimi, John Ortega, Annette Rios, Angela Fan, Ximena Gutierrez-Vasques, Luis Chiruzzo, Gustavo Giménez-Lugo, Ricardo Ramos, et al. 2021. **Findings of the AmericasNLP 2021 shared task on open machine translation for indigenous languages of the Americas**. In *Proceedings of the First Workshop on Natural Language Processing for Indigenous Languages of the Americas*, pages 202–217, Online. Association for Computational Linguistics.
- Peter Makarov and Simon Clematide. 2018a. **Imitation learning for neural morphological string transduction**. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 2877–2882, Brussels, Belgium. Association for Computational Linguistics.
- Peter Makarov and Simon Clematide. 2018b. **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.
- Peter Makarov and Simon Clematide. 2018c. **UZH at CoNLL–SIGMORPHON 2018 shared task on universal morphological reinflection**. In *Proceedings of the CoNLL–SIGMORPHON 2018 Shared Task: Universal Morphological Reinflection*, pages 69–75, Brussels. Association for Computational Linguistics.

- Peter Makarov, Tatiana Ruzsics, and Simon Clematide. 2017. [Align and copy: UZH at SIGMORPHON 2017 shared task for morphological reinflection](#). In *Proceedings of the CoNLL SIGMORPHON 2017 Shared Task: Universal Morphological Reinflection*, pages 49–57, Vancouver. Association for Computational Linguistics.
- 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, Melbourne, Australia. Association for Computational Linguistics.
- Chaitanya Malaviya, Shijie Wu, and Ryan Cotterell. 2019. [A simple joint model for improved contextual neural lemmatization](#). In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 1517–1528, Minneapolis, Minnesota. Association for Computational Linguistics.
- Andreas Maletti and Matthieu Constant. 2011. Proceedings of the 9th international workshop on finite state methods and natural language processing. In *Proceedings of the 9th International Workshop on Finite State Methods and Natural Language Processing*.
- Robert Malouf, Farrell Ackerman, and Arturs Semenuks. 2020. [Lexical databases for computational analyses: A linguistic perspective](#). In *Proceedings of the Society for Computation in Linguistics 2020*, pages 446–456, New York, New York. 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, et al. 2020a. [UniMorph 3.0: Universal Morphology](#). In *Proceedings of the 12th Language Resources and Evaluation Conference*, pages 3922–3931, Marseille, France. European Language Resources Association.
- Arya D. McCarthy, Miikka Silfverberg, Ryan Cotterell, Mans Hulden, and David Yarowsky. 2018. [Marrying Universal Dependencies and Universal Morphology](#). In *Proceedings of the Second Workshop on Universal Dependencies (UDW 2018)*, pages 91–101, Brussels, Belgium. Association for Computational Linguistics.
- Arya D. McCarthy, Ekaterina Vylomova, Shijie Wu, Chaitanya Malaviya, Lawrence Wolf-Sonkin, Garrett Nicolai, Christo Kirov, Miikka Silfverberg, Sabrina J. Mielke, Jeffrey Heinz, et al. 2019. [The SIGMORPHON 2019 shared task: Morphological analysis in context and cross-lingual transfer for inflection](#). In *Proceedings of the 16th Workshop on Computational Research in Phonetics, Phonology, and Morphology*, pages 229–244, Florence, Italy. Association for Computational Linguistics.
- Arya D. McCarthy, Rachel Wicks, Dylan Lewis, Aaron Mueller, Winston Wu, Oliver Adams, Garrett Nicolai, Matt Post, and David Yarowsky. 2020b. [The Johns Hopkins University Bible corpus: 1600+ tongues for typological exploration](#). In *Proceedings of the 12th Language Resources and Evaluation Conference*, pages 2884–2892, Marseille, France. European Language Resources Association.
- Kate McCurdy, Sharon Goldwater, and Adam Lopez. 2021. [Adaptor Grammars for unsupervised paradigm clustering](#). In *Proceedings of the 18th SIGMORPHON Workshop on Computational Research in Phonetics, Phonology, and Morphology*, pages 82–89, Online. Association for Computational Linguistics.
- Eleni Metheniti and Guenter Neumann. 2020. [Wikinflection corpus: A \(better\) multilingual, morpheme-annotated inflectional corpus](#). In *Proceedings of the 12th Language Resources and Evaluation Conference*, pages 3905–3912, Marseille, France. European Language Resources Association.
- Paul Michel and Graham Neubig. 2018. [MTNT: A testbed for machine translation of noisy text](#). In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 543–553, Brussels, Belgium. Association for Computational Linguistics.
- Jeffrey Micher. 2017. [Improving coverage of an Inuktitut morphological analyzer using a segmental recurrent neural network](#). In *Proceedings of the 2nd Workshop on the Use of Computational Methods in the Study of Endangered Languages*, pages 101–106, Honolulu. Association for Computational Linguistics.
- Sarah Moeller, Ling Liu, Changbing Yang, Katharina Kann, and Mans Hulden. 2020. [IGT2P: From interlinear glossed texts to paradigms](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 5251–5262, Online. Association for Computational Linguistics.
- Tumi Moeng, Sheldon Reay, Aaron Daniels, and Jan Buys. 2021. [Canonical and surface morphological segmentation for Nguni languages](#). *CoRR*, abs/2104.00767.
- Christian Monson, Jaime Carbonell, Alon Lavie, and Lori Levin. 2007. [ParaMor: Finding paradigms across morphology](#). In *Workshop of the Cross-Language Evaluation Forum for European Languages*, pages 900–907. Springer.
- Jessica L. Montag, Michael N. Jones, and Linda B. Smith. 2015. [The words children hear: Picture books and the statistics for language learning](#). *Psychological Science*, 26(9):1489–1496. PMID: 26243292.

- Thomas Mueller, 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, Seattle, Washington, USA. Association for Computational Linguistics.
- Karthik Narasimhan, Regina Barzilay, and Tommi Jaakkola. 2015. [An unsupervised method for uncovering morphological chains](#). *Transactions of the Association for Computational Linguistics*, 3:157–167.
- Minh Van Nguyen, Viet Dac Lai, Amir Pouran Ben Veyseh, and Thien Huu Nguyen. 2021. [Trankit: A light-weight transformer-based toolkit for multilingual natural language processing](#). In *Proceedings of the 16th Conference of the European Chapter of the Association for Computational Linguistics: System Demonstrations*, pages 80–90, Online. Association for Computational Linguistics.
- Garrett Nicolai, Colin Cherry, and Grzegorz Kondrak. 2015. [Inflection generation as discriminative string transduction](#). In *Proceedings of the 2015 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 922–931, Denver, Colorado. Association for Computational Linguistics.
- Garrett Nicolai, Kyle Gorman, and Ryan Cotterell, editors. 2021. *Proceedings of the 18th SIGMORPHON Workshop on Computational Research in Phonetics, Phonology, and Morphology*. Association for Computational Linguistics, Online.
- Garrett Nicolai, Dylan Lewis, Arya D. McCarthy, Aaron Mueller, Winston Wu, and David Yarowsky. 2020. [Fine-grained morphosyntactic analysis and generation tools for more than one thousand languages](#). In *Proceedings of the 12th Language Resources and Evaluation Conference*, pages 3963–3972, Marseille, France. European Language Resources Association.
- Myle Ott, Sergey Edunov, Alexei Baevski, Angela Fan, Sam Gross, Nathan Ng, David Grangier, and Michael Auli. 2019. [fairseq: A fast, extensible toolkit for sequence modeling](#). In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics (Demonstrations)*, pages 48–53, Minneapolis, Minnesota. Association for Computational Linguistics.
- Hyunji Hayley Park, Katherine J. Zhang, Coleman Hayley, Kenneth Steimel, Han Liu, and Lane Schwartz. 2021. [Morphology Matters: A Multilingual Language Modeling Analysis](#). *Transactions of the Association for Computational Linguistics*, 9:261–276.
- Tiago Pimentel, Arya D. McCarthy, Damian Blasi, Brian Roark, and Ryan Cotterell. 2019. [Meaning to form: Measuring systematicity as information](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 1751–1764, Florence, Italy. Association for Computational Linguistics.
- Tiago Pimentel, Maria Ryskina, Sabrina J. Mielke, Shijie Wu, Eleanor Chodroff, Brian Leonard, Garrett Nicolai, Yustinus Ghanggo Ate, Salam Khalifa, Nizar Habash, et al. 2021. [Sigmorphon 2021 shared task on morphological reinflection: Generalization across languages](#). In *Proceedings of the 18th SIGMORPHON Workshop on Computational Research in Phonetics, Phonology, and Morphology*, pages 229–259, Online. Association for Computational Linguistics.
- Hoifung Poon, Colin Cherry, and Kristina Toutanova. 2009. [Unsupervised morphological segmentation with log-linear models](#). In *Proceedings of Human Language Technologies: The 2009 Annual Conference of the North American Chapter of the Association for Computational Linguistics*, pages 209–217, Boulder, Colorado. Association for Computational Linguistics.
- Peng Qi, Yuhao Zhang, Yuhui Zhang, Jason Bolton, and Christopher D. Manning. 2020. [Stanza: A python natural language processing toolkit for many human languages](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics: System Demonstrations*, pages 101–108, Online. Association for Computational Linguistics.
- Mohammad Sadegh Rasooli and Michael Collins. 2017. [Cross-lingual syntactic transfer with limited resources](#). *Transactions of the Association for Computational Linguistics*, 5:279–293.
- Teemu Ruokolainen, Oskar Kohonen, Sami Virpioja, and Mikko Kurimo. 2013. [Supervised morphological segmentation in a low-resource learning setting using conditional random fields](#). In *Proceedings of the Seventeenth Conference on Computational Natural Language Learning*, pages 29–37, Sofia, Bulgaria. Association for Computational Linguistics.
- Teemu Ruokolainen, Oskar Kohonen, Sami Virpioja, and Mikko Kurimo. 2014. [Painless semi-supervised morphological segmentation using conditional random fields](#). In *Proceedings of the 14th Conference of the European Chapter of the Association for Computational Linguistics, volume 2: Short Papers*, pages 84–89, Gothenburg, Sweden. Association for Computational Linguistics.
- Tatyana Ruzsics and Tanja Samardžić. 2017. [Neural sequence-to-sequence learning of internal word structure](#). In *Proceedings of the 21st Conference on Computational Natural Language Learning (CoNLL 2017)*, pages 184–194, Vancouver, Canada. Association for Computational Linguistics.
- Patrick Schone and Daniel Jurafsky. 2001. [Knowledge-free induction of inflectional morphologies](#). In *Second Meeting of the North American Chapter of the Association for Computational Linguistics*.
- Wolfgang Seeker and Özlem Çetinoğlu. 2015. [A graph-based lattice dependency parser for joint morpholog-](#)

- ical segmentation and syntactic analysis. *Transactions of the Association for Computational Linguistics*, 3:359–373.
- Amit Seker and Reut Tsarfaty. 2020. A pointer network architecture for joint morphological segmentation and tagging. In *Findings of the Association for Computational Linguistics: EMNLP 2020*, pages 4368–4378, Online. Association for Computational Linguistics.
- Abhishek Sharma, Ganesh Katrapati, and Dipti Misra Sharma. 2018. IIT(BHU)–IIITH at CoNLL–SIGMORPHON 2018 shared task on universal morphological reinflection. In *Proceedings of the CoNLL–SIGMORPHON 2018 Shared Task: Universal Morphological Reinflection*, pages 105–111, Brussels. Association for Computational Linguistics.
- 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.
- Miikka Silfverberg, Adam Wiemerslage, Ling Liu, and Lingshuang Jack Mao. 2017. Data augmentation for morphological reinflection. In *Proceedings of the CoNLL SIGMORPHON 2017 Shared Task: Universal Morphological Reinflection*, pages 90–99, Vancouver. Association for Computational Linguistics.
- Assaf Singer and Katharina Kann. 2020. The NYU-CUBoulder systems for SIGMORPHON 2020 task 0 and task 2. In *Proceedings of the 17th SIGMORPHON Workshop on Computational Research in Phonetics, Phonology, and Morphology*, pages 90–98, Online. Association for Computational Linguistics.
- Peter Smit, Sami Virpioja, Stig-Arne Grönroos, Mikko Kurimo, et al. 2014. Morfessor 2.0: Toolkit for statistical morphological segmentation. In *The 14th Conference of the European Chapter of the Association for Computational Linguistics (EACL), Gothenburg, Sweden, April 26-30, 2014*. Aalto University.
- Noah A. Smith and Jason Eisner. 2005. Contrastive estimation: Training log-linear models on unlabeled data. In *Proceedings of the 43rd Annual Meeting of the Association for Computational Linguistics (ACL’05)*, pages 354–362, Ann Arbor, Michigan. Association for Computational Linguistics.
- Alexey Sorokin. 2019. Convolutional neural networks for low-resource morpheme segmentation: baseline or state-of-the-art? In *Proceedings of the 16th Workshop on Computational Research in Phonetics, Phonology, and Morphology*, pages 154–159, Florence, Italy. Association for Computational Linguistics.
- Milan Straka, Jana Straková, and Jan Hajic. 2019. UD-Pipe at SIGMORPHON 2019: Contextualized embeddings, regularization with morphological categories, corpora merging. In *Proceedings of the 16th Workshop on Computational Research in Phonetics, Phonology, and Morphology*, pages 95–103, Florence, Italy. Association for Computational Linguistics.
- John Sylak-Glassman, Christo Kirov, Matt Post, Roger Que, and David Yarowsky. 2015a. A universal feature schema for rich morphological annotation and fine-grained cross-lingual part-of-speech tagging. In *Systems and Frameworks for Computational Morphology*, pages 72–93, Cham. Springer International Publishing.
- John Sylak-Glassman, Christo Kirov, David Yarowsky, and Roger Que. 2015b. 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. Association for Computational Linguistics.
- Gábor Szolnok, Botond Barta, Dorina Lakatos, and Judit Ács. 2021. Bme submission for sigmorphon 2021 shared task 0. a three step training approach with data augmentation for morphological inflection. In *Proceedings of the 18th SIGMORPHON Workshop on Computational Research in Phonetics, Phonology, and Morphology*, pages 268–273, Online. Association for Computational Linguistics.
- Aleš Tamchyna, Marion Weller-Di Marco, and Alexander Fraser. 2017. Modeling target-side inflection in neural machine translation. In *Proceedings of the Second Conference on Machine Translation*, pages 32–42, Copenhagen, Denmark. Association for Computational Linguistics.
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. 2017. Attention is all you need. *Advances in neural information processing systems*, 30.
- Oriol Vinyals, Meire Fortunato, and Navdeep Jaitly. 2015. Pointer networks. *Advances in neural information processing systems*, 28.
- Sami Virpioja, Oskar Kohonen, and Krista Lagus. 2009. Unsupervised morpheme discovery with Allomorpher. In *CLEF (Working Notes)*.
- Linlin Wang, Zhu Cao, Yu Xia, and Gerard De Melo. 2016. Morphological segmentation with window LSTM neural networks. In *Thirtieth AAAI Conference on Artificial Intelligence*.
- Weihua Wang, Rashel Fam, Feilong Bao, Yves Lepage, and Guanglai Gao. 2019. Neural morphological segmentation model for Mongolian. In *2019*

- International Joint Conference on Neural Networks (IJCNN)*, pages 1–7. IEEE.
- Adam Wiemerslage, Arya D. McCarthy, Alexander Erdmann, Garrett Nicolai, Manex Agirrezabal, Miikka Silfverberg, Mans Hulden, and Katharina Kann. 2021. [Findings of the SIGMORPHON 2021 shared task on unsupervised morphological paradigm clustering](#). In *Proceedings of the 18th SIGMORPHON Workshop on Computational Research in Phonetics, Phonology, and Morphology*, pages 72–81, Online. Association for Computational Linguistics.
- Shijie Wu, Ryan Cotterell, and Mans Hulden. 2021. [Applying the transformer to character-level transduction](#). In *Proceedings of the 16th Conference of the European Chapter of the Association for Computational Linguistics: Main Volume*, pages 1901–1907, Online. Association for Computational Linguistics.
- Shijie Wu, Ryan Cotterell, and Timothy O’Donnell. 2019. [Morphological irregularity correlates with frequency](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 5117–5126, Florence, Italy. Association for Computational Linguistics.
- Hongzhi Xu, Jordan Kodner, Mitchell Marcus, and Charles Yang. 2020. [Modeling morphological typology for unsupervised learning of language morphology](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 6672–6681, Online. Association for Computational Linguistics.
- Ruochen Xu, Yiming Yang, Naoki Otani, and Yuexin Wu. 2018. [Unsupervised cross-lingual transfer of word embedding spaces](#). In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 2465–2474, Brussels, Belgium. Association for Computational Linguistics.
- Changbing Yang, Garrett Nicolai, and Miikka Silfverberg. 2021. [Unsupervised paradigm clustering using transformation rules](#). In *Proceedings of the 18th SIGMORPHON Workshop on Computational Research in Phonetics, Phonology, and Morphology*, pages 98–106, Online. Association for Computational Linguistics.
- David Yarowsky, Grace Ngai, and Richard Wicentowski. 2001. [Inducing multilingual text analysis tools via robust projection across aligned corpora](#). In *Proceedings of the First International Conference on Human Language Technology Research*.
- David Yarowsky and Richard Wicentowski. 2000. [Minimally supervised morphological analysis by multimodal alignment](#). In *Proceedings of the 38th Annual Meeting of the Association for Computational Linguistics*, pages 207–216, Hong Kong. Association for Computational Linguistics.
- Nasser Zalmout and Nizar Habash. 2020. [Utilizing subword entities in character-level sequence-to-sequence lemmatization models](#). In *Proceedings of the 28th International Conference on Computational Linguistics*, pages 4676–4682, Barcelona, Spain (Online). International Committee on Computational Linguistics.
- Daniel Zeman, Joakim Nivre, Mitchell Abrams, Elia Ackermann, Noëmi Aepli, Hamid Aghaei, Željko Agić, Amir Ahmadi, Lars Ahrenberg, Chika Kennedy Ajede, et al. 2021. [Universal dependencies 2.8](#). LINDAT/CLARIAH-CZ digital library at the Institute of Formal and Applied Linguistics (ÚFAL), Faculty of Mathematics and Physics, Charles University.
- Anna Zueva, Anastasia Kuznetsova, and Francis Tyers. 2020. [A finite-state morphological analyser for Evenki](#). In *Proceedings of the 12th Language Resources and Evaluation Conference*, pages 2581–2589, Marseille, France. European Language Resources Association.

A Remaining Results from Main Text

The statistics of the data used in our experiments is given in Table A.1. Paradigm clustering BMF1 is given in Table A.2. Additionally, BMAcc on the two test corpora is given in Figure A.1.

B Hyperparameters

B.1 Morphological Inflection

Training We train all inflection models on the (word, source slot, target slot) triples produced by the slot alignment. Each inflection system considers the word as an input form, and the slots as the tags. We take the hyperparameters from (Makarov and Clematide, 2018a), and (Wu et al., 2021) exactly for each language. For the LSTM, we train a single layer bidirectional encoder with embedding size 100, and LSTM hidden size of 100. The decoder is also a single layer LSTM with hidden size 100. We employ a soft-attention mechanism (Bahdanau et al., 2015), and optimize with Adam (Kingma and Ba, 2014) with a learning rate of 0.001, and a gradient clip of 1.0. We train for up to 30 epochs, and a batch size of 16. We employ a soft attention mechanism (Bahdanau et al., 2015).

B.2 Slot Prediction

The slot prediction model is a character-level Transformer encoder-decoder, where both the encoder and decoder have 3 layers and 4 attention heads. We optimize with Adam with a learning rate of 0.0001, and a clip norm of 0.2 for up to 5 epochs.

C Additional Details Regarding our Datasets

C.1 Statistics of Our Raw-text Corpora

We give dataset statistics in Table A.1, including type-token ratios. Bible sizes vary depending on whether or not the Old Testament is included. In the case of smaller Bibles, we down-sample the child-directed corpus to have a roughly equal number of tokens.

C.2 Test Set Creation

We use lemmas and POS tag annotations to match words from the test corpora with UniMorph entries. We sample sentences from the annotated Wikipedia corpora (Ginter et al., 2017) from the ConLL 2017 shared task on Multilingual Parsing (Hajič and Zeman, 2017). For Icelandic, which is not included in this dataset, we use wikiextractor (Attardi, 2015)

to get the raw Wikipedia text, and acquire lemma and POS annotations with Stanza (Qi et al., 2020). We hypothesize that systems trained on the Bible corpus may not generalize well to the modern language in Wikipedia. We thus additionally sample test sentences from the JW300 corpus, which is more likely to include religious language that resembles that of the bible. For JW300 we rely on the tokenization provided by the authors, but we again use Stanza for lemma and POS annotations.

For a given language and test corpus, we group gold paradigms by POS, and whether at least one form from the paradigm is attested in both training corpora. This means we have two categories for each POS: *seen*, wherein at least one form is attested in both training corpora, and *unseen*, wherein no forms are attested in either training corpus. We sample up to 200 paradigms from each category, ensuring that each category contributes an equal number of paradigms to the gold set. Then one surface form for each gold paradigm is sampled at random, in context, from the test corpus to serve as input to the systems at test time.

D Non-Neural Baseline for tUMPC

Given the set of word form clusters c_1, \dots, c_k , where each cluster $c_i = \{f_1, \dots, f_n\}$ is a collection of forms f_j . We start by extracting all edit trees $t = \text{EditTree}(f, f')$ (Chrupała, 2008), where f and f' belong to the same cluster. Let $\text{Count}(t)$ be the count of tree t across the entire training set. Further, let $\text{MLen}(t)$ be the total number of characters which have to match in the input string, when we apply edit tree t . For example, for an edit tree t which maps **walking** to **walks**, a suffix **ing** must match, so $\text{MLen}(t) = 3$. Finally, let $\text{MStr}(t) = u$ be the string consisting of all insertions performed by the edit tree. For the given example t , $\text{MStr}(t) = \text{s}$

When generating outputs for a given form f , we first form the set of all edit trees which can be applied to f . We then order them in the following way: $t > t'$ if $\text{MLen}(t) > \text{MLen}(t')$, or if the precondition lengths are equal, $\text{Count}(t) > \text{Count}(t')$. We then apply the top- N trees to f to generate all remaining forms in the inflectional paradigm of f . We set N to the 95th percentile of paradigm sizes in our input cluster data, not counting singleton paradigms. Each slot labeled is assigned based on t as $\text{MStr}(t)$. Note that this will typically not generate a slot label for the input form

Corpus	Language	Lines	Tokens	Types	Type-Token Ratio
Bible	German	31102	813317	20644	0.025
	Greek	7914	194135	15541	0.080
	Icelandic	7860	185995	13050	0.070
	Russian	31102	714828	43542	0.061
Child Directed	German	26592	633229	31384	0.050
	Greek	8513	196344	18424	0.090
	Icelandic	8380	181687	17767	0.101
	Russian	26592	586274	44823	0.077

Table A.1: Statistics for raw text corpora used for morphology learning

System	Bible					Child-Directed				
	DEU	ELL	ISL	RUS	Average	DEU	ELL	ISL	RUS	Average
McC	79.19	81.91	81.66	82.01	81.19	87.72	73.68	84.65	86.28	83.08
Xu	63.90	65.14	67.81	52.80	63.91	70.02	46.14	55.22	63.48	58.72
STG	46.04	57.22	47.24	45.10	48.90	45.69	47.04	43.08	47.80	45.90

Table A.2: Paradigm clustering BMF1 scores for a sample of clusters attested in UniMorph.

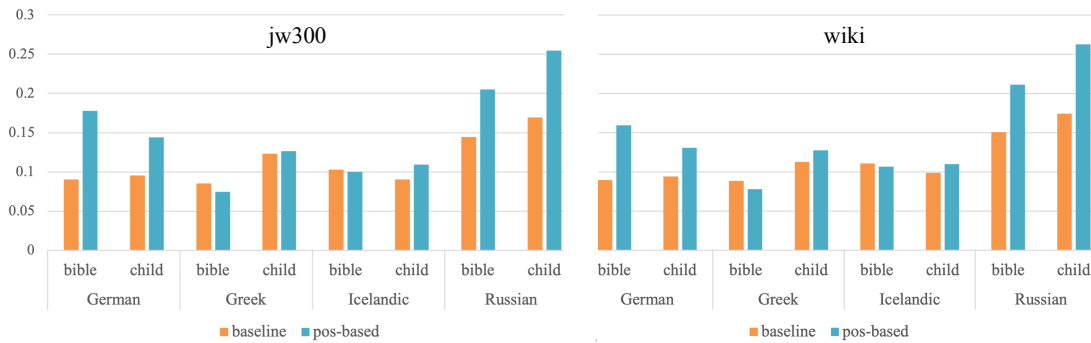


Figure A.1: BMF1 for both slot alignment systems on each test corpus, averaged over results for all input clusters. The POS-based system is also averaged over each inflection system.

f . We, therefore, find the maximal edit tree t (in the sense that it has maximal precondition length and count) which translates one of the generated forms f' back into the original input form f . The slot label for form f is then $MStr(t)$.

A comparison between baseline and our proposed POS-based system is shown in Figure A.1. The latter outperforms baseline in the majority of settings, often by a large margin.