

The Paradigm Discovery Problem

Alexander Erdmann^{ⓐ,ⓓ} Micha Elsner[ⓐ] Shijie Wu[ⓓ]

Ryan Cotterell^{ⓕ,ⓓ} Nizar Habash[ⓓ]

[ⓐ]Ohio State University [ⓓ]New York University Abu Dhabi

[ⓕ]Johns Hopkins University [ⓓ]University of Cambridge [ⓕ]ETH Zürich

{ae1541,nizar.habash}@nyu.edu elsner.14@osu.edu

shijie.wu@jhu.edu ryan.cotterell@inf.ethz.ch

Abstract

This work treats the paradigm discovery problem (PDP)—the task of learning an inflectional morphological system from unannotated sentences. We formalize the PDP and develop evaluation metrics for judging systems. Using currently available resources, we construct datasets for the task. We also devise a heuristic benchmark for the PDP and report empirical results on five diverse languages. Our benchmark system first makes use of word embeddings and string similarity to cluster forms by cell and by paradigm. Then, we bootstrap a neural transducer on top of the clustered data to predict words to realize the empty paradigm slots. An error analysis of our system suggests clustering by cell across different inflection classes is the most pressing challenge for future work. Our code and data are available at <https://github.com/alexerdmann/ParadigmDiscovery>.

1 Introduction

In childhood, we induce our native language’s morphological system from unannotated input. For instance, we learn that *ring* and *rang* belong to the same inflectional paradigm. We also learn that *rings* and *bangs* belong to the same cell, i.e., they realize the same morphosyntactic properties 3.SG.PRES, but in different paradigms. Acquiring such paradigmatic knowledge enables us to produce unseen inflectional variants of new vocabulary items, i.e. to complete morphological paradigms. Much work has addressed this task, which Ackerman et al. (2009) call the paradigm cell filling problem (PCFP),¹ but few have discussed inducing paradigmatic knowledge from scratch, which we call the **paradigm discovery problem (PDP)**.²

¹In the NLP literature, this task is called morphological reinflection or morphological inflection generation (Cotterell et al., 2016a); this is only a difference in nomenclature.

²Elsner et al. (2019) call the task the paradigm cell discovery problem; we drop *cell* to distinguish our task from

As an unsupervised task, the PDP poses challenges for modeling and evaluation and has yet to be attempted in its full form (Elsner et al., 2019). However, we contend there is much to be gained from formalizing and studying the PDP. There are insights for cognitive modeling to be won (Pinker, 2001; Goldwater, 2007) and intuitions on combating sparse data for language generation (King and White, 2018) to be accrued. Unsupervised language processing also has natural applications in the documentation of endangered languages (Zamaraeva et al., 2019) where a lot of annotated data is never likely to exist. Our formalization of the PDP offers a starting point for future work on unsupervised morphological paradigm completion.

Our paper presents a concrete formalization of the PDP. Then, as a baseline for future work, we introduce a heuristic benchmark system. Our benchmark system takes an unannotated text corpus and a lexicon of words from the corpus to be analyzed. It first clusters the lexicon by cell and then by paradigm making use of distributional semantics and string similarity. Finally, it uses this clustering as silver-standard supervision to bootstrap a neural transducer (Vaswani et al., 2017) that generates the desired target inflections. That is, the model posits forms to realize unoccupied cell slots in each proposed paradigm. Even though our benchmark system models only one part of speech (POS) at a time, our framework extends to the full PDP to support future, more intricate systems. We propose two separate metrics to evaluate both the clustering of attested forms into paradigms and cells and the prediction of unseen inflected forms. Our metrics handle non-canonical morphological behavior discussed in theoretical literature (Corbett, 2005) and extend to the full PDP.

For three of the five languages we consider, our benchmark system predicts unattested inflections

one of its subtasks which Boyé and Schalchli (2019) call the paradigm cell finding problem (see §2.2).

of lexicon forms with accuracy within 20% of a fully supervised system. However, our analysis suggests clustering forms into cells consistently across paradigms is still a very pressing challenge.

2 Previous Work in Morphology

This section couches our work on the PDP in terms of previous trends in morphological modeling.

2.1 Unsupervised Morphology

Much work on unsupervised morphological modeling focuses on segmentation (Gaussier, 1999; Goldsmith, 2001; Creutz and Lagus, 2005; Narasimhan et al., 2015; Bergmanis and Goldwater, 2017; Xu et al., 2018). While morphological segmenters can distinguish real from spurious affixes (e.g., *bring* \neq *br* + *ing*) with high accuracy, they do not attempt to solve the PDP. They do, however, reveal which forms take the same affixes (e.g., *walked*, *talked*), not which forms occupy the same cell (e.g., *walked*, *brought*). Indeed, they explicitly struggle with irregular morphology. Segmenters also cannot easily model non-concatenative phenomena like ablaut, vowel harmony and templatic processes.

Two works have proposed tasks which can be considered alternative formulations of the PDP, using either minimal or indirect supervision to bootstrap their models. We discuss each in turn. First, Dreyer and Eisner (2011) use a generative model to cluster forms into paradigms and cells with a Bayesian non-parametric mixture of weighted finite-state transducers. They present a PDP framework which, in principle, could be fully unsupervised, but their model requires a small seed of labeled data to get key information like the number of cells distinguished, making it less relevant cognitively. In contrast, our task is not directly supervised and focuses on distributional context. Second, contemporaneous to our work, Jin et al. (2020) propose a similar framework for SIGMORPHON 2020’s shared task on unsupervised morphological paradigm completion. Given only a small corpus and lexicon of verbal lemmata, participating systems must propose full paradigms for each lemma. By contrast, our framework does not reveal how many paradigms should be generated, nor do we privilege a specific form as the lemma, but we do use a larger lexicon of exclusively verbal or nominal forms. Their proposed baseline uses distributional context for POS tagging and features, but does not train embeddings as the corpus is small.

2.2 Subtasks of Paradigm Discovery

A few works address subtasks of the PDP. Erdmann and Habash (2018) learn paradigm membership from raw text, but do not sort paradigms into cells. Boyé and Schalchli (2019) discuss the paradigm cell finding problem, identifying the cell (but not paradigm) realized by a given form. Lee (2015) clusters forms into cells across inflection classes. Beniamine et al. (2018) group paradigms into inflection classes, and Eskander et al. (2013) induce inflection classes and lemmata from cell labels.

2.3 The Paradigm Cell Filling Problem

The PCFP is the task of predicting unseen inflected forms given morphologically labeled input. PCFP models can guess a word’s plural having only seen its singular, but the child must bootstrap morphological knowledge from scratch, first learning that singular–plural is a relevant distinction. Thus, the PDP must be at least partially solved before the PCFP can be attempted. Yet, as a supervised task, the PCFP is more easily studied, and has received much attention on its own, especially from the word-and-paradigm camp of morphological theory.

Some cognitive works suggest the PCFP cannot be too difficult for any language (Dale et al., 1998; Ackerman and Malouf, 2013, 2015; Blevins et al., 2017; Cotterell et al., 2019). Neural models can test and extend such proposals (Cotterell et al., 2018a; Silfverberg and Hulden, 2018). A related vein of work discusses how speakers inflect nonce words (Berko, 1958; Plunkett and Juola, 1999; Yang, 2015), e.g., is the past tense of *sping*, *spinged* or *spung*? There is a long tradition of modeling past-tense generation with neural networks (Rumelhart and McClelland, 1986; Kirov and Cotterell, 2018; Corkery et al., 2019).

On the engineering side, Durrett and DeNero (2013) inspired much recent work, which has since benefited from large inflectional datasets (Kirov et al., 2018) and advances in neural sequence modeling (Bahdanau et al., 2015). Shared tasks have drawn extra attention to the PCFP (Cotterell et al., 2016a, 2017, 2018c; McCarthy et al., 2019).

3 The Paradigm Discovery Problem

Paradigm discovery is a natural next step in computational morphology, building on related minimally or indirectly supervised works (§2.2) to bridge the gap between unsupervised traditions (§2.1) and supervised work on the PCFP (§2.3). In the PCFP,

Corpus	<i>The cat <u>watched</u> me <u>watching</u> it .</i> <i>I <u>followed</u> the show but she had n't <u>seen</u> it .</i> <i>Let 's <u>see</u> who <u>follows</u> your logic .</i>				
Lexicon	<i>watching, seen, follows, watched, followed, see</i>				
Gold Grid	cell 1	cell 2	cell 3	cell 4	cell 5
paradigm 1	« <i>watch</i> »	« <i>watches</i> »	<i>watching</i>	<i>watched</i>	<i>watched</i>
paradigm 2	« <i>follow</i> »	<i>follows</i>	« <i>following</i> »	<i>followed</i>	<i>followed</i>
paradigm 3	<i>see</i>	« <i>sees</i> »	« <i>seeing</i> »	« <i>saw</i> »	<i>seen</i>

Table 1: An example corpus, lexicon, and gold analyses. All lexicon entries appear in the corpus and, for our experiments, they will all share a POS, here, verb. The grid reflects all possible analyses of syncretic forms (e.g., *walked*, *followed*), even though these only occur in the corpus as PST realizations, like *saw* in Cell 4, not as PST.PTCP, like *seen* in Cell 5. Bracketed «forms» are paradigm mates of attested forms, not attested in the lexicon.

each input form is labeled with its morphosyntactic property set, i.e., the **cell** in the paradigm which it realizes, and its lexeme, i.e., the **paradigm** of related forms to which it belongs. By contrast, to solve the PDP, unlabeled input forms must be assigned cells and paradigms. This task requires learning what syntactic and semantic factors distinguish cells, what combinations of cells can co-occur in a paradigm, and what aspects of a surface form reflect its paradigm and its cell, respectively.

3.1 Task Setup

Table 1 provides an overview of our PDP setup. The first two rows show input data: an unannotated **corpus** and a **lexicon** of forms attested in that corpus. Given only these data, the task is to output a **grid** such that (i) all lexicon forms and all their (potentially unseen) inflectional variants appear in the grid, (ii) all forms appearing in the same column realize the same morphosyntactic cell, and (iii) all forms appearing in the same row belong to the same paradigm. Unattested «forms» to be generated are depicted in brackets in Table 1’s **gold grid**, which shows the ideal output of the system.

Our setup permits multiple forms realizing the same **slot**, i.e., a specific cell in a specific paradigm, a single form realizing multiple slots, and unrealizable empty slots. This supports overabundance (Thornton, 2010, 2011), defectiveness (Sims, 2015), and syncretism (Blevins, 1995; Cotterell et al., 2018b). See Corbett (2005) for more on these phenomena. Experimentally, we constrain the PDP by limiting the lexicon to forms from one POS, but our formalization is more general.

3.2 Data for the PDP

For a given language and POS, we create a corpus, lexicon, and gold grid based on a Universal Depen-

dencies (UD) corpus (Nivre et al., 2016). At a high level, the corpus includes raw, non-UD sentences, and UD sentences stripped of annotations. The lexicon includes all forms occurring in the UD sentences with the specified POS (potentially including variant spellings and typographical errors). The gold grid consists of full paradigms for every word which co-occurs in UD and the UniMorph lexicon (Kirov et al., 2018) with a matching lemma–cell analysis; this is similar to the corpus created by Vylomova et al. (2019). As a system does not know which lexicon forms will be evaluated in the gold grid, it must model the entire lexicon, which should contain a realistic distribution over rare words and inflection classes having been directly extracted from distributional data (Bybee, 2003; Lignos and Yang, 2018).

To ensure the gold grid is reasonably clean, we take all word–lemma–feature tuples from the UD portion of the corpus matching the specified POS and convert the features to a morphosyntactic cell identifier compatible with UniMorph representation as in McCarthy et al. (2018).³ Then we check which word–lemma–cell tuples also occur in UniMorph. For each unique lemma in this intersection, the full paradigm is added as a row to the gold grid. To filter typos and annotation discrepancies, we identify any overabundant slots, i.e., slots realized by multiple forms, and remove all but the most frequently attested realization in UD. While some languages permit overabundance (Thornton, 2010), it often indicates typographical or annotation errors

³Aligning UniMorph and UD requires removing diacritics in (Latin and Arabic) UniMorph corpora to match UD. This can obscure some morphosyntactic distinctions but is more consistent with natural orthography in distributional data. The use of orthographic data for morphological tasks is problematic, but standard in the field, due to scarcity of phonologically transcribed data (Malouf et al., 2020).

Predictions	cell 1	cell 2	cell 3	cell 4
paradigm 1	<i>watched</i>	<i>watching</i>	« <i>watches</i> »	« <i>watch</i> »
paradigm 2	<i>followed</i>	« <i>following</i> »	<i>follows</i>	« <i>follow</i> »
paradigm 3	« <i>seed</i> »	« <i>seeing</i> »	« <i>sees</i> »	<i>see</i>
paradigm 4	« <i>seened</i> »	« <i>seening</i> »	« <i>seens</i> »	<i>seen</i>

Table 2: Toy predictions made from the corpus and lexicon in Table 1, to be evaluated against the toy gold grid. Again, bracketed «forms» are those not occurring in the lexicon.

in UD and UniMorph (Gorman et al., 2019; Malouf et al., 2020). Unlike the gold grid, the lexicon retains overabundant realizations, requiring systems to handle such phenomena.

For each language, the raw sentences used to augment the corpus add over 1 million additional words. For German and Russian, we sample sentences from OpenSubtitles (Lison and Tiedemann, 2016), for Latin, the Latin Library (Johnson et al., 2016), and for English and Arabic, Gigaword (Parker et al., 2011a,b). Supplementary sentences are preprocessed via Moses (Koehn et al., 2007) to split punctuation, and, for supported languages, clitics. Table 3 shows corpus and lexicon sizes.

3.3 Metrics

A system attempting the PDP is expected to output a morphologically organized grid in which rows and columns are arbitrarily ordered, but ideally, each row corresponds to a gold paradigm and each column to a gold cell. Aligning rows to paradigms and columns to cells is non-trivial, making it difficult to simply compute accuracy over gold grid slots. Furthermore, cluster-based metrics (Rosenberg and Hirschberg, 2007) are difficult to apply as forms can appear in multiple columns or rows. Thus, we propose novel metrics that are *lexical*, based on analogical relationships between forms. We propose a set of PDP metrics, to measure how well organized lexicon forms are in the grid, and a set of PCFP metrics, to measure how well the system anticipates unattested inflectional variants. All metrics support non-canonical phenomena such as defective paradigms and overabundant slots.

3.3.1 PDP Metrics

A form f 's **paradigm mates** are all those forms that co-occur in at least one paradigm with f . f 's **paradigm F-score** is the harmonic mean of precision and recall of how well we predicted its paradigm mates when viewed as an information retrieval problem (Manning et al., 2008). We macro-average all forms' paradigm F-scores to compute F_{par} . Qualitatively, F_{par} tells us how well we clus-

ter words that belong to the *same paradigm*. A form f 's **cell mates** are all those forms that co-occur in at least one cell with f . f 's **cell F-score** is the harmonic mean of precision and recall of how well we predicted its cell mates. As before, we macro-average all forms' cell F-scores to compute F_{cell} . Qualitatively, F_{cell} tells us how well we cluster words that belong to the *same cell*. Finally, we propose the F_{grid} metric as the harmonic mean of F_{par} and F_{cell} . F_{grid} is a single number that reflects a system's ability to cluster forms into *both* paradigms and cells. Because we designate separate PCFP metrics to evaluate gold grid forms not in the lexicon, we restrict f 's mates to only include forms that occur in the lexicon.

Consider the proposed grid in Table 2. There are 6 lexicon forms in the gold grid. Starting with *watched*, we correctly propose its only attested paradigm mate, *watching*. Thus, *watched*'s paradigm F-score is 100%. For *see*, we propose no attested paradigm mates, but we should have proposed *seen*. 0 correct out of 1 true paradigm mate from 0 predictions results in an F-score of 0% for *seen*. We continue like this for all 6 attested forms in the gold grid and average their scores to get F_{par} . As for F_{cell} , we correctly predict that *watched*'s only cell mate is *followed*, yielding an F-score of 100%. However, we incorrectly predict that *see* has a cell mate, *seen*, yielding an F-score of 0%; we average each word's F-score to get F_{cell} ; the harmonic mean of F_{par} and F_{cell} gives us F_{grid} .

While F_{grid} handles syncretism, overabundance, defectiveness and mismatched grid dimensions, it is exploitable by focusing exclusively on the best attested cells realized by the most unique forms, since attested cells tend to exhibit a Zipfian distribution (Blevins et al., 2017; Lignos and Yang, 2018). Exploiting F_{grid} in this manner propagates errors when bootstrapping to predict unattested forms and, thus, will be punished by PCFP metrics.

3.3.2 PCFP Metrics

We cannot evaluate the PCFP as in supervised settings (Cotterell et al., 2016a) because proposed

	Lexicon	Corpus	UD
	Types	Tokens	Tokens
Arabic	8,732	1,050,336	223,881
German	19,481	1,270,650	263,804
English	3,330	1,212,986	204,608
Latin	6,903	1,218,377	171,928
Russian	36,321	1,885,302	871,548

Table 3: Statistics regarding the input corpus and lexicon. UD tokens refers to tokens in the corpus originally extracted from UD sentences.

cells and paradigms cannot be trivially aligned to gold cells and paradigms. Instead, we create a test set by sampling 2,000 four-way analogies from the gold grid. The first and second forms must share a row, as must the third and fourth; the first three forms must be attested and the fourth unattested, e.g., *watched* : *watching* :: *seen* : «*seeing*».

From this test set and a proposed grid, we compute a strict **analogy accuracy** (An) metric and a more lenient **lexicon expansion accuracy** (LE) metric. An counts predictions as correct if *all* analogy directions hold in the proposed grid (i.e., *watched*, *watching* and *seen*, «*seeing*» share rows and *watched*, *seen* and *watching*, «*seeing*» share columns). LE counts predictions as correct if the unattested fourth form appears anywhere in the grid. That is, LE asks, for each gold form, if it was predicted in *any* slot in *any* paradigm.

Like the PDP metrics, our PCFP metrics support syncretism, overabundance, defectiveness, etc. One can, however, exploit them by proposing a gratuitous number of cells, paradigms, and syncretisms, increasing the likelihood of completing analogies by chance, though this will reduce F_{grid} . As both PDP and PCFP metrics can be exploited independently but not jointly, we argue that both types of metrics should be considered when evaluating an unsupervised system.

4 Building a Benchmark

This section presents a benchmark system for proposing a morphologically organized grid given a corpus and lexicon. First, we cluster lexicon forms into cells. Then we cluster forms into paradigms given their fixed cell membership. To maintain tractability, clustering assumes a one-to-one mapping of forms to slots. Following cell and paradigm clustering, we predict forms to realize empty slots given one of the lexicon forms assigned to a cell in

the same paradigm. This allows forms to appear in multiple slots, but does not support overabundance, defectiveness, or multi-word inflections.

4.1 Clustering into Cells

We use a heuristic method to determine the number of cells and what lexicon forms to assign to each. Inspired by work on inductive biases in word embeddings (Pennington et al., 2014; Trask et al., 2015; Goldberg, 2016; Avraham and Goldberg, 2017; Tu et al., 2017), we train morphosyntactically biased embeddings on the corpus and use them to k -means cluster lexicon forms into cells. Following Erdmann et al. (2018), we emphasize morphosyntactically salient dimensions in embedding space by manipulating hyperparameters in fastText (Bojanowski et al., 2017). Specifically, to encourage grouping of morphologically related words, fastText computes a word’s embedding as the sum of its subword embeddings for all subword sequences between 3 and 6 characters long (Schütze, 1993). We shorten this range to 2 to 4 to bias the grouping toward shared affixes rather than (usually longer) shared stems. This helps recognize that the same affix is likely to realize the same cell, e.g., *watch* +*ed* and *follow* +*ed*. We limit the context window size to 1; small windows encourage a morphosyntactic bias in embeddings (Erk, 2016).

We determine the number of cells to cluster lexicon forms into, k , via the **elbow method**, which progressively considers adding clusters until the reduction in dispersion levels off (Kodinariya and Makwana, 2013; Bholowalia and Kumar, 2014).⁴ Since Tibshirani et al. (2001)’s popular formalism of the method does not converge on our data, we implement a simpler technique that works in our case. We incrementally increase k , each time recording clustering dispersion, d_k (for consistency, we average d_k over 25 iterations). Starting at $k = 2$, we calculate dispersion deceleration as the difference between the current and previous dispersions:

$$\text{decel}(k) = d_{k-1} - 2(d_k) + d_{k+1} \quad (1)$$

Once $\text{decel}(k)$ decreases below $\sqrt{\text{decel}(2)}$, we take the k^{th} clustering: the $(k + 1)^{\text{th}}$ cluster did not explain enough variation in the embedding space to justify an additional morphosyntactic distinction.

⁴Clustering dispersion is the squared distance of a point from its cluster’s centroid, summed over all points clustered.

4.2 Clustering into Paradigms

Given a clustering of lexicon forms into k cells, denoted as C_1, \dots, C_k , we heuristically cluster each form f into a paradigm, π , as a function of f 's cell, c . For tractability, we assume paradigms are pairwise disjoint and no paradigm contains multiple forms from the same cell. Our algorithm greedily builds paradigms cell by cell. To gauge the quality of a candidate paradigm, we first identify its **base** and **exponents**. Following Beniamine et al. (2018), we define π 's base, b_π , as the longest common subsequence shared by all forms in π .⁵⁶ For each form f in π , we define the exponent x_f as the subsequences of f that remain after removing b_π , i.e., x_f is a tuple of affixes. For example, if π contains words $wxyxz$ and axx , b_π is xx and the exponents are $(\langle w, y, z \rangle)$ and $(\langle a \rangle)$, respectively.⁷ Inspired by unsupervised maximum matching in greedy tokenization (Guo, 1997; Erdmann et al., 2019), we define the following paradigm score function:

$$\text{score}(\pi) = \sum_{(c,f) \in \pi} (|b_\pi| - |x_f|) \quad (2)$$

which scores a candidate paradigm according to the number of base characters minus the number of exponent characters; it can be negative.

Algorithm 1 then details our heuristic clustering approach. We greedily select one or zero forms from each cell to add (via the list concatenation operator \circ) to each paradigm such that the paradigm's score is maximized.⁸ After performing a first pass of paradigm clustering with Algorithm 1, we estimate an unsmoothed probability distribution $p(x | c)$ as follows: we take the number of times each exponent (tuple of affixes) realizes a cell in the output of Algorithm 1 and divide by the number of occurrences of that cell. We use this distribution $p(x | c)$ to construct an exponent penalty:

⁵The fact that we use a *subsequence*, instead of a *substring*, means that we can handle non-concatenative morphology.

⁶We note that the longest common subsequence may be found with a polynomial-time dynamic program; however, there will not exist an algorithm whose runtime is polynomial in the number of strings unless $P = NP$ (Maier, 1978).

⁷We use word start (\langle) and end (\rangle) tokens to distinguish exponents; they do not count as exponent characters in eq. (2).

⁸Algorithm 1 has complexity $\mathcal{O}(|L|^2)$ where $|L|$ is lexicon size. In practice, to make Algorithm 1 tractable, we limit the candidates for f'_j (line 8) to the $n = 250$ forms from cell j nearest to f_i in pre-trained embedding space (trained via FastText with default parameters). This achieves a complexity upper bounded by $\mathcal{O}(|L|nk)$.

Algorithm 1 Paradigm Clustering Algorithm

```

1: input  $C_1, \dots, C_k$ 
2:  $\pi \leftarrow []$ 
3: for  $C_i \in \{C_1, \dots, C_k\}$  do
4:   for  $f_i \in C_i$  do
5:      $\pi \leftarrow [\langle i, f_i \rangle]$ 
6:      $s \leftarrow \text{score}(\pi)$ 
7:     for  $C_j \in \{C_{i+1}, \dots, C_k\}$  do
8:        $f_j \leftarrow \underset{f'_j \in C_j}{\text{argmax}} \text{score}(\pi \circ [\langle j, f'_j \rangle])$ 
9:        $s_{f_j} \leftarrow \text{score}(\pi \circ [\langle j, f_j \rangle])$ 
10:      if  $s_{f_j} > s$  then
11:         $\pi \leftarrow \pi \circ [\langle j, f_j \rangle]$ 
12:         $s \leftarrow s_{f_j}$ 
13:         $C_j.\text{remove}(f_j)$ 
14:       $\pi \leftarrow \pi \circ [\pi]$ 
15: return  $\pi$ 

```

$$\omega(x_f, c) = \begin{cases} 0 & \text{if } \underset{x}{\text{argmax}} p(x | c) = x_f \\ 2 - \frac{p(x_f | c)}{\max_x p(x | c)} & \text{otherwise} \end{cases} \quad (3)$$

Intuitively, if an exponent is the most likely exponent in the cell to which it belongs, the penalty weight is zero and its characters are not subtracted from the score. Otherwise, the weight is in the interval $[1, 2]$ such that each exponent character is penalized at least as harshly but no more than twice as harshly than in the first pass, according to the exponent's likelihood. We use this exponent penalty weight to define a penalized score function:

$$\text{score}_\omega(\pi) = \sum_{(c,f) \in \pi} (|b_\pi| - |x_f| \omega(x_f, c)) \quad (4)$$

We then re-run Algorithm 1, swapping out $\text{score}(\cdot)$ for $\text{score}_\omega(\cdot)$, to re-cluster forms into paradigms. Empirically, we find that harsher exponent penalties—i.e., forcing weights to be greater than 1 for suboptimal exponents—lead to higher paradigm precision in this second pass. For an example, consider candidate paradigm $[\langle \langle \rangle, \text{watched}, \langle \rangle, \langle \rangle, \langle \rangle]$. If we add nothing, each character of *watched* can be analyzed as part of the base, yielding a score of 7. What if we attempt to add *watching*—pre-determined to belong to column 5 during cell clustering? Candidate paradigm $[\langle \rangle, \text{watched}, \langle \rangle, \langle \rangle, \text{watching}]$ increases the number of base characters to 10 (*watch* shared by 2 words), but yields a score of 5 after subtracting

the characters from both exponents, (*ed*>) and (*ing*>). Hence, we do not get this paradigm right on our first pass, as $5 < 7$. Yet, after the first pass, should (*ed*>) and (*ing*>) be the most frequent exponents in the second and fifth cells, the second pass will be different. Candidate paradigm [\llcorner , *watched*, \llcorner , \llcorner , *watching*] is not penalized for either exponent, yielding a score of 10, thereby allowing *watching* to be added to the paradigm.

4.3 Reinflection

We now use the output of the clustering by cell and paradigm to bootstrap the PCFP. We use a Transformer (Vaswani et al., 2017) to predict the forms that realize empty slots. Transformer-based neural transducers constitute the state of the art for the PCFP.⁹ In Cotterell et al. (2016b)’s terms, we reinflect the target from one of the non-empty source cells in the same paradigm. We select the source from which we can most reliably reinflect the target. We quantify this reliability by calculating the accuracy with which each target cell’s realizations were predicted from each source cell’s realizations in our development set. For each target cell, we rank our preferred source cells according to accuracy.

To generate train and development sets, we create instances for every possible pair of realizations occurring in the same paradigm (90% train, 10% development). We pass these instances into the Transformer, flattening cells and characters into a single sequence. Neural models for reinflection often perform poorly when the training data are noisy. We mitigate this via the harsh exponent penalty weights (eq. (3)) which encourage high paradigm precision during clustering.

5 Results and Discussion

Table 4 shows results for two versions of our benchmark system: BENCH, as described in §4, and GOLD k , with the number of cells oracularly set to the ground truth. For reference, we also report a supervised benchmark, SUP, which assumes a gold grid as input, then solves the PCFP exactly as the benchmark does. In terms of the PDP, clustering assigns lexicon forms to paradigms (46–82%) more accurately than to cells (26–80%). Results are high for English, which has the fewest gold cells, and

⁹We use the following hyperparameters: $N = 4$, $d_{model} = 128$, $d_{ff} = 512$. Remaining hyperparameters retain their default values as specified in Vaswani et al. (2017). Our models are trained for 100 epochs in batches of 64. We stop early after 20 epochs without improvement on the development set.

		PDP			PCFP			
		Cells	Paradigms	F_{cell}	F_{par}	F_{grid}	An	LE
Arabic nouns – 8,732 forms								
SUP	27	4,283					85.9	87.0
BENCH	12.8	5,279.3	39.9	48.5	43.7		16.8	49.5
GOLD k	27	4,930.3	25.9	46.4	33.1		16.1	57.2
German nouns – 19,481 forms								
SUP	8	17,018					72.2	74.9
BENCH	7.3	17,073.3	35.2	59.4	43.3		14.2	56.7
GOLD k	8	16,836.0	29.4	66.6	40.8		14.8	60.4
English verbs – 3,330 forms								
SUP	5	1,801					80.4	80.7
BENCH	7.5	1,949.5	64.0	80.1	71.1		52.0	67.5
GOLD k	5	1,977.3	79.6	82.1	80.8		54.7	69.4
Latin nouns – 6,903 forms								
SUP	12	3,013					80.0	88.0
BENCH	13.0	3,746.5	38.8	73.2	50.6		17.2	72.9
GOLD k	12	3,749.0	39.9	71.6	51.3		17.5	72.6
Russian nouns – 36,321 forms								
SUP	14	14,502					94.7	96.8
BENCH	16.5	19,792.0	44.5	72.2	55.0		31.9	86.2
GOLD k	14	20,944.0	45.7	69.1	55.0		31.6	84.3

Table 4: PDP and PCFP results for all languages and models, averaged over 4 runs. Metrics are defined in §3.3. An refers to analogy accuracy and LE to the lexicon expansion accuracy.

lower elsewhere. In German, Latin, and Russian, our benchmark proposes nearly as many cells as GOLD k , thus performing similarly. For English, it overestimates the true number and performs worse. For Arabic, it severely underestimates k but performs better, likely due to the orthography: without diacritics, the three case distinctions become obscured in almost all instances. In general, fixing the true number of cells can be unhelpful because syncretism and the Zipfian distribution of cells creates situations where certain gold cells are too difficult to detect. Allowing the system to choose its own number of cells lets it focus on distinctions for which there is sufficient distributional evidence.

As for the PCFP, our benchmark system does well on lexicon expansion accuracy and poorly on the analogy task. While lexicon expansion accuracy (50–86% compared to 72–97% for SUP) shows that the benchmark captures meaningful inflectional trends, analogy accuracy demonstrates vast room for improvement in terms of consistently organizing cell-realizations across paradigms. English is the only language where analogy accuracy is within half of SUP’s upper bound. A major reason for low analogy accuracy is that forms, despite being clustered into paradigms well, get assigned

SG					PL					Gloss
NOM	GEN	DAT	ACC	ABL	NOM	GEN	DAT	ACC	ABL	
<i>serv-us</i>	<i>i</i>	<i>o</i>	<i>um</i>	<i>o</i>	<i>i</i>	<i>orum</i>	<i>is</i>	<i>os</i>	<i>is</i>	“slave.M”
<i>serv-a</i>	<i>ae</i>	<i>ae</i>	<i>am</i>	<i>a</i>	<i>ae</i>	<i>arum</i>	<i>is</i>	<i>as</i>	<i>is</i>	“slave.F”
<i>frat-er</i>	<i>ris</i>	<i>ri</i>	<i>rem</i>	<i>re</i>	<i>res</i>	<i>rum</i>	<i>ribus</i>	<i>res</i>	<i>ribus</i>	“brother”

Table 5: Suffixal exponents for each cell in the paradigm of three Latin nouns from different inflection classes.

Cell	Interpretations	Suffix
0	ACC.SG (.51), GEN.PL (.45)	<i>um</i>
1	ACC.PL (.71), NOM.PL (.27)	<i>s</i>
2	ACC.SG (.99)	<i>m</i>
3	ABL.PL (.52), GEN.SG (.40)	<i>is</i>
4	NOM.SG (.39), ABL.SG (.36)	<i>a</i>
5	ABL.SG (.62), NOM.SG (.36)	<i>o</i>
6	GEN.SG (.46), DAT.SG (.30)	<i>i</i>
7	ABL.PL (.77), DAT.PL (.25)	<i>s</i>
8	NOM.SG (.67), ABL.SG (.22)	\emptyset
9	ABL.SG (.936)	<i>e</i>
10	ABL.SG (.5), GEN.SG (.28)	<i>e</i>
11	NOM.SG (.87), ACC.PL (.16)	<i>us</i>

Table 6: System clustering of Latin nouns.

to the wrong cell, or the same gold cell gets misaligned across paradigms from different inflection classes. We discuss this phenomenon in more detail below.

5.1 Latin Noun Error Analysis

A detailed analysis of Latin nouns (also analyzed by Stump and Finkel (2015) and Beniamine et al. (2018)) reveals challenges for our system. Table 5 shows the inflectional paradigms for three Latin nouns exemplifying different inflection classes, which are mentioned throughout the analysis. In keeping with the UD standard, there are no diacritics for long vowels in the table.

One major challenge for our system is that similar affixes can mark different cells in different inflection classes, e.g. the ACC.SG of *servus* “slave.M” ends in *um*, as does the GEN.PL of *frater* “brother”. Table 6 shows system-posed cells, the gold cells they best match to, and the longest suffix shared by 90% of their members. The system is often misled by shared affixes, e.g., cell 0 is evenly split between ACC.SG and GEN.PL, driven by the suffix *um* (cells 3 (*is*) and 4 (*a*) suffer from this as well). This kind of confusion could be resolved with better context modeling, as each distinct underlying cell, despite sharing a surface affix, occurs in distinct distributional contexts. We observe that

the current system often fails to make use of context to handle some misleading suffixes. However, Cell 7 correctly groups ABL.PL forms marked with both *is* and *ibus*, excluding other suffixes ending in *s*. Similarly, cell 8 contains NOM.SG forms with heterogeneous endings, e.g., *r*, *ix* and *ns*.

In some cases, the system misinterprets derivational processes as inflectional, combining gold paradigms. Derivational relatives *servus* and *serva*, male and female variants of “slave”, are grouped into one paradigm, as are *philosophos* “philosopher” and *philosophia* “philosophy.” In other cases, cell clustering errors due to shared suffixes create spurious paradigms. After falsely clustering gold paradigm mates *servum* (ACC.SG) and *servorum* (GEN.PL) into the same cell, we must assign each to separate paradigms during paradigm clustering. This suggests clustering cells and paradigms jointly might avoid error propagation in future work.

We also find that clustering errors lead to PCFP errors. For *servus/a*, the neural reinlector predicts *servibus* in cell 8 with a suffix from the wrong inflection class, yet the slot should not be empty in the first place. The correct form, *servis*, is attested, but was mistakenly clustered into cell 3.

5.2 Benchmark Analysis

Table 7 evaluates variants of the benchmark to determine the contribution of several system–task components in Arabic and Latin. We consider augmenting and shrinking the corpus. We also reset the fastText hyperparameters used to achieve a morphosyntactic inductive bias to their default values (no affix or window bias) and consider two constant exponent penalty weights ($\omega(x_f, c) = 1$ and $\omega(x_f, c) = 0$) instead of our heuristic weight defined in eq. (3). Finally, we consider selecting random sources for PCFP reinlection instead of identifying reliable sources. For all variants, the number of cells is fixed to the ground truth.

Corpus Size We consider either using a smaller corpus containing only the UD subset, or using a larger corpus containing 15 (Latin) or 100 (Ara-

Paradigms	PDP			PCFP		
	F_{cell}	F_{par}	F_{grid}	An	LE	
Arabic nouns – 27 cells						
GOLD k	4,930.3	25.9	46.4	33.1	16.1	57.2
larger corpus	5,039.5	29.1	37.5	32.8	20.4	49.2
smaller corpus	5,004.0	18.8	37.7	24.9	9.5	42.1
no affix bias	4,860.3	21.5	47.7	29.7	16.3	43.5
no window bias	4,978.5	24.0	47.5	31.8	17.6	55.8
$\omega(x, c) = 1$	3,685.0		34.4	28.8	5.2	35.5
$\omega(x, c) = 0$	1,310.5		10.0	13.9	0.1	5.8
random sources					16.3	55.9
Latin nouns – 12 cells						
GOLD k	3,749.0	39.9	71.6	51.3	17.5	72.6
larger corpus	3,529.5	42.8	79.1	55.5	16.2	69.9
smaller corpus	4,381.5	30.7	49.1	37.8	14.6	51.1
no affix bias	3,906.8	37.1	68.2	48.1	22.7	66.6
no window bias	3,756.5	42.0	71.2	52.8	17.9	70.9
$\omega(x, c) = 1$	3,262.5		67.1	49.6	11.0	52.9
$\omega(x, c) = 0$	1,333.3		26.3	31.7	0.7	7.1
random sources					16.5	72.3

Table 7: Benchmark variations demonstrating the effects of various factors, averaged over 4 runs.

bic) million words from additional supplementary sentences. As expected, performance decreases for smaller corpora, but it does not always increase for larger ones, potentially due to domain differences between UD and the supplemental sentences. Interestingly, F_{cell} always increases with larger corpora, yet this can lead to worse F_{par} scores, more evidence of error propagation that might be avoided with joint cell–paradigm clustering.

Embedding Morphosyntactic Biases Targeting affix embeddings by shrinking the default fastText character n -gram sizes seems to yield a much more significant effect than shrinking the context window. In Latin, small context windows can even hurt performance slightly, likely due to extremely flexible word order, where agreement is often realized over non-adjacent words.

Exponent Penalties When clustering paradigms with the penalty weight $\omega(x, c) = 1$, (which is equivalent to just running the first pass of paradigm clustering), we see a steep decline in performance as opposed to the proposed heuristic weighting. It is even more detrimental to not penalize exponents at all (i.e., $\omega(x, c) = 0$), but maximize the base characters in paradigms without concern for size or likelihoods of exponents. Given allomorphic variation and multiple inflection classes, we ideally want a penalty weight which is lenient to more than just the single most likely exponent, but without

supervised data, it is difficult to determine when to stop being lenient and start being harsh in a language agnostic manner. Our choice to be harsh by default proposes fewer false paradigm mates, yielding less noisy input to train the reinflection model. In a post-hoc study, we calculated GOLD k PCFP scores on *pure* analogies only, where the first three attested forms were assigned correctly during clustering. Pure analogy PCFP scores were still closer to GOLD k 's performance than SUP's for all languages. This suggests most of the gap between GOLD k and SUP is due to noisy training on bad clustering assignments, not impossible test instances created by bad clustering assignments. This supports our choice of harsh penalties and suggests future work might reconsider clustering decisions given the reinflection model's confidence.

Reinflection Source Selection During reinflection, feeding the Transformer random sources instead of learning the most reliable source cell for each target cell slightly hurts performance. The margin is small, though, as most paradigms have only one attested form. In preliminary experiments, we also tried jointly encoding all available sources instead of just the most reliable, but this drastically lowers performance.

6 Conclusion

We present a framework for the paradigm discovery problem, in which words attested in an unannotated corpus are analyzed according to the morphosyntactic property set they realize and the paradigm to which they belong. Additionally, unseen inflectional variants of seen forms are to be predicted. We discuss the data required to undertake this task, a benchmark for solving it, and multiple evaluation metrics. We believe our benchmark system represents a reasonable approach to solving the problem based on past work and highlights many directions for improvement, e.g. joint modeling and making better use of distributional semantic information.

Acknowledgments

The authors would like to thank the members of New York University Abu Dhabi's CAMEL Lab, Marie-Catherine de Marneffe, Eleanor Chodroff, Katharina Kann, and Markus Dreyer. We acknowledge the support of the High Performance Computing Center at New York University Abu Dhabi. Finally, we wish to thank the anonymous reviewers at EMNLP 2019 and ACL 2020 for their feedback.

References

- Farrell Ackerman, James P. Blevins, and Robert Malouf. 2009. Parts and wholes: Implicative patterns in inflectional paradigms. In *Analogy in Grammar: Form and Acquisition*, pages 54–82. Oxford University Press.
- Farrell Ackerman and Robert Malouf. 2013. Morphological organization: The low conditional entropy conjecture. *Language*, 89(3):429–464.
- Farrell Ackerman and Robert Malouf. 2015. The No Blur Principle effects as an emergent property of language systems. In Anna E. Jurgensen, Hannah Sande, Spencer Lamoureux, Kenny Baclawski, and Alison Zerbe, editors, *Proceedings of the Forty-First Annual Meeting of the Berkeley Linguistics Society*, pages 1–14. Berkeley Linguistics Society.
- Oded Avraham and Yoav Goldberg. 2017. [The interplay of semantics and morphology in word embeddings](#). In *Proceedings of the 15th Conference of the European Chapter of the Association for Computational Linguistics: Volume 2, Short Papers*, pages 422–426, Valencia, Spain. Association for Computational Linguistics.
- Dzmitry Bahdanau, Kyunghyun Cho, and Yoshua Bengio. 2015. Neural machine translation by jointly learning to align and translate. In *3rd International Conference on Learning Representations, ICLR 2015*.
- Sacha Beniamine, Olivier Bonami, and Benoît Sagot. 2018. Inferring inflection classes with description length. *Journal of Language Modelling*, 5(3):465–525.
- 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.
- Jean Berko. 1958. The child’s learning of English morphology. *Word*, 14(2-3):150–177.
- Purnima Bholowalia and Arvind Kumar. 2014. EBK-means: A clustering technique based on elbow method and k -means in WSN. *International Journal of Computer Applications*, 105(9).
- James P. Blevins. 1995. Syncretism and paradigmatic opposition. *Linguistics and Philosophy*, 18(2):113–152.
- James P. Blevins, Petar Milin, and Michael Ramscar. 2017. The Zipfian paradigm cell filling problem. In *Perspectives on Morphological Organization*, pages 139–158. Brill.
- 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.
- Gilles Boyé and Gauvain Schaalchli. 2019. Realistic data and paradigms: The paradigm cell finding problem. *Morphology*, 29(2):199–248.
- Joan Bybee. 2003. Mechanisms of change in grammaticalization: The role of frequency. In Brian D. Joseph and Richard D. Janda, editors, *The Handbook of Historical Linguistics*, pages 602–623. Blackwell.
- Greville G. Corbett. 2005. The canonical approach in typology. *Linguistic Diversity and Language Theories*, 25:49.
- Maria Corkery, Yevgen Matuskevych, and Sharon Goldwater. 2019. [Are we there yet? Encoder-decoder neural networks as cognitive models of English past tense inflection](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 3868–3877, Florence, Italy. Association for Computational Linguistics.
- Ryan Cotterell, Christo Kirov, Mans Hulden, and Jason Eisner. 2018a. On the diachronic stability of irregularity in inflectional morphology. *arXiv preprint arXiv:1804.08262*.
- Ryan Cotterell, Christo Kirov, Mans Hulden, and Jason Eisner. 2019. [On the complexity and typology of inflectional morphological systems](#). *Transactions of the Association for Computational Linguistics*, 7:327–342.
- Ryan Cotterell, Christo Kirov, Sabrina J. Mielke, and Jason Eisner. 2018b. [Unsupervised disambiguation of syncretism in inflected lexicons](#). In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 2 (Short Papers)*, pages 548–553, New Orleans, Louisiana. 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, David Yarowsky, Jason Eisner, and Mans Hulden. 2018c. [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. 2017. [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, Hinrich Schütze, and Jason Eisner. 2016b. [Morphological smoothing and extrapolation of word embeddings](#). In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1651–1660, Berlin, Germany. Association for Computational Linguistics.
- Mathias Creutz and Krista Lagus. 2005. *Unsupervised morpheme segmentation and morphology induction from text corpora using Morfessor 1.0*. Helsinki University of Technology.
- Robert Dale, Barbara Di Eugenio, and Donia Scott. 1998. [Introduction to the special issue on natural language generation](#). *Computational Linguistics*, 24(3):345–353.
- Markus Dreyer and Jason Eisner. 2011. [Discovering morphological paradigms from plain text using a Dirichlet process mixture model](#). In *Proceedings of the 2011 Conference on Empirical Methods in Natural Language Processing*, pages 616–627, Edinburgh, Scotland, UK. 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.
- 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 and Nizar Habash. 2018. [Complementary strategies for low resourced morphological modeling](#). In *Proceedings of the Fifteenth Workshop on Computational Research in Phonetics, Phonology, and Morphology*, pages 54–65, Brussels, Belgium. Association for Computational Linguistics.
- Alexander Erdmann, Salam Khalifa, Mai Oudah, Nizar Habash, and Houda Bouamor. 2019. [A little linguistics goes a long way: Unsupervised segmentation with limited language specific guidance](#). In *Proceedings of the 16th Workshop on Computational Research in Phonetics, Phonology, and Morphology*, pages 113–124, Florence, Italy. Association for Computational Linguistics.
- Alexander Erdmann, Nasser Zalmout, and Nizar Habash. 2018. [Addressing noise in multidialectal word embeddings](#). In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, pages 558–565, Melbourne, Australia. Association for Computational Linguistics.
- Katrin Erk. 2016. [What do you know about an alligator when you know the company it keeps?](#) *Semantics and Pragmatics*, 9:17–1.
- Ramy Eskander, Nizar Habash, and Owen Rambow. 2013. [Automatic extraction of morphological lexicons from morphologically annotated corpora](#). In *Proceedings of the 2013 Conference on Empirical Methods in Natural Language Processing*, pages 1032–1043, Seattle, Washington, USA. Association for Computational Linguistics.
- Eric Gaussier. 1999. [Unsupervised learning of derivational morphology from inflectional lexicons](#). In *Unsupervised Learning in Natural Language Processing*.
- Yoav Goldberg. 2016. [A primer on neural network models for natural language processing](#). *Journal of Artificial Intelligence Research*, 57:345–420.
- John Goldsmith. 2001. [Unsupervised learning of the morphology of a natural language](#). *Computational Linguistics*, 27(2):153–198.
- Sharon Goldwater. 2007. *Nonparametric Bayesian Models of Lexical Acquisition*. Ph.D. thesis.
- 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.
- Jin Guo. 1997. [Critical tokenization and its properties](#). *Computational Linguistics*, 23(4):569–596.
- 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 (ACL)*.
- Kyle P. Johnson, P. J. Burns, L. Hollis, M. Pozzi, A. Shilo, S. Margheim, G. Badger, and E. Bell. 2016. CLTK: The classical languages toolkit.
- David King and Michael White. 2018. [The OSU realizer for SRST ‘18: Neural sequence-to-sequence inflection and incremental locality-based linearization](#).

- In *Proceedings of the First Workshop on Multilingual Surface Realisation*, pages 39–48, Melbourne, Australia. Association for Computational Linguistics.
- Christo Kirov and Ryan Cotterell. 2018. [Recurrent neural networks in linguistic theory: Revisiting Pinker and Prince \(1988\) and the past tense debate](#). *Transactions of the Association for Computational Linguistics*, 6:651–665.
- Christo Kirov, Ryan Cotterell, John Sylak-Glassman, Géraldine Walther, Ekaterina Vylomova, Patrick Xia, Manaal Faruqui, Sabrina J. Mielke, Arya McCarthy, Sandra Kübler, David Yarowsky, Jason Eisner, and Mans Hulden. 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).
- Trupti M. Kodinariya and Prashant R. Makwana. 2013. Review on determining number of cluster in k -means clustering. *International Journal of Advance Research in Computer Science and Management Studies*, 1(6):90–95.
- Philipp Koehn, Hieu Hoang, Alexandra Birch, Chris Callison-Burch, Marcello Federico, Nicola Bertoldi, Brooke Cowan, Wade Shen, Christine Moran, Richard Zens, Chris Dyer, Ondřej Bojar, Alexandra Constantin, and Evan Herbst. 2007. [Moses: Open source toolkit for statistical machine translation](#). In *Proceedings of the 45th Annual Meeting of the Association for Computational Linguistics Companion Volume Proceedings of the Demo and Poster Sessions*, pages 177–180, Prague, Czech Republic. Association for Computational Linguistics.
- Jackson Lee. 2015. [Morphological paradigms: Computational structure and unsupervised learning](#). In *Proceedings of the 2015 Conference of the North American Chapter of the Association for Computational Linguistics: Student Research Workshop*, pages 161–167, Denver, Colorado. Association for Computational Linguistics.
- Constantine Lignos and Charles Yang. 2018. Morphology and language acquisition. In Andrew Hippiusley and Gregory T. Stump, editors, *Cambridge Handbook of Morphology*, pages 765–791. Cambridge University Press.
- Pierre Lison and Jörg Tiedemann. 2016. [OpenSubtitles2016: Extracting large parallel corpora from movie and TV subtitles](#). In *Proceedings of the Tenth International Conference on Language Resources and Evaluation (LREC’16)*, pages 923–929, Portorož, Slovenia. European Language Resources Association (ELRA).
- David Maier. 1978. The complexity of some problems on subsequences and supersequences. *Journal of the ACM*, 25(2):322–336.
- Robert Malouf, Farrell Ackerman, and Artrus Semenuks. 2020. Lexical databases for computational analyses: A linguistic perspective. *Proceedings of the Society for Computation in Linguistics*, 3(1):297–307.
- Christopher D. Manning, Prabhakar Raghavan, and Hinrich Schütze. 2008. *Introduction to Information Retrieval*. Cambridge University Press.
- 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, Ryan Cotterell, and Mans Hulden. 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.
- 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.
- Joakim Nivre, Marie-Catherine de Marneffe, Filip Ginter, Yoav Goldberg, Jan Hajič, Christopher D. Manning, Ryan McDonald, Slav Petrov, Sampo Pyysalo, Natalia Silveira, Reut Tsarfaty, and Daniel Zeman. 2016. [Universal dependencies v1: A multilingual treebank collection](#). In *Proceedings of the Tenth International Conference on Language Resources and Evaluation (LREC’16)*, pages 1659–1666, Portorož, Slovenia. European Language Resources Association (ELRA).
- Robert Parker, David Graff, Ke Chen, Junbo Kong, and Kazuaki Maeda. 2011a. Arabic Gigaword Fifth Edition. LDC catalog number No. LDC2011T11, ISBN 1-58563-595-2.
- Robert Parker, David Graff, Junbo Kong, Ke Chen, and Kazuaki Maeda. 2011b. English gigaword. *Linguistic Data Consortium*.
- Jeffrey Pennington, Richard Socher, and Christopher Manning. 2014. [GloVe: Global vectors for word representation](#). In *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 1532–1543, Doha, Qatar. Association for Computational Linguistics.
- Steven Pinker. 2001. Four decades of rules and associations, or whatever happened to the past tense debate. *Language, the Brain, and Cognitive Development: Papers in Honor of Jacques Mehler*, pages 157–179.

- Kim Plunkett and Patrick Juola. 1999. A connectionist model of English past tense and plural morphology. *Cognitive Science*, 23(4):463–490.
- Andrew Rosenberg and Julia Hirschberg. 2007. **V-measure: A conditional entropy-based external cluster evaluation measure**. In *Proceedings of the 2007 Joint Conference on Empirical Methods in Natural Language Processing and Computational Natural Language Learning (EMNLP-CoNLL)*, pages 410–420, Prague, Czech Republic. Association for Computational Linguistics.
- David E. Rumelhart and James L. McClelland. 1986. On learning the past tenses of English verbs.
- Hinrich Schütze. 1993. Word space. In *Advances in Neural Information Processing Systems*, pages 895–902.
- 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.
- Andrea D. Sims. 2015. *Inflectional Defectiveness*, volume 148. Cambridge University Press.
- Gregory Stump and Raphael A. Finkel. 2015. The complexity of inflectional systems. *Linguistics Vanguard*, 1(1):101–117.
- Anna Thornton. 2011. Overabundance (multiple forms realizing the same cell): A non-canonical phenomenon in Italian verb morphology.
- Anna M. Thornton. 2010. Towards a typology of overabundance. In *Décembrettes 7: International Conference on Morphology, University of Toulouse*, pages 2–3.
- Robert Tibshirani, Guenther Walther, and Trevor Hastie. 2001. Estimating the number of clusters in a data set via the gap statistic. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 63(2):411–423.
- Andrew Trask, David Gilmore, and Matthew Russell. 2015. **Modeling order in neural word embeddings at scale**. In *Proceedings of the 32nd International Conference on Machine Learning*, volume 37 of *Proceedings of Machine Learning Research*, pages 2266–2275, Lille, France. PMLR.
- Lifu Tu, Kevin Gimpel, and Karen Livescu. 2017. **Learning to embed words in context for syntactic tasks**. In *Proceedings of the 2nd Workshop on Representation Learning for NLP*, pages 265–275, Vancouver, Canada. 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. In *Advances in Neural Information Processing Systems*, pages 5998–6008.
- Ekaterina Vylomova, Ryan Cotterell, Trevor Cohn, Timothy Baldwin, and Jason Eisner. 2019. **Contextualization of morphological inflection**. 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 2018–2024, Minneapolis, Minnesota. 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.
- Charles Yang. 2015. For and against frequencies. *Journal of Child Language*, 42(2):287–293.
- Olga Zamaraeva, Kristen Howell, and Emily M. Bender. 2019. **Handling cross-cutting properties in automatic inference of lexical classes: A case study of Chintang**. In *Proceedings of the 3rd Workshop on the Use of Computational Methods in the Study of Endangered Languages Volume 1 (Papers)*, pages 28–38, Honolulu. Association for Computational Linguistics.