

An Investigation of Noise in Morphological Inflection

Adam Wiemerslage[◇] Changbing Yang[♣] Garrett Nicolai[♣]
Miikka Silfverberg[♣] Katharina Kann[◇]

[◇]University of Colorado Boulder

[♣]University of British Columbia

{adam.wiemerslage, katharina.kann}@colorado.edu

Abstract

With a growing focus on morphological inflection systems for languages where high-quality data is scarce, training data noise is a serious but so far largely ignored concern. We aim at closing this gap by investigating the types of noise encountered within a pipeline for truly unsupervised morphological paradigm completion and its impact on morphological inflection systems: First, we propose an error taxonomy and annotation pipeline for inflection training data. Then, we compare the effect of different types of noise on multiple state-of-the-art inflection models. Finally, we propose a novel character-level masked language modeling (CMLM) pretraining objective and explore its impact on the models' resistance to noise. Our experiments show that various architectures are impacted differently by separate types of noise, but encoder-decoders tend to be more robust to noise than models trained with a copy bias. CMLM pretraining helps transformers, but has lower impact on LSTMs.

1 Introduction

Neural morphological inflection has shown impressive results for a huge variety of languages (Cotterell et al., 2016; Kodner et al., 2022). Performance is impressive even for languages with very little supervised inflection data, and often generalizes to unseen lemmas. However, the language settings that arguably stand to benefit the most from these tools, those with extremely sparse normalized texts, are less likely to have clean, gold-standard data. Despite this, inflection training data noise is rarely addressed nor evaluated in popular benchmarks. Noise, like incorrect annotations, or mixed dialects or orthographies, can arise in inflection data from web-scraping issues (Gorman et al., 2019; McCarthy et al., 2020a), human error or changes in writing standards in field data (Moeller et al., 2020), or system errors when bootstrapping silver-standard data in an unsupervised fashion

(Kann et al., 2020; Erdmann et al., 2020; Wiemerslage et al., 2022). Unsupervised systems are also prone to *over-regularization*, where the dataset contains few or no irregular samples. Datasets derived from FSTs or textbook examples can also display over-regularization (Vylomova et al., 2020).

In this work, we investigate the impact of noise on inflection generation systems. We build an automatic pipeline for annotating inflection noise and explore the noise distribution that arises in an unsupervised system for bootstrapping inflection data. We measure the impact of noise on several state-of-the-art neural inflection generation systems that we benchmark on the SIGMORPHON 2017 shared task development set (Cotterell et al., 2017). Finally, we explore a novel character-level masked language modeling (CMLM) pretraining objective to mitigate the impact of noise during training. By this, we aim to shed light on which architectures and training methods are how robust to noise, which types of inflection noise should be targeted in filtering approaches, and how conservatively unsupervised systems should sample inflection pairs.

We find that noise related to slot alignment issues is more common, but also less impactful than noise related to paradigm induction issues. Architectures with an inductive bias towards copying from the lemma are more effective on datasets that lack sample diversity, but more typical encoder-decoder models are more robust to noise. Standard encoder-decoders display better performance on noisy data when pretrained with CMLM, which is especially effective for the Transformer. Our code and data are publicly available.¹

2 Noisy Training Data

Our data comprises four languages: Icelandic, German, Swedish, and Russian. We describe the source of training data below.

¹<https://github.com/Adamits/morphological-inflection-noise>

	Deu	Isl	Rus	Swe
Original	40956	32225	128326	59814
Annotated	6444	4735	14171	6054

Table 1: Size of training set for each language before and after filtering out pairs that cannot be annotated.

tUMPC We focus on training data based on Wiemerslage et al. (2022), who propose a system for truly unsupervised morphological paradigm completion (tUMPC). Starting with the Bible corpus (McCarthy et al., 2020b) for a given language, tUMPC first clusters data into paradigms using the system from McCurdy et al. (2021). Second, paradigms are clustered into parts-of-speech. Finally, tUMPC aligns similarly inflected forms across paradigms that belong to the same POS. The result is a dataset of paradigms for which each type is marked with its inflectional slot. We take all possible pairs of words from the tUMPC paradigms to form inflection training data. This dataset is likely to contain noise due to errors in the learning process. In §4, we discuss this noise in detail.

UniMorph tUMPC relies on frequency thresholds to find productive inflection transformations, which makes it likely that morpho-phonological variations and irregular inflections may not be well-attested in the data. To control for the potential lack of sample diversity in our training data, we create a second training dataset wherein all samples marked as correct are replaced by a pair sampled from UniMorph (Sylak-Glassman et al., 2015; Batsuren et al., 2022), a database of morphological paradigms covering hundreds of languages. We sample pairs from the same MSDs as the tUMPC correct samples. Following recent results (Liu and Hulden, 2022; Goldman et al., 2022; Kodner et al., 2022), we also ensure that the lemma overlap with the evaluation set is the same as in the original tUMPC data. The resulting dataset has higher lemma diversity, without introducing missing MSDs, and maintains lemma overlap with the evaluation set. We combine this data with the noisy training data from tUMPC to investigate the impact of noise when sample diversity is less pervasive.

Slot Mapping tUMPC slots are arbitrary identifiers with no grammatical meaning. In order to compare tUMPC to UniMorph, and because we evaluate on a benchmark sampled from UniMorph, we need to map each tUMPC slot to a unique Uni-

Morph slot. UniMorph slots are bundles of morphological tags that express categories like person or tense. We refer to these bundles as morpho-syntactic descriptions (MSDs). Since the tUMPC training pairs have low overlap with UniMorph, we require a morphological analyzer to get ground truth morphological tags of tUMPC types. Here, we use morphological analyzers from Apertium (Forcada et al., 2011), which we map to UniMorph MSDs. Then, we apply the mapped Apertium to all tUMPC types, resulting in ground truth MSDs for every training sample in tUMPC. We also want the slots predicted by tUMPC expressed as UniMorph MSDs during training. For this, we compare the unsupervised slot and gold MSD of all tUMPC types and find the alignment that maximizes their overlap following the evaluation metric from Jin et al. (2020). Every training sample can then be assigned the (noisy) MSD from this alignment. For more details see Section A.1 in the appendix.

Many inflection pairs created by tUMPC could not be reliably annotated by our pipeline and are thus filtered out. Table 1 shows the amount of training data before and after filtering. These are inflection pairs that are not in the Apertium lexicon, but are also not determined to be noise. We additionally filter pairs containing a form with an Apertium analysis that we could not reliably map to UniMorph due to disagreements between the two resources. We compare the distribution of morphological tags for our dataset according to Apertium before and after applying this filtering step, and find no systematic difference. We conclude that removing forms due to mapping errors does not bias our data towards certain inflections. These filtering steps remove a majority of data, but still leave reasonably sized training sets for each language.

3 Related Work

Morphological Inflection Morphological inflection is the task of generating a word form given a lemma and target MSD. For example, given the verb *laugh* and the MSD expressing past tense, the goal is to generate *laughed*. Encoder-decoder neural approaches have largely dominated morphological inflection in recent years (Faruqui et al., 2016; Kann and Schütze, 2016), where the full target string is decoded from a neural network. But neural models that bias the task towards transducing input strings have been shown to be successful in low-resource data settings (Aharoni and Gold-

berg, 2017; Makarov et al., 2017; Makarov and Clemenide, 2018; Sharma et al., 2018). Shared tasks on morphological inflection (Cotterell et al., 2016) have spurred large interest in the task, and also serve as an evaluation benchmark.

Learning from Noisy Data Several approaches have been proposed for mitigating the impact of noise in machine learning, for example: confidence weighting (Rebbapragada and Brodley, 2007), loss correction (Patrini et al., 2017), and noise-contrastive estimation (Gutmann and Hyvärinen, 2010). For a recent survey on noise robust neural networks, see Song et al. (2022). Most approaches consider classification, framing noise as label corruption. However, there is also work exploring noise in tasks like machine translation (Khayrallah and Koehn, 2018; Michel and Neubig, 2018), as well as morphological disambiguation (Zalmout et al., 2018). We focus on morphological inflection, a conditional generation problem.

Morphological Inflection with Training Noise

Morphological inflection with noisy data has largely been unexplored. Moeller et al. (2020) report gains in performance after manually cleaning training data that were bootstrapped from interlinear glossed texts. Nicolai and Silfverberg (2020) find that exposing an inflection model to its own mistakes leads to better generalizability. Our work explores the impact of different types of noise that occur in the training data in detail.

4 Noise Taxonomy

We develop an automated pipeline to annotate each inflection pair from a taxonomy of noise, primarily relying on rule-based morphological analyzers for each language from Apertium. Here we describe each type of noise in our taxonomy, which we organize into three categories: lemma errors, paired errors, and MSD errors. For a description of how each type of noise is detected, see Section Appendix A.2 in the appendix.

4.1 Lemma Errors

We first describe noise in which a training sample includes lexical items that should not be in the inflection training data at all. Lemma noise arises due to issues inherent to the corpus (e.g. misspellings, etc), or issues in which lexical items that were sampled from the corpus (e.g. punctuation is not




Figure 1: tUMPC noise distribution for all languages.

inflection data). We describe two types of noise that fall into this category.

Lexicon Noise Any word type that is not in the standard vocabulary of a given language is considered lexicon noise. We expect this noise to come from archaisms, borrowings, and biblical references. This means that lexicon noise *could* follow the regular inflections of a language, and in that case may have a low impact on downstream inflection systems. It could also have a positive impact by increasing training data size, likely with higher lemma diversity. However, lexicon noise could also entail archaic or borrowed *inflections*, which would introduce non-existent transformations. In general, lexicon noise could also occur due to language, dialect, or orthography mixing – though we expect this to be somewhat rare in the Bible corpora that we build upon.

POS Noise Not all parts of speech inflect, and this varies by language. However, tUMPC sometimes induces spurious inflection pairs from words that do not inflect, like conjunctions in most languages. Any word of a POS that does not inflect in the given language is thus considered POS noise. This is detected with the Apertium POS, but ultimately these samples will be assigned an MSD from UniMorph with a POS that does inflect. See Table 5 for the POS that inflect according to this study.

While inflection pairs with POS noise may sometimes have real transformations, it is highly likely that these add spurious inflections to the training data. For this reason, we expect POS noise to consistently cause mistakes in each system.

4.2 Paired Errors

Next we describe noise types in which valid lexical items erroneously form an inflection pair. These types of noise occur due to issues in the paradigm induction algorithm that produced the inflection data.

POS Pair Noise Two words from different POS forming an inflection pair constitutes POS pair noise. This can occur when tUMPC puts two completely unrelated words in the same paradigm, but this could also arise due to valid derivations. The latter case could be considered a morphological transformation that our model should learn. There is debate around whether or not there is a clearly defined distinction between inflection and derivation (Haspelmath, 2023), but since our goal is to evaluate on SIGMORPHON shared task data that does not include derivation, we consider this noise in our pipeline. POS pair noise should have a similar impact as POS noise, except that due to derivations we expect several of the induced transformations to appear much more commonly in languages where productive derivation is pervasive.

Paradigm Noise Pairs of forms that do not belong to the same paradigm but share a POS constitute paradigm noise. An example in English is *warp* → *wraps*, where *wraps* comes from a different paradigm than *warp*. Because paradigm noise should contain target words expressing a valid inflection, it is possible that it will not have much negative impact on an inflection system’s decoder. However, since paradigm noise also contains source forms that the target is not actually inflected from, the full transformation from source to target has the potential to be spurious. For instance, in the *warp* → *wraps* example above, consider the apparent *a, r* metathesis. This could cause models relying on a bias towards transduction to struggle more with paradigm noise than models relying predominantly on a decoder.

4.3 MSD Errors

Finally we describe noise consisting of assigning the wrong MSD to an inflection. We have only a single noise type here. It occurs due to issues in the slot alignment algorithm.

Slot Noise When the target word in an inflection pair has an incorrect MSD, we mark this as slot noise. For example, consider an inflection *cry* → *cried*, which is a valid inflection pair, but if, for example, *cried* is incorrectly marked as the third person present, this would be slot noise. Our training setup makes use of *only* the target slot – if only the source form has an incorrect slot then a pair is not marked as slot noise. This error results in a mismatch of MSDs and inflectional transformations. So, for some paradigms, the output form for

an inflection can be thought of as swapped with another output form in the same paradigm. If these two forms are different, this is likely to confuse the inflection system by presenting counter-evidence to the correct inflection. Additionally, it is possible that there are few to no correctly tagged instances of rare or irregular inflection classes, causing a system to confidently learn that, e.g., the past tense inflection is the third person present tense.

4.4 Analysis

Figure 1 presents the noise distribution in the training data for each language according to our annotation pipeline. In German and Russian, there is more correct data than noise, but in Icelandic and Swedish, there is more noise than correct data. We can also consider noise by its source of error in tUMPC: either resulting from an error in slot alignment, i.e. slot noise; an error in the corpus, i.e. lexicon noise; or a paradigm induction error, i.e. all other noise. Slot alignment issues are the most common source of noise, though, for German and Swedish, there are nearly as many paradigm induction errors. More specifically, POS pair noise is the most common noise type after slot noise, and lexicon noise is the least common, most of which occurs in Swedish.

5 Experiments

5.1 Models

We compare four neural inflection generation systems. We implement a bidirectional LSTM with attention (Kann and Schütze, 2016, LSTM), a Transformer (Wu et al., 2021, Trm), a pointer generator LSTM (Sharma et al., 2018, PtrGen), and we use the Dynet implementation of (Makarov and Clematide, 2018, M&C): a transducer optimized with minimum risk training. For all other models, our implementation is based on yoyodyne,² which is built on pytorch (Paszke et al., 2019). All models follow the hyperparameters reported in the original papers, with minor increases in epochs to ensure that they converge. Explicit hyperparameters are listed in the appendix in Table 6.

This gives us two classes of models: general encoder-decoders (LSTM and Trm), which may struggle for low-resource scenarios and under a lack of sample diversity; and transducer-like models with a bias towards copying from the lemma (PtrGen and M&C), which are known to perform

²<https://github.com/CUNY-CL/yoyodyne>

Lang	M&C	LSTM	PtrGen	Trm
Deu	28.38	23.14	31.34	21.42
Isl	21.90	16.60	19.03	16.27
Rus	38.92	31.20	33.07	32.67
Swe	35.88	22.34	26.05	25.68
Avg	31.27	23.32	27.37	24.01

Table 2: Accuracy for the tUMPC training data.

Lang	M&C	LSTM	PtrGen	Trm
Deu	40.12	45.98	42.06	34.75
Isl	31.28	28.64	30.00	32.00
Rus	61.80	65.58	61.96	67.23
Swe	54.20	49.32	47.57	53.43
Avg	46.85	47.38	45.40	46.85

Table 3: Accuracy for the UniMorph training data.

better in low-resource scenarios by relying on modeling character transduction. For all results that we report, we train five of each model on the same five random seeds and report the mean.

5.2 Evaluation

We evaluate on development sets from the SIGMORPHON 2017 shared task (Cotterell et al., 2017). During training, we reuse the target MSDs found through the mapping described in §2, which match the SIGMORPHON target MSDs. Though our training setup considers *reinflection*, we evaluate on inflection from a lemma. This small mismatch in task has a minor negative impact on accuracy (Cotterell et al., 2016).

5.3 Experiment 1: Training on Noisy Data

We first benchmark each model on all four languages when trained on the full noisy dataset.

tUMPC In Table 2 we present the results when models are trained on all data that we were able to classify in our annotation pipeline, cf. Table 1. M&C performs best on average, with PtrGen performing second best. Notably, both of these models are designed with an inductive bias towards copying from the lemma. LSTM and Trm, which both rely on an encoder-decoder that generates from the vocabulary at every time-step, perform similarly to one another. The largest training dataset, Russian, is the most accurate for every model, and Icelandic, the smallest dataset, is the least accurate.

UniMorph Table 3 shows a large increase in accuracy for every language and model when compared to tUMPC. This indicates that the tUMPC

data lacks diversity. Accuracy is similar for all models on average, but LSTM’s accuracy is highest, and PtrGen’s is lowest. In a second experiment, we sample UniMorph pairs from the tUMPC word-length distribution in order to control for the fact that many UniMorph words tend to be uncharacteristically long. This lowers the type frequency of the dataset compared to the original UniMorph sampling — which we interpret as reducing the diversity — and results in a very small increase in performance compared to training on the tUMPC dataset. For results on UniMorph, we focus on the initial, more diverse dataset.

5.4 Experiment 2: Quantity of Noise

We investigate how introducing randomly sampled noise into a training dataset affects model performance as the amount of noise increases. This characterizes the behavior of models as each dataset becomes noisier, and it also shows us at which quantities noise becomes most problematic. We first train models on data comprising only samples that were annotated as correct in order to benchmark performance in the absence of noise. We then partition all of the noisy samples into ten equally sized splits. The x-axis in Figure 2 represents the number of partitions that have been added to a given dataset, so, at 10, we get the results in Tables 2 and 3, and, at 0, we have the results when all noise is removed. Notice that the amount of noisy data in a given partition depends on the language, and represents one tenth of the total noise — we focus our analysis on trends more so than performance at any particular point. Here we analyze the dotted lines, the solid lines represent the results of Experiment 4 on CMLM.

tUMPC Every architecture is negatively impacted by noise in German, with a huge negative impact on M&C. PtrGen behaves sporadically, increasing performance when the last few noise partitions are introduced, and LSTM suffers the least from noise. In Swedish, Icelandic, and Russian, model performance tends to increase as noise is introduced. This is consistent in Trm and LSTM and not always true for M&C and PtrGen.

The negative impact of noise on M&C in German, and the somewhat sporadic behavior of PtrGen indicate that the classic encoder-decoder may be more robust to noise. However, the ranking of models by accuracy does not tend to change over different amounts of noise, demonstrating that they un-




Figure 2: Change in accuracy as each dataset is augmented with noise. One tenth of the noisy data for a given dataset is added to the dataset at each point, until all noise is added at point 10.

derperform on tUMPC. We see consistent increases in performance in Swedish for all models. One explanation for this could be that the large amount of lexicon noise in Swedish tends to contain real inflections that our models learn from. We additionally see several other cases where LSTM and Trm increase in accuracy as more noise is added. This may simply be because they suffer from a lack of sample diversity and more data helps even though it is noisy.

UniMorph In the UniMorph dataset, M&C performance steadily, and sometimes drastically, decreases as noise is introduced for every language. In German, PtrGen behaves sporadically again, and both Trm and LSTM decrease in performance as noise is added, with LSTM decreasing very slowly. The same is true for Swedish and Icelandic, where Trm and LSTM are less impacted by noise than the copy models. All models besides M&C seem relatively robust to noise in Russian.

This suggests that LSTM may be the most robust to noise in our data on average, and that the models with an inductive bias towards copying from the lemma are less robust. In all languages but Russian, there is a distinct downward trend as noise is added for every architecture. Russian is the largest dataset and has more correct than noisy data. This suggests that, with sufficient correct training data, noise is less of a problem. Additionally, a large portion of Russian noise are slot and paradigm noise, which could feasibly have a less negative impact on learning. This also shows that filtering out noisy samples is most useful for M&C, and the negative impact of noise is more severe for models trained on higher sample diversity.

5.5 Experiment 3: Type of Noise

We investigate the impact of each noise type in our annotation schema to see if certain noise types are more important than others and if they affect architectures differently. We produce k training data sets, where each set comprises all of the correct data and all of the samples that have been annotated with one particular annotation so that we can measure that annotation’s impact in isolation. Thus, the impact of each annotation is both a function of the errors entailed by it and the frequency of samples with that annotation. Because noise types can co-occur on a single training sample, we consider the unique combination of noise types as a single annotation. Figure 3 presents the percent change in accuracy over training on only the correct data for each dataset, averaged over all languages. To represent the effect of data size, we add the red line to track the average training set size when a given annotation is included.

tUMPC M&C, the best performing model, is most negatively impacted by any single noise type in both the tUMPC and the UniMorph dataset, which supports the findings of the previous experiments. POS noise, especially in combination with slot noise has the largest negative impact. One explanation for this is that M&C is trained to learn a policy over edits on the input, framing inflection as a true string transduction task. Since POS noise is likely to add inflections that are not actually meaningful, this may result in an edit policy that generates made-up words.

Every model *besides* M&C gets better when the slot annotation is included in the tUMPC dataset.

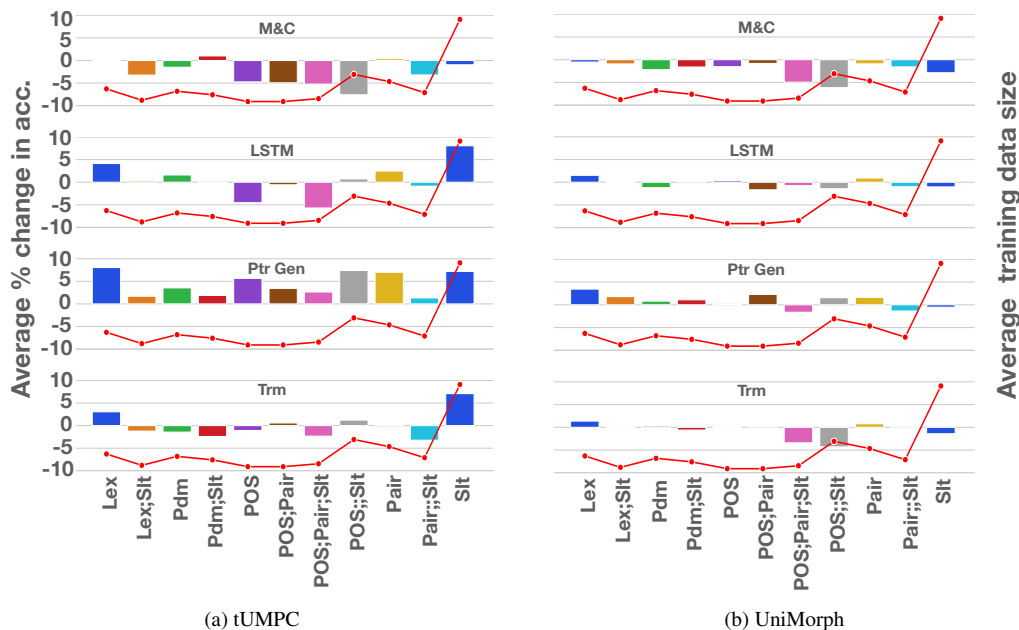


Figure 3: Impact of noise by type. Each bar shows the % increase in accuracy when samples with the corresponding annotation are added to a training set comprising only correct data. The red line represents the training data size after adding those samples.

This may be due to the fact that there are a huge number of slot annotations to learn from, meaning each model is trained on a larger training set. Additionally, M&C may be less sensitive to slot errors as it has a bias towards copying from the lemma – perhaps the best M&C model trained on tUMPC ignores some parts of the MSD. However, PtrGen, which also has a bias towards copying from the lemma, increases in accuracy when *any* noise, especially lexicon noise, is added to the data.

UniMorph M&C behaves similarly here, except the slot annotations have a negative impact. PtrGen still learns from some noise types, but is generally less affected by any single noise type. LSTM is almost completely unaffected by any single noise type, which corroborates Experiment 2. Trm behaves similarly to M&C, though the impact on accuracy is smaller. This indicates that with sufficiently diverse training data, LSTM is more robust to noise than other models, and that Trm learns an inductive bias that resembles M&C.

5.6 Experiment 4: CMLM Pretraining

We experiment with a simple character masked language modeling (CMLM) pretraining objective as a method for improving noise-robustness of each model. Auto-encoding without masking has been shown to impart a helpful inductive bias towards copying from the lemma in low-resource morpho-




Figure 4: % change in accuracy from pretraining.

logical inflection (Kann and Schütze, 2017). Masking can be thought of as additionally adding a de-noising objective to auto-encoding, which may contribute to more robust learning from noisy training samples. We experiment with CMLM for every model and analyze its effect on noisy training.

Implementation We follow the BERT masked language modeling objective (Devlin et al., 2019) with only minor adjustments. First, we mask characters rather than subwords. Second, our sequences will be shorter on average than BERT, so we increase the mask probability from 0.15 to 0.2. In practice, we follow RoBERTa (Liu et al., 2019) and generate a new mask for each sample dynamically at each epoch. Every unique type in a given training set, without their MSDs, comprise the pre-training dataset. We do not include additional types in order to test whether CMLM itself is effective, rather than the addition of data. Each model is

otherwise trained in the exact same way as fine-tuning with the same hyperparameters other than number of epochs. The only exception is that we train `PtrGen` with a warm-up scheduler in order to avoid over-fitting to copying the lemma. Exact hyperparameters are in the appendix in Table 7.

Results Figure 4 shows the average percent change in accuracy for each model in both datasets when adding the pretraining objective. We find that it has little effect for `PtrGen` on average, and actually has a negative impact on `M&C`. We attribute this to the fact that these models have an inductive copy bias, which may be a large part of what CMLM adds to the models. `M&C` might additionally be negatively impacted if the imitation learning objective overfits to the edits during pretraining. More experimentation with pretraining objectives are needed to understand this negative result, however. Our focus is on the impact of pretraining on noise robustness.

`LSTM` is barely affected by pretraining on average. `Trm`, however, increases in accuracy in both datasets when pretrained. It additionally increases in accuracy when training on correct data only, suggesting that CMLM may be generally beneficial to `Trm`, not just in noisy settings. We note, however, that the increase in accuracy is greater when noise is included in both datasets.

CMLM by Noise Quantity We also reproduce Experiment 2 in order to look more closely at the impact of pretraining. We focus on the solid lines in Figure 2, and compare them to the dotted lines that represent the models that were not pretrained.

tUMPC For `M&C`, we see the negative impact of pretraining is especially strong in German and Swedish, and pretraining has little effect in other languages. Pretraining benefits all other architectures in Swedish and Icelandic, especially as noise is added. In German and Russian, pretraining is not beneficial to `PtrGen`, but still helps `Trm` and, at small amounts of noise, `LSTM`. Pretraining is also beneficial when there is no noise in the dataset, but to a lesser extent and not in every language.

UniMorph `M&C` shows performance increases in German under small amounts of noise only. But in every other language pretraining still impacts `M&C` negatively. Similar to the tUMPC data, pretraining benefits all other models in Icelandic and Swedish. Even without noise, `LSTM` and `Trm` increase in performance in Icelandic and to some extent in

Swedish and Russian. In German, `Trm` increases massively in accuracy from pretraining, including when there is no noise. All other architectures also benefit from pretraining in German, but largely only at small amounts of noise. Still, pretrained `LSTM` is the best overall model for German in the UniMorph dataset. In Russian, `Trm` is the only pretrained architecture that performs better as noise is introduced. Pretraining benefits `LSTM` in Russian noise when there is no noise, however.

`Trm` is the only architecture for which CMLM pretraining helps in every language and dataset. Overall, pretraining is also clearly beneficial to `LSTM`, and sometimes to `PtrGen`. Pretraining is particularly effective for these three architectures in Swedish and Icelandic, especially as noise is introduced to the datasets. Combined with the fact that these two languages have more noisy than correct data, this implies that pretraining *is* effective for noise-robust learning. Still, in many cases pretraining leads to a gain in accuracy when there is no noise in the dataset, implying that this learning strategy is generally beneficial, especially for `Trm`.

6 Discussion

We find that low sample diversity has a strong impact on performance of all models. The tUMPC training setup favors architectures with a copy bias, and demonstrates that models can learn from noisy training samples when the dataset is not diverse. We find that the low `LSTM` and `Trm` performance is largely explained by low sample diversity. On the other hand, they seem to be more robust to noisy data, particularly `LSTM`, which has stable performance as noise is added to the UniMorph dataset. Pretraining with CMLM leads to further gains in performance for `LSTM`, but the largest gains from CMLM are for `Trm`. On average, `Trm` pretrained with CMLM is the best performing model under noise, when we have sufficient sample diversity. When we look at specific noise in the data, we find that slot alignment issues in tUMPC tend to have low impact on every model. This could be for several reasons: slot noise should be irrelevant under high amounts of syncretism, which is abundant in German, Icelandic, and Swedish. Additionally, as models become more corrupted by noise, they may rely on a bias towards copying from the source form – which is unaffected by the slot.

We find that certain noise that results from errors in paradigm induction are particularly impactful for

M&C. This is especially true for POS noise, which may motivate better POS induction in unsupervised morphology systems. Under greater sample diversity, Trm is similarly impacted by paradigm induction errors. This suggests that Trm begins to learn an inductive bias similar to M&C on noisy data. The addition of almost any single noise annotation leads to an increase in accuracy in PtrGen , and removing noise from the training data often impacts PtrGen negatively. However, Experiment 2 suggests that PtrGen is not as robust to increasing amounts of noise as other models. A lot of reduction in accuracy may come from combinations of different noise types, which is not captured by Experiment 3. Future work could investigate noise distributions by type with particular focus on PtrGen behavior. We additionally find that M&C is negatively impacted by CMLM pretraining. We believe this may be due to overfitting its copy bias and learning spurious transductions from the masking objective. However, future work should consider alternate pretraining strategies for M&C.

Overall this implies that, although copy models are preferred for training on low sample diversity, classic encoder-decoders are a good choice for noisy datasets with more diversity. Our results indicate that POS and paradigm induction components are more important for training data quality than slot alignment in unsupervised systems and that bootstrapping inflection pairs should prioritize lemma diversity, even if it may induce noise.

7 Conclusion

We have investigated the impact of noise on state-of-the-art neural morphological inflection models. We find that the noise that arises in an unsupervised system for bootstrapping inflection pairs is frequently related to slot alignment errors, but that those also have less impact on the models. We have also compared two inflection architectures with a copy bias to two typical encoder-decoder models. We find that, though copy bias is helpful under low sample diversity, the encoder-decoders are more robust to noise. Finally, we find that a simple masked pretraining objective makes encoder-decoders, and especially Transformers, more accurate under noise.

8 Limitations

The largest limitation of this study is that our annotation pipeline is automated. This makes it possible

that there are errors in the noise annotations that we base our analysis on. Additionally, since we capture a naturally occurring noise distribution, our findings are coupled to the datasets we study here. Our findings may not generalize to distributions of noise in other datasets.

Acknowledgments

We thank the anonymous reviewers for their helpful feedback.

References

- Roe 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.
- Khuyagbaatar Batsuren, Omer Goldman, Salam Khalifa, Nizar Habash, Witold Kieraś, Gábor Bella, Brian Leonard, Garrett Nicolai, Kyle Gorman, Yustinus Ghanggo Ate, et al. 2022. Unimorph 4.0: Universal morphology. *arXiv preprint arXiv:2205.03608*.
- Ryan Cotterell, Christo Kirov, John Sylak-Glassman, Géraldine Walther, Ekaterina Vylomova, Patrick Xia, Manaal Faruqui, Sandra Kübler, David Yarowsky, Jason Eisner, et al. 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.
- Ryan Cotterell, Christo Kirov, John Sylak-Glassman, David Yarowsky, Jason Eisner, and Mans Hulden. 2016. 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.
- 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.
- 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.

- 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.
- Mikel L Forcada, Mireia Ginestí-Rosell, Jacob Nordfalk, Jim O’Regan, Sergio Ortiz-Rojas, Juan Antonio Pérez-Ortiz, Felipe Sánchez-Martínez, Gema Ramírez-Sánchez, and Francis M Tyers. 2011. AperiTium: a free/open-source platform for rule-based machine translation. *Machine translation*, 25(2):127–144.
- Omer Goldman, David Guriel, and Reut Tsarfaty. 2022. (un)solving morphological inflection: Lemma overlap artificially inflates models’ performance. In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, pages 864–870, Dublin, Ireland. Association for Computational Linguistics.
- 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.
- Michael Gutmann and Aapo Hyvärinen. 2010. Noise-contrastive estimation: A new estimation principle for unnormalized statistical models. In *Proceedings of the thirteenth international conference on artificial intelligence and statistics*, pages 297–304. JMLR Workshop and Conference Proceedings.
- Martin Haspelmath. 2023. Inflection and derivation as traditional comparative concepts. *Linguistics (to appear)*.
- Huiming Jin, Liwei Cai, Yihui Peng, Chen Xia, Arya D 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.
- 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. 2016. **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.
- Richard M Karp. 1980. An algorithm to solve the $m \times n$ assignment problem in expected time $O(mn \log n)$. *Networks*, 10(2):143–152.
- Huda Khayrallah and Philipp Koehn. 2018. **On the impact of various types of noise on neural machine translation**. In *Proceedings of the 2nd Workshop on Neural Machine Translation and Generation*, pages 74–83, Melbourne, Australia. Association for Computational Linguistics.
- Jordan Kodner, Salam Khalifa, Khuyagbaatar Batsuren, Hossep Dolatian, Ryan Cotterell, Faruk Akkus, Antonios Anastasopoulos, Taras Andrushko, Aryaman Arora, Nona Atanalov, Gábor Bella, Elena Budianskaya, Yustinus Ghanggo Ate, Omer Goldman, David Guriel, Simon Guriel, Silvia Guriel-Agiashvili, Witold Kieraś, Andrew Krizhanovsky, Natalia Krizhanovsky, Igor Marchenko, Magdalena Markowska, Polina Mashkovtseva, Maria Nepomniashchaya, Daria Rodionova, Karina Scheifer, Alexandra Sorova, Anastasia Yemelina, Jeremiah Young, and Ekaterina Vylomova. 2022. **SIGMORPHON–UniMorph 2022 shared task 0: Generalization and typologically diverse morphological inflection**. In *Proceedings of the 19th SIGMORPHON Workshop on Computational Research in Phonetics, Phonology, and Morphology*, pages 176–203, Seattle, Washington. Association for Computational Linguistics.
- Ling Liu and Mans Hulden. 2022. **Can a transformer pass the wug test? tuning copying bias in neural morphological inflection models**. In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, pages 739–749, Dublin, Ireland. Association for Computational Linguistics.
- Yinhan Liu, Myle Ott, Naman Goyal, Jingfei Du, Mandar Joshi, Danqi Chen, Omer Levy, Mike Lewis, Luke Zettlemoyer, and Veselin Stoyanov. 2019. Roberta: A robustly optimized bert pretraining approach. *arXiv preprint arXiv:1907.11692*.
- Peter Makarov and Simon Clematide. 2018. **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, Tatyana 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.

- Arya D. McCarthy, Christo Kirov, Matteo Grella, Amrit Nidhi, Patrick Xia, Kyle Gorman, Ekaterina Vylomova, Sabrina J. Mielke, Garrett Nicolai, Miikka Silfverberg, Timofey Arkhangelskiy, Natalya Krizhanovskiy, Andrew Krizhanovskiy, Elena Klyachko, Alexey Sorokin, John Mansfield, Valts Ernštreits, Yuval Pinter, Cassandra L. Jacobs, Ryan Cotterell, Mans Hulden, and David Yarowsky. 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, 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.
- 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.
- 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.
- Garrett Nicolai and Miikka Silfverberg. 2020. [Noise isn't always negative: Countering exposure bias in sequence-to-sequence inflection models](#). In *Proceedings of the 28th International Conference on Computational Linguistics*, pages 2837–2846, Barcelona, Spain (Online). International Committee on Computational Linguistics.
- Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, Alban Desmaison, Andreas Kopf, Edward Yang, Zachary DeVito, Martin Raison, Alykhan Tejani, Sasank Chilamkurthy, Benoit Steiner, Lu Fang, Junjie Bai, and Soumith Chintala. 2019. [Pytorch: An imperative style, high-performance deep learning library](#). In *Advances in Neural Information Processing Systems 32*, pages 8024–8035. Curran Associates, Inc.
- Giorgio Patrini, Alessandro Rozza, Aditya Krishna Menon, Richard Nock, and Lizhen Qu. 2017. [Making deep neural networks robust to label noise: A loss correction approach](#). In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 1944–1952.
- Umaa Rebbapragada and Carla E Brodley. 2007. [Class noise mitigation through instance weighting](#). In *European conference on machine learning*, pages 708–715. Springer.
- 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.
- Hwanjun Song, Minseok Kim, Dongmin Park, Yooju Shin, and Jae-Gil Lee. 2022. [Learning from noisy labels with deep neural networks: A survey](#). *IEEE Transactions on Neural Networks and Learning Systems*.
- John Sylak-Glassman, Christo Kirov, Matt Post, Roger Que, and David Yarowsky. 2015. [A universal feature schema for rich morphological annotation and fine-grained cross-lingual part-of-speech tagging](#). In *International Workshop on Systems and Frameworks for Computational Morphology*, pages 72–93. Springer.
- Ekaterina Vylomova, Jennifer White, Elizabeth Salesky, Sabrina J. Mielke, Shijie Wu, Edoardo Maria Ponti, Rowan Hall Maudslay, Ran Zmigrod, Josef Valvoda, Svetlana Toldova, Francis Tyers, Elena Klyachko, Ilya Yegorov, Natalia Krizhanovskiy, Paula Czarnowska, Irene Nikkarinen, Andrew Krizhanovskiy, Tiago Pimentel, Lucas Torroba Hennigen, Christo Kirov, Garrett Nicolai, Adina Williams, Antonios Anastasopoulos, Hilaria Cruz, Eleanor Chodroff, Ryan Cotterell, Miikka Silfverberg, and Mans Hulden. 2020. [SIGMORPHON 2020 shared task 0: Typologically diverse morphological inflection](#). In *Proceedings of the 17th SIGMORPHON Workshop on Computational Research in Phonetics, Phonology, and Morphology*, pages 1–39, Online. Association for Computational Linguistics.
- Adam Wiemerslage, Miikka Silfverberg, Changbing Yang, Arya D McCarthy, Garrett Nicolai, Eliana Colunga, and Katharina Kann. 2022. [Morphological processing of low-resource languages: Where we are and what's next](#). In *Findings of the Association for Computational Linguistics: ACL 2022*, pages 988–1007.
- 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.

Nasser Zalmout, Alexander Erdmann, and Nizar Habash. 2018. [Noise-robust morphological disambiguation for dialectal Arabic](#). 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 953–964, New Orleans, Louisiana. Association for Computational Linguistics.

A Appendix

A.1 Slot Mapping Details

We begin with each type processed by tUMPC, which has a slot: an arbitrary identifier for its POS and inflection category, and a (not disambiguated) morphological analysis from Apertium. For example, given the German verb *tragt*, we have a tUMPC slot 2. We additionally have the analysis from Apertium with two possibilities: an imperative plural verb (`<vblex><imp><pl>`) or a second person present indicative plural verb (`<vblex><pri><p2><pl>`).

Each possibility in the Apertium analysis is then mapped to a UniMorph MSD via a mapping we create that translates each tag one at a time. For example, the tag `<vblex>` becomes V, in order to match the UniMorph schema. After some language specific post-processing heuristics, we get a set of UniMorph MSDs from every Apertium analysis.

We can use these mapped analyses to align tUMPC slots with UniMorph MSDs. Consider our example above, *tragt*, where we would end up with two possible MSDs: V;IMP;2;PL and V;IND;PRS;2;PL. This forms a mapping from the slot 2, to both of these MSDS. Due to tUMPC errors, we may also erroneously get a mapping from slot 2 to N;ACC;PL via some other word. This gives us a mapping from three differing UniMorph MSDs to one tUMPC slot. Over all such mappings, this forms a bipartite graph between tUMPC slots and UniMorph MSDs, where the same UniMorph MSD may correspond to multiple tUMPC slots. However, one tUMPC slot represents exactly one MSD. We thus follow [Jin et al. \(2020\)](#) and attain a one-to-one mapping from tUMPC slots to MSDs by finding the matching that maximizes the overlap of word types with an aligned slot and MSD. Like them, we use the algorithm from [Karp \(1980\)](#) to optimize this matching. Finally, the slot for every training sample can be mapped to its MSD according to this matching.

A.2 Noise Detection

Here we describe how our annotation pipeline detects each type of noise. We rely on Apertium for the entire pipeline. Though most noise is found with original Apertium analysis, the analyses that have been mapped to UniMorph MSDs are used for detecting slot noise. .

Language	C	LEX	PDGM	POS	POS Pair	SLOT	Lemma Overlap	MSD Overlap	Tag Overlap
deu	61.37	1.52	0.68	8.07	8.60	19.75	3.84	24.32	83.33
isl	37.95	3.91	2.11	15.59	8.17	32.27	3.24	60.61	95.65
rus	50.77	0.98	6.01	4.21	4.25	33.79	3.31	68.18	100.0
swe	27.90	11.36	5.48	8.39	15.66	31.20	1.94	69.33	93.1

Table 4: Statistics for each language’s training data. The % of samples with a given annotation (left). The % overlap with the evaluation set for lemmas, MSDs, or individual tags (right).

Isl	verb	adjective	noun			
Deu	verb	adjective	noun	prn		
Rus	verb	adjective	noun	prn	numeral	
Swe	verb	adjective	noun			

Table 5: Parts of speech that inflect in our annotation schema.

	M&C	PtrGen	LSTM	Trm
Epochs	60	60	60	800
Batch size	1	32	20	400
Optimizer	AD	adam	AD	adam
LR	1.0	0.001	1.0	0.001
Scheduler	-	-	-	inv. sqr root
Warmup	-	-	-	4000
Total Params	260k	1.1M	370k	7.4M

Table 6: Hyperparameters for each architecture. All other hyperparameters follow their respective publications exactly. AD=ADADELTA.

Lexicon Noise We use three resources to form the lexicon of each language: Apertium, Wikipedia, and a python spellchecker³ based on hunspell⁴. Any word not in any of these three is lexicon noise.

POS Noise Here we produce lists of POS that inflect for each language. Any word whose Apertium analysis does not contain any valid POS according to this list is considered POS noise. The valid POS for each language are listed in Table 5.

POS Pair Here we consider all POS from Apertium analysis of both words in a pair. If they have no POS in common, then it is POS pair noise.

Paradigm For a given inflection pair, if both words have no overlapping lemmas of the same POS (but *do* have a shared POS) according to Apertium, we consider them to be from separate paradigms, and thus paradigm noise.

Slot Slot noise occurs when the slot assigned to a word by tUMPC is not in the set of slots from Apertium analysis. We rely on the version of tUMPC

	M&C	PtrGen	LSTM	Trm
Pretrain epochs	40	40	40	200
Pretrain scheduler	-	inv. sqr root	-	inv. sqr root
Pretrain warmup	-	100	-	1000

Table 7: Hyperparameters for each architecture during pretraining.

and Apertium slots that have been mapped to UniMorph for this part of the annotation. Slot noise considers only the predicted and gold MSD for a slot, so it can co-occur on a single sample with any other noise in the taxonomy.

³<https://github.com/dvwright/phunspell>

⁴<http://hunspell.github.io/>

ACL 2023 Responsible NLP Checklist

A For every submission:

- A1. Did you describe the limitations of your work?
8
- A2. Did you discuss any potential risks of your work?
We do not see any clear ethical risks with the work.
- A3. Do the abstract and introduction summarize the paper’s main claims?
1
- A4. Have you used AI writing assistants when working on this paper?
Left blank.

B Did you use or create scientific artifacts?

2, 4, 5.1

- B1. Did you cite the creators of artifacts you used?
1,2,4,5.1
- B2. Did you discuss the license or terms for use and / or distribution of any artifacts?
We use standard tools and datasets for research and appropriately cite the publications they were released with.
- B3. Did you discuss if your use of existing artifact(s) was consistent with their intended use, provided that it was specified? For the artifacts you create, do you specify intended use and whether that is compatible with the original access conditions (in particular, derivatives of data accessed for research purposes should not be used outside of research contexts)?
The only resource for which this is not extremely obvious is Apertium morphological analyzers. They are intended for research, but typically used in an MT pipeline. We discuss that we use them just for morphological analysis in section 2.
- B4. Did you discuss the steps taken to check whether the data that was collected / used contains any information that names or uniquely identifies individual people or offensive content, and the steps taken to protect / anonymize it?
We use well-established datasets (the bible and UniMorph).
- B5. Did you provide documentation of the artifacts, e.g., coverage of domains, languages, and linguistic phenomena, demographic groups represented, etc.?
Section 4, Figure 1, Table 4 in the appendix
- B6. Did you report relevant statistics like the number of examples, details of train / test / dev splits, etc. for the data that you used / created? Even for commonly-used benchmark datasets, include the number of examples in train / validation / test splits, as these provide necessary context for a reader to understand experimental results. For example, small differences in accuracy on large test sets may be significant, while on small test sets they may not be.
Table 1, Figure 1, Table 4 in the appendix

The Responsible NLP Checklist used at ACL 2023 is adopted from NAACL 2022, with the addition of a [question on AI writing assistance](#).

C Did you run computational experiments?

5

- C1. Did you report the number of parameters in the models used, the total computational budget (e.g., GPU hours), and computing infrastructure used?
Tables 5 & 6. We do not discuss GPU hours/infrastructure. These models are very small compared to most research – and are reimplementations of previous work. Many of our models are trained on CPU. Even so, we run thousands of experiments on varying data sizes, and run them on a cluster for which a job scheduler allocates varying models of CPU and GPU as they become available. It would be very cumbersome to list all possible hardware any given experiment could have been trained on, especially when all of our architectures can feasibly be trained on a laptop in a few hours.
- C2. Did you discuss the experimental setup, including hyperparameter search and best-found hyperparameter values?
5.1, Tables 5&6
- C3. Did you report descriptive statistics about your results (e.g., error bars around results, summary statistics from sets of experiments), and is it transparent whether you are reporting the max, mean, etc. or just a single run?
5.3, 5.4, 5.5, 5.6
- C4. If you used existing packages (e.g., for preprocessing, for normalization, or for evaluation), did you report the implementation, model, and parameter settings used (e.g., NLTK, Spacy, ROUGE, etc.)?
5.1

D Did you use human annotators (e.g., crowdworkers) or research with human participants?

Left blank.

- D1. Did you report the full text of instructions given to participants, including e.g., screenshots, disclaimers of any risks to participants or annotators, etc.?
Not applicable. Left blank.
- D2. Did you report information about how you recruited (e.g., crowdsourcing platform, students) and paid participants, and discuss if such payment is adequate given the participants' demographic (e.g., country of residence)?
Not applicable. Left blank.
- D3. Did you discuss whether and how consent was obtained from people whose data you're using/curating? For example, if you collected data via crowdsourcing, did your instructions to crowdworkers explain how the data would be used?
Not applicable. Left blank.
- D4. Was the data collection protocol approved (or determined exempt) by an ethics review board?
Not applicable. Left blank.
- D5. Did you report the basic demographic and geographic characteristics of the annotator population that is the source of the data?
Not applicable. Left blank.