

# Character-aware Transformers Learn an Irregular Morphological Pattern Yet None Generalize Like Humans

Akhilesh Kakolu Ramarao<sup>1</sup>, Kevin Tang<sup>1,3</sup>, Dinah Baer-Henney<sup>2</sup>

<sup>1</sup>Department of English Language and Linguistics, Heinrich Heine University Düsseldorf

<sup>2</sup>Institut für Germanistik, Philologische Fakultät, Ruhr-Universität Bochum

<sup>3</sup>Department of Linguistics, College of Liberal Arts and Sciences, University of Florida  
{akhilesh.kakolu.ramarao, kevin.tang}@uni-duesseldorf.de, dinah.baer-henney@rub.de

## Abstract

Whether neural networks can serve as cognitive models of morphological learning remains an open question. Recent work has shown that encoder-decoder models can acquire irregular patterns, but evidence that they generalize these patterns like humans is mixed. We investigate this using the Spanish *L-shaped morpheme*, where only the first-person singular indicative (e.g., *pongo* ‘I put’) shares its stem with all subjunctive forms (e.g., *ponga*, *pongas*) despite lacking apparent phonological, semantic, or syntactic motivation. We compare five encoder-decoder transformers varying along two dimensions: sequential vs. position-invariant positional encoding, and atomic vs. decomposed tag representations. Positional encoding proves decisive: position-invariant models recover the correct L-shaped paradigm clustering even when L-shaped verbs are scarce in training, whereas sequential positional encoding models only partially capture the pattern. Yet none of the models productively generalize this pattern to novel forms. Position-invariant models generalize the L-shaped stem across subjunctive cells but fail to extend it to the first-person singular indicative, producing a mood-based generalization rather than the L-shaped morphomic pattern. Humans do the opposite, generalizing preferentially to the first-person singular indicative over subjunctive forms. None of the models reproduce the human pattern, highlighting the gap between statistical pattern reproduction and morphological abstraction.

**Keywords:** Cognitive Modeling, Morphological Inflection, Morphemes, Transformers, Positional Encoding

## 1. Introduction

Inflectional paradigms are not merely lists of forms but constitute a central organizing principle of morphological knowledge (Stump, 2001). This claim, that the paradigm rather than the morpheme is the primary object of morphological analysis (Matthews, 1972), has profound implications for how we understand the mental lexicon. If paradigms are cognitively primitive, then speakers do not assemble inflected words by concatenating stored morphemes; instead, they exploit implicational relations among whole forms, inferring unattested forms from partial paradigmatic knowledge (Ackerman et al., 2009). Yet how far does this organization extend to irregular verbs? Can paradigmatic patterns that lack phonological, semantic, or syntactic motivation be learned and productively generalized? In this paper, we use computational modeling to probe these questions, training neural network models that differ in architectural choices on inflectional paradigms, comparing how these choices affect pattern acquisition, and evaluating their generalization behavior against human speakers tested on the same items.<sup>1</sup>

We focus on *morphomic* patterns (Spencer and

Aronoff, 1994), which provide the strongest test of paradigmatic learning because they lack any apparent phonological, semantic, or syntactic motivation. They challenge theories that require all morphological structure to be grounded in meaning or function, representing instead autonomous organization within the morphological system itself. Such patterns are well attested across Romance languages (Maiden, 2011, 2018, 2021), making them ideal candidates for investigating whether computational models can acquire purely morphological generalizations.

In Spanish, the *L-shaped morpheme* is one such pattern: the first-person singular indicative (*pongo*) shares its stem with every subjunctive form (*ponga*, *pongas*, ...), while all other indicative forms use the regular stem (*pones*, *pone*, ...). Table 1 illustrates this pattern for the verb *poner* ‘to put’. The defining property is the grouping of 1SG.IND patterns with all subjunctive forms rather than with its own mood category.

Whether morphomic patterns constitute a psychological reality for speakers remains contested. Nevins et al. (2015) argue against the cognitive status of the L-shaped morpheme with evidence from wug-test experiments on Spanish, Italian, and Portuguese, but subsequent work challenges this view. Cappellaro et al. (2024) provide experimental evidence that Italian speakers exhibit sensitivity to L-shaped patterns in processing and production

<sup>1</sup>Code and Data: [https://github.com/hnuslamlab/feature\\_informed\\_transformer\\_cmcl2026](https://github.com/hnuslamlab/feature_informed_transformer_cmcl2026)

	Indicative	Subjunctive
1SG	<b>pong</b> -o	<b>pong</b> -a
2SG	pon-es	<b>pong</b> -as
3SG	pon-e	<b>pong</b> -a
1PL	pon-emos	<b>pong</b> -amos
2PL	pon-éis	<b>pong</b> -áis
3PL	pon-en	<b>pong</b> -an

Table 1: Present tense paradigm of *poner* ‘to put’. Hyphens mark stem-suffix boundaries. Bold marks the L-shaped stem.

tasks. Similarly, O’Neill (2025) offer corpus-based evidence that gaps in Portuguese defective verbs cluster specifically in L-pattern cells rather than being randomly distributed across the paradigm. Computational modeling offers a complementary approach to this debate. Kakolu Ramarao et al. (2025) train neural sequence-to-sequence models on Spanish verbal inflection, providing form-morphological feature pairs as input and generating inflected forms as output. They demonstrate that these models can reproduce L-shaped stem alternations, with frequency of L-shaped verbs playing a crucial role. However, their study employed a single architecture. We extend this line of work by comparing five model architectures to ask whether design choices affect the acquisition of morphomic patterns, and then evaluating their generalization behavior against human experimental data.

We focus on two architectural dimensions that could plausibly affect the acquisition of morphomic patterns: sequential vs. position-invariant positional encoding, and atomic vs. decomposed morphosyntactic tag representations. Character-level encoder-decoder transformers have become the dominant approach for morphological inflection (Wu et al., 2021), but the Unimorph schema (Sylak-Glassman, 2016), which provides feature specifications for inflectional morphology, treats morphosyntactic features as unordered bundles:  $V; IND; PRS; 1; SG$  and  $V; PRS; IND; SG; 1$  specify the same cell. This suggests that a tag should carry the same meaning regardless of its position in the sequence. We therefore compare architectures that assign tags sequential positions against ones that give all tags a fixed positional encoding, making their representation position-invariant. On the representation side, treating tags as atomic tokens requires the model to learn feature structure implicitly, whereas decomposed representations such as one-hot vectors or linguistically-informed feature geometries (Harley and Ritter, 2002) encode this structure directly. To test whether these choices affect morphomic pattern acquisition, we implement five encoder-decoder transformer variants that differ along these two dimensions. Because L-shaped verbs constitute a small minority of the Spanish lexi-

con, we additionally vary the proportion of L-shaped verbs in training across three frequency conditions. This allows us to assess whether architectural inductive biases compensate for the natural scarcity of L-shaped verbs or only matter when such data is abundant, and whether increased exposure to L-shaped patterns during training improves generalization to novel forms. We address the following research questions:

- **RQ1** How do positional encoding and morphosyntactic tag representation affect the acquisition of L-shaped morphomic patterns?
- **RQ2** Do models acquire the L-shaped paradigm as a structural pattern?
- **RQ3** Do models productively generalize the L-shaped pattern to nonce verbs, and how does their behavior compare to that of human speakers?

## 2. Background

The SIGMORPHON shared task series has driven rapid progress in neural morphological inflection since 2016 (Cotterell et al., 2016), with successive iterations expanding to typologically diverse languages (Cotterell et al., 2017, 2018; Vylomova et al., 2020). Character-level encoder-decoder transformers have emerged as the dominant approach (Wu et al., 2021), and design choices such as separate encoders for characters and morphosyntactic tags have produced better results (Anastasopoulos and Neubig, 2019; Peters and Martins, 2019). Yet despite high reported accuracies, productive generalization remains elusive: when models are evaluated on unseen lemmas rather than forms from familiar paradigms, performance drops sharply (Goldman et al., 2022; Kodner et al., 2023), and predicting inflections for unseen feature combinations proves particularly challenging (Kodner and Khalifa, 2022). This gap between memorization and productive generalization suggests that current transformer-based architectures may lack the inductive biases needed to acquire the abstract representations that morphological knowledge demands.

Whether neural networks can serve as cognitive models of morphological learning has been debated since connectionist models of the English past tense (Rumelhart et al., 1986; Kirov and Cotterell, 2018). But so far, encoder-decoder models correlate poorly with human wug-test data (Corkery et al., 2019; McCurdy et al., 2020), and transformers struggle to generalize in a wug-test setting (Liu and Hulden, 2022; Weissweiler et al., 2023). This gap between model and human behavior has motivated benchmarks, such as the BabyLM Challenge,

which trains models on developmentally plausible datasets to probe language acquisition (Warstadt et al., 2023; Hu et al., 2024).

The transformer architecture has no built-in sense of order, i.e., without positional encodings, it cannot, for example, distinguish *pongo* from *gopon* (Dufter et al., 2022). The original transformer introduced sinusoidal positional encodings (Vaswani et al., 2017). Since then, numerous alternatives have been proposed, including relative positional encoding (Shaw et al., 2018), learned absolute embeddings (Devlin et al., 2019), and rotary positional encoding (Su et al., 2024).

However, the necessity of positional encoding varies across tasks: Wang and Chen (2020) show that BERT without positional encoding functions essentially as a bag-of-words model in understanding tasks (classification, QA), while Haviv et al. (2022) find that causal language models remain competitive even without explicit positional encoding in generation tasks, suggesting some tasks are less sensitive to position.

How morphosyntactic tags are represented has received comparatively less attention. Since Kann and Schütze (2016), most systems concatenate morphological features as individual tokens to the character sequence and encode both through a single encoder. Other approaches separate the two inputs: Aharoni and Goldberg (2017) concatenate learned feature embeddings into a single vector that conditions the decoder at every step, Peters and Martins (2019) learn separate attention distributions for lemma characters and tags, and Anastasopoulos and Neubig (2019) use separate encoders for lemma characters and tags. However, decomposing feature bundles into structured vectors, such as one-hot encodings or binary features from feature geometry (Harley and Ritter, 2002), remains unexplored in the literature.

## 3. Methodology

### 3.1. Data

#### 3.1.1. Real verbs

We use Spanish verbal paradigms in IPA transcription from Kakolu Ramarao et al. (2025). Each verb has a 12-cell present-tense paradigm crossing three persons (1, 2, 3), two numbers (singular, plural), and two moods (indicative, subjunctive). Verbs are classified as *L-shaped* if the 1SG.IND stem matches all subjunctive stems but differs from the remaining indicative stems, and as *NL-shaped* otherwise. The full dataset contains 299 L-shaped and 4,859 NL-shaped verbs.

To investigate the effect of input frequency on learning, we create three training conditions by varying the proportion of L-shaped verbs: 10%L-90%NL,

50%L-50%NL, and 90%L-10%NL. The 10%L condition most closely approximates the natural distribution of L-shaped verbs in the Spanish lexicon, where they constitute roughly 7% of verb types. For each condition, we sample 332 lemmas from the full pool and split them into training (70%, 232 lemmas), development (10%, 34 lemmas), and test (20%, 66 lemmas) sets. Following Goldman et al. (2022) and Kodner et al. (2023), we ensure that there is no lemma overlap between the training, development, and test sets, so that models must generalize to unseen lemmas. The L-shaped to NL-shaped ratio is maintained across all three splits.

Each training example consists of two source form-tag pairs and one target tag (see Section 3.2). From a 12-cell paradigm, we generate all 660 unique (unordered source pair, target) triples per lemma and subsample 25% per training seed. We repeat the procedure with three lemma splits and four data subsamples, yielding 12 independently trained models per condition per architecture.

#### 3.1.2. Nonce verbs

To compare model behavior with human generalization, we use behavioral data from the wug test experiment by Nevins et al. (2015) on 107 native Spanish speakers. In a paradigm completion task, participants were presented two forms of 15 nonce verbs exhibiting novel stem-final consonant alternations between fricatives and stops (e.g., *bus-* vs. *but-*, *mif-* vs. *mip-*) and asked to produce a third form.

Participant responses were converted from Spanish orthography to IPA transcription, and only responses matching one of the expected paradigm cells were retained. For each nonce verb, we evaluate model predictions on the same three paradigm cells tested in the human study (1SG.IND, 2SG.SBJV, and 3SG.SBJV), using the source pairs in both orderings: two source pairs for 1SG.IND and one each for the two subjunctive cells, yielding 8 test items per verb (8 x 15 verbs = 120 test items per model run).

### 3.2. Task

To mirror the above wug test, where participants are given two forms and must produce a third, we frame the task as two-source morphological re-inflection (Kann et al., 2017): given two source form-tag pairs from a verb's paradigm and a target feature bundle, the model must produce the corresponding target form. For example, given (*pongo*, V; IND; PRS; 1; SG) and (*ponga*, V; SBJV; PRS; 1; SG) with target features V; SBJV; PRS; 2; SG, the model should produce *pongas*. Following Wu et al. (2021), we model this

as character-level sequence-to-sequence transduction. During training, models observe all source-target combinations for each lemma; at test time, they must generalize to entirely unseen lemmas.

### 3.3. Model architectures

All five models share a common encoder-decoder transformer backbone with 4 layers, 4 attention heads, embedding dimension  $d = 256$ , and feedforward dimension 1024. The models differ in how the encoder handles morphosyntactic tags along two dimensions: (i) whether tags receive sequential positional encoding or are assigned a fixed position, and (ii) how tags are represented. We organize the models into two classes based on positional encoding: the first class assigns tags sequential positions as it is done with characters, while the second treats tags as position-invariant. The three architectures in the position-invariant class further differ in tag representation.

**Training.** We use the same hyperparameter settings as [Kakolu Ramarao et al. \(2025\)](#). All models are optimized using Adam ([Kingma and Ba, 2015](#)) with a learning rate of 0.001 and a linear warmup schedule over 4,000 steps. We use a batch size of 400, dropout of 0.3, and label smoothing of 0.1. Training proceeds for a maximum of 10,000 update steps. At inference, we use beam search with a beam size of 5. For each of the three frequency conditions (10%L, 50%L, 90%L), we train 12 models per architecture.

#### 3.3.1. Class 1: Uniform positional encoding

**Vanilla.** The baseline concatenates form characters and feature tags into a single flat sequence. Characters are tokenized individually (e.g., *pongo*  $\rightarrow$  p o n g o), while each morphosyntactic feature bundle is a single vocabulary entry (e.g., V; IND; PRS; 1; SG). Both characters and tags are embedded through a shared learned embedding matrix. Both token types receive sequential sinusoidal positional encoding based on their position in the concatenated sequence. The model has no mechanism to distinguish tag tokens from character tokens ([Kakolu Ramarao et al., 2025](#)).

**Character-separated.** Identical to VANILLA, except that feature tags are decomposed into individual characters separated by whitespace (e.g., V; IND; PRS; 1; SG  $\rightarrow$  V IND PRS 1 SG). All tokens share the same embedding table and receive uniform sequential positional encoding.

#### 3.3.2. Class 2: Position-invariant tags

The UniMorph schema ([Sylak-Glassman, 2016](#)) treats morphosyntactic features as unordered bundles, suggesting that sequential positional encoding may be inappropriate for morphosyntactic specifications. The following three architectures assign tags a fixed positional index of 0 while characters retain sequential positional encoding. They differ in how they represent tag content.

**Feature-invariant.** Following [Wu et al. \(2021\)](#), tags are embedded through the same shared matrix as in VANILLA but are assigned a fixed positional index of 0, while characters retain sequential positional encoding.

**Feature-onehot.** Tags are assigned a fixed positional index of 0 and represented not as atomic tokens but as structured feature vectors. Each tag is parsed into its component features and encoded as a one-hot vector over feature categories: mood (IND, SBJV), person (1, 2, 3), and number (SG, PL), yielding a sparse 7-dimensional binary vector. A learned linear projection maps this vector to the same  $d$ -dimensional embedding space used by character tokens. At each position, the model selects between the character embedding and the feature projection based on token type.

**Feature-geometric.** Like FEATURE-ONEHOT, tags are assigned a fixed positional index of 0 and represented as structured vectors. However, instead of one-hot category vectors, tags are encoded as binary feature vectors following the feature-geometric analysis of [Harley and Ritter \(2002\)](#): person is decomposed into [ $\pm$ participant,  $\pm$ author] (1st = [+,+], 2nd = [+,-], 3rd = [-,-]), number into [ $\pm$ plural], and mood into [ $\pm$ indicative], yielding a 4-dimensional binary vector. This vector is projected to the embedding space through a learned linear transformation.

### 3.4. Evaluation

#### 3.4.1. Sequence accuracy

We evaluate overall sequence accuracy as the percentage of predictions that exactly match the target form, and additionally report it separately for L-shaped and NL-shaped verbs.

#### 3.4.2. Stem accuracy by verb type

The distinction between L-shaped and NL-shaped verbs is defined by the stem-final consonant alternations: L-shaped verbs use one stem in the 1SG.IND and all subjunctive cells and a different stem elsewhere. Whether a model has learned the alternation pattern is therefore most directly

assessed at the stem level. We extract stems by stripping the conjugation suffixes (for *-ar*, *-er*, and *-ir* verb classes) from both predicted and target forms. A prediction is stem-correct if the extracted stem matches the target stem. We report stem accuracy separately for L-shaped and NL-shaped verbs.

### 3.4.3. Paradigm shape analysis

Beyond aggregate accuracy, we ask whether models treat the 12 paradigm cells in a way that reflects the L-shaped pattern, i.e., whether they achieve *structural generalization* (Hupkes et al., 2023). If a model has learned the L-shape, its stem predictions should systematically differ between L-cells (1SG.IND + all subjunctive) and the remaining indicative cells. We first report raw mean stem accuracy per paradigm cell across all test lemmas to examine whether models produce different stems across cells.

To quantify whether this per-cell pattern reflects the L-shape, we apply a scoring transformation to stem accuracies aggregated across all test lemmas, regardless of verb type: L-cell accuracies are kept as is, while the remaining indicative cell accuracies are inverted ( $1 - accuracy$ ), since for L-shaped verbs an incorrect stem in a non-L indicative cell corresponds to the model producing the L-shaped stem where the regular stem is expected. After this transformation, a high score indicates that the model produced the L-shaped stem for that cell. If L-cells score high while non-L cells score low, this suggests L-shaped behavior. We then apply  $k$ -means clustering ( $k = 2$ ) to group the 12 cells based on these transformed scores. A model that has acquired the L-shaped paradigm should produce two clusters: one containing 1SG.IND and all subjunctive cells, and another containing the five remaining indicative cells.

### 3.4.4. Nonce verb evaluation

To compare model generalization with human behavior, we evaluate all models on the nonce verb stimuli from the wug test experiment of Nevins et al. (2015). For each of the 15 nonce verbs, the model is tested on the same three paradigm cells as in the human study (1SG.IND, 2SG.SBJV, 3SG.SBJV). We report stem accuracy and per-cell paradigm shape for models alongside the human behavioral data.

## 4. Results

### 4.1. Real verbs

#### 4.1.1. Sequence accuracy

Figure 1 compares sequence-level accuracy across the five architectures and three frequency con-




Figure 1: Overall sequence accuracy for all five architectures across the three frequency conditions.

ditions. The three position-invariant models (Class 2: FEATURE-INVARIANT, FEATURE-ONEHOT, FEATURE-GEOMETRIC) consistently outperform the two sequential-PE models (Class 1: VANILLA, CHARACTER-SEPARATED). In the 10%L condition, where L-shaped examples are scarce, the gap is largest: Class 2 models achieve mean accuracies of 68-70%, whereas VANILLA reaches 59% and CHARACTER-SEPARATED only 53%. As L-shaped training data increases, overall accuracy improves for all architectures and the gap narrows: in the 90%L condition, all five models converge to the 81-83% range. Class 1 models also exhibit higher variance in the 10%L and 50%L conditions, whereas Class 2 models maintain low variability even when L-shaped examples are scarce. This indicates that position-invariant tag encoding provides a stronger and more stable inductive bias when L-shaped examples are scarce, but with sufficient data even simpler architectures achieve comparable performance.

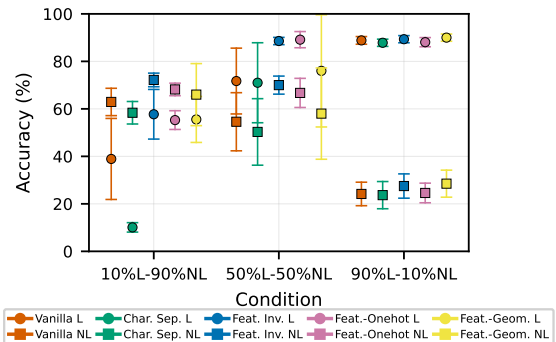


Figure 2: L-shape and NL-shape accuracy across five model architectures. Circles denote L-shape accuracy, squares denote NL-shape accuracy.

Furthermore, we examine how accurately models predict L-shaped forms compared to NL-shaped forms across the three test conditions in Figure 2. As L-shaped training frequency increases, L-shaped accuracy rises, reaching 88-90% in the

90%L condition across all models. Conversely, NL-shaped accuracy drops as NL-shaped verbs become scarcer, falling to 24-29% in the 90%L condition. This symmetric pattern indicates that all models, regardless of architecture, are more strongly driven by training frequency than by the regularity of the pattern: forms from the minority verb class are consistently predicted poorly, even when that class is the regular one.

In the critical 10%L condition, NL-shaped accuracy exceeds L-shaped accuracy for all models, reflecting the dominance of NL-shaped verbs in training. The gap is most extreme for CHARACTER-SEPARATED (L: 10%, NL: 58%) and VANILLA (L: 37%, NL: 62%), whereas Class 2 models maintain more balanced performance (FEATURE-INVARIANT: L: 58%, NL: 72%; FEATURE-ONEHOT: L: 55%, NL: 68%; FEATURE-GEOMETRIC: L: 56%, NL: 67%).

#### 4.1.2. Stem accuracy by verb type

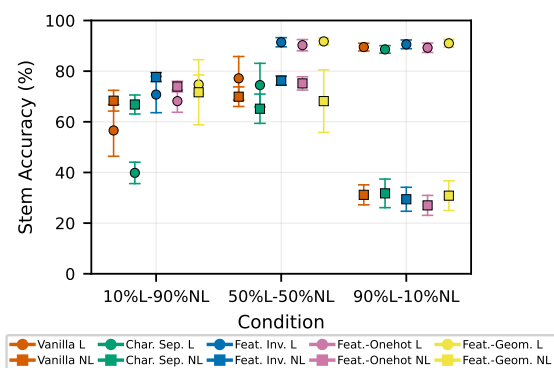


Figure 3: Stem accuracy separated by verb type. Circles denote L-shaped stem accuracy; squares denote NL-shaped stem accuracy.

Figure 3 evaluates stem accuracy, isolating the stem-final consonant alternation that defines the L-shaped pattern from suffix errors. Stem accuracy is consistently higher than sequence accuracy.

All models reach 89-92% L-shaped stem accuracy in the 90%L condition. NL-shaped stem accuracy decreases with decreasing NL training frequency, falling to 27-31% in the 90%L condition, mirroring the frequency dependence observed at the sequence level.

In the critical 10%L condition, Class 2 models show a clear advantage in L-shaped stem accuracy: FEATURE-GEOMETRIC (75%), FEATURE-INVARIANT (71%), and FEATURE-ONEHOT (68%) all outperform VANILLA (57%) and CHARACTER-SEPARATED (40%). This indicates that position-invariant tag encoding helps models learn the correct stem alternation even from limited L-shaped exposure.

#### 4.1.3. Paradigm shape analysis

We apply the paradigm shape analysis described in Section 3.4.3. Figures 4-8 show raw (untransformed) per-cell mean stem accuracy ( $\pm$  SD) across all test lemmas for the three conditions. Black-dashed borders indicate which cells belong to the L-shape for reference. Class 1 models show a clear accuracy contrast between L-cells and the remaining indicative cells even at 10%L, whereas Class 2 models show more uniform accuracy across cells.

To quantify whether this pattern reflects the L-shape, we apply the scoring transformation and  $k$ -means clustering described in Section 3.4.3. We report results across all three frequency conditions because the critical question is not merely *whether* models can recover the L-shaped clustering, but *how much L-shaped exposure they require to do so*. A perfect L-shape would yield high transformed scores for L-cells (1SG.IND + all subjunctive) and low scores for the remaining indicative cells, producing two distinct clusters. In the 50%L and 90%L conditions, all five architectures recover the correct L-shaped clustering. Table 2 shows the cluster assignments in the critical 10%L condition, where L-shaped verbs are rare in training.

Cell	Exp.	Van	C-Sep	F-Inv	F-1H	F-Geo
1SG.IND	L	L	L	L	L	L
2SG.IND	NL	NL	NL	NL	NL	NL
3SG.IND	NL	NL	NL	NL	NL	NL
1PL.IND	NL	NL	NL	NL	NL	NL
2PL.IND	NL	NL	NL	NL	NL	NL
3PL.IND	NL	NL	NL	NL	NL	NL
1SG.SBJV	L	L	L	L	L	L
2SG.SBJV	L	L	L	L	L	L
3SG.SBJV	L	L	L	L	L	L
1PL.SBJV	L	NL*	L	L	L	L
2PL.SBJV	L	NL*	L	L	L	L
3PL.SBJV	L	L	L	L	L	L

Table 2:  $K$ -means cluster assignments in the 10%L condition. Column abbreviations: Van = VANILLA, C-Sep = CHARACTER-SEPARATED, F-Inv = FEATURE-INVARIANT, F-1H = FEATURE-ONEHOT, F-Geo = FEATURE-GEOMETRIC, Exp. = expected cluster under a perfect L-shape. L = L-shape cluster, NL = non-L cluster. \*Misclassified cells.

All three Class 2 models and CHARACTER-SEPARATED correctly cluster 1SG.IND with all six subjunctive cells. VANILLA achieves only a partial L-shape: 1PL.SBJV and 2PL.SBJV are misclassified with the non-L indicative cells. Although CHARACTER-SEPARATED achieves the correct grouping, its L-shape scores are lower than those of the Class 2 models, indicating weaker paradigm differentiation.

## 4.2. Nonce verb evaluation

### 4.2.1. Stem accuracies: models vs. humans

Figure 9 compares relaxed stem accuracy on the 15 nonce verbs from Nevins et al. (2015) across




Figure 4: VANILLA: per-cell stem accuracy across the three frequency conditions.




Figure 5: CHARACTER-SEPARATED: per-cell stem accuracy across the three frequency conditions.

all five architectures and human participants. All models outperform the human mean of 33%: in the 10%L condition, CHARACTER-SEPARATED and VANILLA reach 49% and 48% respectively, while the three Class 2 models cluster around 43-44%.

Unlike the real-verb results, increasing L-shaped training frequency does not improve nonce-verb stem accuracy. All architectures decline from 10%L to 90%L, with Class 2 models dropping more sharply (from 43-44% to 32-37%) than Class 1 models (from 48-49% to 37-42%). This pattern, consistent with Kakolu Ramarao et al. (2025), suggests that models trained on higher proportions of L-shaped verbs generalize less effectively to novel stems.

#### 4.2.2. Paradigm shape: models vs. humans

Figures 10 and 11 show per-cell nonce-verb stem accuracy for models and humans, respectively. Class 1 models, Class 2 models, and humans each exhibit distinct paradigm shapes. Human participants (Figure 11) show higher accuracy on 1SG.IND (37%) than on subjunctive cells (27-29%), suggesting that they preferentially generalize the stem alternation to 1SG.IND, the most frequent and salient L-shaped cell, rather than to subjunctive forms.

Class 2 models show the opposite pattern: subjunctive stem accuracy is high (79-88% at 10%L) while 1SG.IND accuracy is near zero (0-4%). These models have learned to produce the L-shaped stem in subjunctive cells but fail to extend it to 1SG.IND for novel items, yielding a partial L-shape that captures




Figure 6: FEATURE-INVARIANT: per-cell stem accuracy across the three frequency conditions.




Figure 7: FEATURE-ONEHOT: per-cell stem accuracy across the three frequency conditions.

the mood distinction but misses the defining feature of the morpheme, the inclusion of 1SG.IND. Class 1 models (VANILLA, CHARACTER-SEPARATED) show near-uniform accuracy across all cells (47-50% at 10%L), with no preference for indicative or subjunctive. These models do not exhibit any paradigm shape on nonce verbs.

## 5. Discussion

**Positional encoding and tag representation (RQ1).** Of the two architectural dimensions we manipulated, positional encoding proves decisive. When L-shaped verbs are scarce in training, position-invariant models recover higher L-shaped accuracy than sequential-PE models across all metrics. How tags are represented, whether as single indivisible vocabulary entries (e.g.,  $v$ ; IND; PRS; 1; SG), one-hot feature vectors, or linguistically motivated feature geometries, matters far less: the three position-invariant architectures perform comparably despite substantial differences in how they encode morphosyntactic information. As L-shaped training frequency increases, all architectures converge, suggesting that position-invariant encoding acts as an inductive bias that compensates for data sparsity. Across all architectures, the minority verb class in training is consistently predicted poorly, confirming the frequency dependence shown by Kakolu Ramarao et al. (2025). However, position-invariant models maintain higher L-shaped accuracy than sequential-PE models un-




Figure 8: FEATURE-GEOMETRIC: per-cell stem accuracy across the three frequency conditions.




Figure 9: Stem accuracy on nonce verbs from Nevins et al. (2015). Human results (left) are shown alongside model results across the three frequency conditions.

der the same low-frequency conditions, suggesting that the reliance on input frequency that Kakolu Ramarao et al. (2025) observed is a property of the vanilla architecture rather than of neural inflection models in general.

**Structural pattern acquisition (RQ2).** Position-invariant models recover the correct L-shaped clustering, grouping 1SG.IND with all subjunctive cells, even from limited exposure. Among sequential-PE models, VANILLA only partially captures the pattern, misclassifying 1PL.SBJV and 2PL.SBJV, while CHARACTER-SEPARATED achieves the correct grouping but with lower L-shape scores, indicating weaker paradigm differentiation. The per-cell stem accuracy heatmaps corroborate this: Class 2 models achieve higher and more uniform accuracy across all cells. Class 1 models show lower overall accuracy with a wider gap between L-cells and the remaining indicative cells, reflecting greater difficulty with the pattern. This suggests that position-invariant models acquire the L-shape as a structural paradigmatic pattern.

**Productive generalization and comparison with humans (RQ3).** Yet none of the models extend the L-shaped pattern productively to novel stems. Position-invariant models generalize the L-shaped stem to subjunctive cells but fail almost entirely on 1SG.IND, while humans do the opposite (Nevins

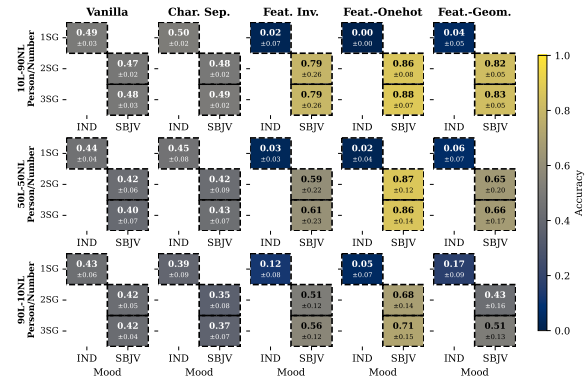


Figure 10: Per-cell mean stem accuracy on nonce verbs for all five architectures across the three frequency conditions.




Figure 11: Per-cell mean stem accuracy on nonce verbs for human participants from Nevins et al. (2015).

et al., 2015). This likely reflects a difference between the two test sets: the real verb test set shares its stem alternation patterns with training, whereas the nonce verbs introduce consonant alternations absent from Spanish. Models succeed when alternation patterns are familiar but fail when both stems and alternations are novel, consistent with prior evidence that neural inflection models struggle to match human wug-test behavior (Corkery et al., 2019; Weissweiler et al., 2023). Neither model class reproduces the human paradigm shape, indicating that human morphological generalization differs qualitatively from that of the transformer architectures tested here.

## 6. Conclusion

We compared five encoder-decoder transformers on their ability to acquire and generalize the Spanish L-shaped morpheme. The critical architectural distinction is between the two classes of models, not between individual architectures within each class: models within Class 1 (sequential PE) and within Class 2 (position-invariant PE) behave similarly regardless of how morphosyntactic tags are represented. Position-invariant positional encoding enables learning the L-shaped paradigm even when L-shaped verbs are scarce. However, no ar-

chitecture productively generalizes this pattern to novel stems, and all models diverge qualitatively from human generalization behavior. These findings contribute to the debate on the cognitive status of morphomic patterns: the L-shape is statistically learnable from distributional evidence, but the generalization mechanisms of character-aware transformers differ from those of human speakers. Future work should probe models' internal representations to determine whether the L-shape is encoded but not realized in production, extend the analysis to Italian and Portuguese where experimental data exists (Nevins et al., 2015; Cappellaro et al., 2024), and explore training regimes that better support generalization of minority patterns.

## 7. Limitations

Whether our findings extend to other Romance languages or to the complete paradigm remains untested. The nonce verb stimuli from Nevins et al. (2015) introduce fricative-stop alternations absent from Spanish, making it difficult to separate failures of paradigm generalization from failures on novel phonological alternations. All five models share the same hyperparameters from Kakolu Ramarao et al. (2025), which may not be optimal for each architecture individually. Additionally, our models are trained on paradigm-balanced data where each cell appears equally often, whereas humans encounter indicative forms, particularly first-person singular, far more frequently than subjunctive forms in natural language. Finally, we compare only transformer variants. Non-neural models such as linear discriminative learning have shown competitive performance on cognitively plausible morphological generalization tasks (Jeong et al., 2023) and may behave differently on the acquisition and generalization of morphomic patterns.

## CRedit authorship contribution statement

We follow the CRedit taxonomy<sup>2</sup>. Conceptualization: [AKR]; Data curation: [AKR]; Formal Analysis: [AKR]; Investigation: [AKR, KT, DB]; Methodology: [AKR, KT, DB]; Supervision: [KT, DB]; Visualization: [AKR]; and Writing – original draft: [AKR] and Writing – review & editing: [KT, DB, AKR].

## Acknowledgements

We gratefully acknowledge the support of the central HPC system “HILBERT” at Heinrich Heine University Düsseldorf. We thank Martin Krämer and Niklas Wiskandt for their valuable feedback. We

<sup>2</sup><https://credit.niso.org/>

also thank the reviewers for their constructive feedback, which improved our work.

## 8. Bibliographical 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*. Oxford University Press.
- 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.
- Antonios Anastasopoulos and Graham Neubig. 2019. *Pushing the limits of low-resource morphological inflection*. In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 984–996, Hong Kong, China. Association for Computational Linguistics.
- Chiara Cappellaro, Nina Dumrukic, Isabella Fritz, Francesca Franzon, and Martin Maiden. 2024. *The cognitive reality of morphemes: evidence from Italian*. *Morphology*, 34(1):33–71.
- 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, 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. 2018. *The CoNLL–SIGMORPHON 2018 shared task: Universal morphological reinflection*. In *Proceedings of the CoNLL–SIGMORPHON 2018 Shared Task: Universal Morphological Reinflection*, pages 1–27, Brussels. Association for Computational Linguistics.
- Ryan Cotterell, Christo Kirov, John Sylak-Glassman, Géraldine Walther, Ekaterina Vylomova, Patrick Xia, Manaal Faruqui, Sandra

- Kübler, David Yarowsky, Jason Eisner, and Mans Hulden. 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. 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, Berlin, Germany. Association for Computational Linguistics.
- 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.
- Philipp Dufter, Martin Schmitt, and Hinrich Schütze. 2022. [Position information in transformers: An overview](#). *Computational Linguistics*, 48(3):733–763.
- Omer Goldman, David Guriel, and Reut Tsarfay. 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), ACL 2022, Dublin, Ireland, May 22-27, 2022*, pages 864–870. Association for Computational Linguistics.
- Heidi Harley and Elizabeth Ritter. 2002. [Person and number in pronouns: A feature-geometric analysis](#). *Language*, 78(3):482–526.
- Adi Haviv, Ori Ram, Ofir Press, Peter Izsak, and Omer Levy. 2022. [Transformer language models without positional encodings still learn positional information](#). In *Findings of the Association for Computational Linguistics: EMNLP 2022*, pages 1382–1390, Abu Dhabi, United Arab Emirates. Association for Computational Linguistics.
- Michael Y. Hu, Aaron Mueller, Candace Ross, Adina Williams, Tal Linzen, Chengxu Zhuang, Ryan Cotterell, Leshem Choshen, Alex Warstadt, and Ethan Gotlieb Wilcox. 2024. [Findings of the second BabyLM challenge: Sample-efficient pre-training on developmentally plausible corpora](#). In *The 2nd BabyLM Challenge at the 28th Conference on Computational Natural Language Learning*, pages 1–21, Miami, FL, USA. Association for Computational Linguistics.
- Dieuwke Hupkes, Mario Giulianelli, Verna Dankers, Mikel Artetxe, Yanai Elazar, Tiago Pimentel, Christos Christodoulopoulos, Karim Lasri, Naomi Saphra, Arabella Sinclair, Dennis Ulmer, Florian Schottmann, Khuyagbaatar Batsuren, Kaiser Sun, Koustuv Sinha, Leila Khalatbari, Maria Ryskina, Rita Frieske, Ryan Cotterell, and Zhijing Jin. 2023. [A taxonomy and review of generalization research in nlp](#). *Nature Machine Intelligence*, 5(10):1161–1174.
- Cheonkam Jeong, Dominic Schmitz, Akhilesh Kakolu Ramarao, Anna Stein, and Kevin Tang. 2023. [Linear discriminative learning: a competitive non-neural baseline for morphological inflection](#). In *Proceedings of the 20th SIGMORPHON workshop on Computational Research in Phonetics, Phonology, and Morphology*, pages 138–150, Toronto, Canada. Association for Computational Linguistics.
- Akhilesh Kakolu Ramarao, Kevin Tang, and Dinah Baer-Henney. 2025. [Frequency matters: Modeling irregular morphological patterns in Spanish with transformers](#). In *Findings of the Association for Computational Linguistics: ACL 2025*, pages 4474–4489, Vienna, Austria. Association for Computational Linguistics.
- Katharina Kann, Ryan Cotterell, and Hinrich Schütze. 2017. [Neural multi-source morphological reinflection](#). In *Proceedings of the 15th Conference of the European Chapter of the Association for Computational Linguistics, EACL 2017, Valencia, Spain, April 3-7, 2017, Volume 1: Long Papers*, pages 514–524. 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.
- Diederik P. Kingma and Jimmy Ba. 2015. Adam: A method for stochastic optimization. *CoRR*, abs/1412.6980.
- 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.

- Jordan Kodner and Salam Khalifa. 2022. [SIGMORPHON–UniMorph 2022 shared task 0: Modeling inflection in language acquisition](#). In *Proceedings of the 19th SIGMORPHON Workshop on Computational Research in Phonetics, Phonology, and Morphology*, pages 157–175, Seattle, Washington. Association for Computational Linguistics.
- Jordan Kodner, Sarah Payne, Salam Khalifa, and Zoey Liu. 2023. [Morphological inflection: A reality check](#). In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 6082–6101, Toronto, Canada. 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.
- Martin Maiden. 2011. [Allomorphy, autonomous morphology and phonological conditioning in the history of the Daco-Romance present and subjunctive](#). *Transactions of the Philological Society*, 109(1):59–91.
- Martin Maiden. 2018. *The Romance verb: Morphomic structure and diachrony*. Oxford University Press.
- Martin Maiden. 2021. [The morphome](#). *Annual Review of Linguistics*, 7:89–108.
- P H Matthews. 1972. *Cambridge studies in linguistics: Inflectional morphology: A theoretical study based on aspects of Latin verb conjugation series number 6*. Cambridge studies in linguistics. Cambridge University Press, Cambridge, England.
- Kate McCurdy, Sharon Goldwater, and Adam Lopez. 2020. [Inflecting when there's no majority: Limitations of encoder-decoder neural networks as cognitive models for German plurals](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 1745–1756, Online. Association for Computational Linguistics.
- Andrew Nevins, Cilene Rodrigues, and Kevin Tang. 2015. [The rise and fall of the L-shaped morphome: diachronic and experimental studies](#). *Probus: International Journal of Latin and Romance Linguistics*, 27(1):101–155.
- Paul O'Neill. 2025. [Defective verbs in portuguese: a morphomic approach](#). *Cognitive Linguistics*, 36(2):153–181.
- Ben Peters and André F. T. Martins. 2019. [IT–IST at the SIGMORPHON 2019 shared task: Sparse two-headed models for inflection](#). In *Proceedings of the 16th Workshop on Computational Research in Phonetics, Phonology, and Morphology*, pages 50–56, Florence, Italy. Association for Computational Linguistics.
- David E. Rumelhart, James L. MollClelland, and the PDP Research Group. 1986. *Parallel Distributed Processing: Explorations in the Microstructure of Cognition: Foundations*, volume 1. The MIT Press.
- Peter Shaw, Jakob Uszkoreit, and Ashish Vaswani. 2018. [Self-attention with relative position representations](#). 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 464–468, New Orleans, Louisiana. Association for Computational Linguistics.
- Andrew Spencer and Mark Aronoff. 1994. [Morphology by itself: Stems and inflectional classes](#). *Language*, 70:811.
- Gregory T Stump. 2001. *Inflectional morphology: A theory of paradigm structure*, volume 93. Cambridge University Press.
- Jianlin Su, Murtadha Ahmed, Yu Lu, Shengfeng Pan, Wen Bo, and Yunfeng Liu. 2024. [Roformer: Enhanced transformer with rotary position embedding](#). *Neurocomputing*, 568:127063.
- John Sylak-Glassman. 2016. [The composition and use of the universal morphological feature schema \(uni-morph schema\)](#). Technical report, Center for Language and Speech Processing, Johns Hopkins.
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N. Gomez, Lukasz Kaiser, and Illia Polosukhin. 2017. [Attention is all you need](#). *CoRR*, abs/1706.03762.
- 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.

Yu-An Wang and Yun-Nung Chen. 2020. [What do position embeddings learn? an empirical study of pre-trained language model positional encoding](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 6840–6849, Online. Association for Computational Linguistics.

Alex Warstadt, Aaron Mueller, Leshem Choshen, Ethan Wilcox, Chengxu Zhuang, Juan Ciro, Rafael Mosquera, Bhargavi Paranjabe, Adina Williams, Tal Linzen, and Ryan Cotterell. 2023. [Findings of the BabyLM challenge: Sample-efficient pretraining on developmentally plausible corpora](#). In *Proceedings of the BabyLM Challenge at the 27th Conference on Computational Natural Language Learning*, pages 1–34, Singapore. Association for Computational Linguistics.

Leonie Weissweiler, Valentin Hofmann, Anjali Kantharuban, Anna Cai, Ritam Dutt, Amey Hengle, Anubha Kabra, Atharva Kulkarni, Abhishek Vijayakumar, Haofei Yu, Hinrich Schuetze, Kemal Oflazer, and David Mortensen. 2023. [Counting the bugs in ChatGPT’s wugs: A multilingual investigation into the morphological capabilities of a large language model](#). In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pages 6508–6524, Singapore. Association for Computational Linguistics.

Shijie Wu, Ryan Cotterell, and Mans Hulden. 2021. [Applying the transformer to character-level transduction](#). In *Proceedings of the 16th Conference of the European Chapter of the Association for Computational Linguistics: Main Volume*, pages 1901–1907, Online. Association for Computational Linguistics.