Kyle Gorman and Daniel Yakubov. NELS 55, 2024.

Acquiring inflectional gaps with indirect negative evidence: evidence from Russian

Introduction In Russian there are roughly forty nouns which lack a well-formed genitive plural (gen.pl.; Pertsova 2005) and roughly sixty-five verbs without an acceptable 1st person singular (1sg.) non-past (Halle 1973, Pertsova 2016); examples are given below. In each case, one can imagine one or more possible words which might fill the hole in the paradigm, but speakers judge all such possibilities unacceptable.

(1)

nom.sg.	gen.pl.	
dno	*dn, *den	'bottom (of a container or body of water)'
kočerga	*kočerg	'fire poker'
mečta	*mečt	'dream; fantasy'

(2)

infinitive	1sg. non-past	
derzit'	*deržu	'be insolent'
pobedit'	*pobežu, *pobeždu	'win; overcome'
pylesosit'	*pylesošu	'vacuum'

Inflectional gaps (or paradigm gaps, or morphological defectivity) were noted as far back as late antiquity, and have been documented in several dozen languages. Yet there is little consensus about how inflectional gaps are represented grammatically (e.g., Rice and Blaho 2009, Baerman et al. 2010, Gorman and Yang 2019, Sims 2015), and even less is known about how they are acquired by children. To put it simply, children must learn not only what they can say, but also must learn what they cannot say. One can be reasonably certain inflectional gaps are not acquired via "direct" negative evidence (in the sense of Marcus 1993) alone: there is little evidence that Russian speakers are ever told—by parents, teachers, or peers—not to say **mečt* or **pobežu*. However, it has been suggested that speakers might make use of some form of **indirect** (or **weak**) **negative evidence** to learn inflectional gaps (e.g., Orgun and Sprouse 1999, Pertsova 2005, Daland et al. 2007, Orgun and Sprouse 1999). In this study we attempt to exhaustively evaluate this idea using data from Russian. We focus on Russian because it has a reasonably large number of inflectional gaps, they are overall very well characterized, and because this language is supported by high-quality automated morphological analysis tools needed to perform this evaluation at scale.

Modeling indirect negative evidence We propose three models of how a child learner might infer the presence of an inflectional gap. Suppose that for each (possible or attested) inflected word w, there is a corresponding lemma l which represents the lexeme, and a morphosyntactic feature bundle f which represents that word's cell in the paradigm.[1] Let p_w, p_l, and p_f denote the probability of a word, lemma, and feature bundle, respectively, computed from a morphologically annotated corpus using ordinary maximum likelihood estimation. All else held equal (i.e., under a simple independence assumption), p_w should be proportional to the product of p_l and p_f, and if p_w is substantially less than $p_l p_f$ leads one to expect, the learner might posit that lemma l has an inflectional gap in paradigm cell f. We consider three ways one might measure

[1] Here the terms **lexeme**, **cell**, and **paradigm** are used pre-theoretically; the models proposed are general enough to encompass both inferential (i.e., paradigmatic) and realizational theories of morphology.

how far a given cell has diverged from its expected probability. **Absolute divergence** (3a) is the arithmetic difference between observed and expected frequencies; the **log-odds ratio** (3b) is a monotonic transform of the traditional observed-over-expected ratio, and the **standard score** (3c) converts absolute divergence into a z-score, which takes into account the greater uncertainty associated with probability estimation on rarer lexemes and/or feature bundles.

(3)
a. absolute divergence: $\delta_w = p_w - p_l p_f$
b. log-odds ratio: $L_w = \log p_w - \log p_l - \log p_f$
c. standard score: $z_w = \frac{p_w - p_l p_f}{\sqrt{s_l^2 s_f^2 + s_l^2 p_f^2 + s_f^2 p_l^2}}$ (s^2: sample variance)

Data Pertsova (2016) furnishes a list of Russian verbs defective in the 1sg. non-past; following Pertsova (2005) we derive a list of nouns from those labeled either as either ‘difficult’ or ‘do not exist’ in the gen.pl. in an authoritative morphological dictionary (Zaliznyak 1977). Ideally one would compute the relevant statistics from a corpus manually annotated for morphological information. However, on average the defective lexemes occur quite rarely in the largest manually annotated corpus of Russian, the 1.5m sentence SynTagRus corpus (Droganova et al. 2018). Something larger is needed to get around the inherent sparsity of this richly inflected language. Therefore, we use a large corpus of Russian web text from the CC-100 corpus (Wenzek et al. 2020); this yields 232m sentences and over 3.3b tokens after some filtering. We then perform automatic morphological analysis, yielding part of speech, lemma, and a morphosyntactic feature bundle for each token, using a state-of-the-art neural network-based morphological analyzer (Straka et al. 2019, Kondratyuk and Straka 2019) trained on the SynTagRus data.

Evaluation We first compare the grammatical descriptions of Russian gaps with the morphologically analyzed corpus. With a few exceptions, we find that the claimed inflectional forms are actually avoided in modern written Russian, though nearly all the defective lexemes are attested in other forms. The automatically annotated corpus is then used to compute the statistics in (3) for the gen.pl. for nouns and the 1sg. non-past for verbs. For each of the three statistics, we rank the noun and verb lexeme according to that statistic, and determine whether there is any threshold which separates defective and non-defective lexemes.[2]

Results We find, however, that for both nouns and verbs, and for all three statistics, no such threshold exists. While most of the defective lexemes have low scores for these statistics, as expected, so do many other lexemes which are not defective. This is perhaps unsurprising: whatever signal is present is swamped by unrelated pragmatic pressures—for example, the English word *moons* is perfectly acceptable, but quite rare, since our planet only has one moon—and/or the inherent sparsity of the data.

Conclusions Indirect negative evidence, as formalized above, is unlikely to be the primary mechanism by which inflectional gaps are acquired. A hypothetical child provided with a veritable glut of primary linguistic data—billions of words—would still not be able to conclude that *mečta* lacks a gen.pl. on a purely statistical basis, since defective paradigm cells have statistical profiles similar to non-defective nouns and verbs. Children must use of some additional information about inflectional patterns; for instance, they might notice the existence of different inflectional patterns competing to realize the relevant cells, such as the various gen.pl. allomorphs and/or stem allomorphy introduced by yer realization in certain classees of nouns.[3] We release our data—a database of counts of words, lemmas, and feature bundles—under a Creative Commons license to support future work, including the development of new models of indirect negative evidence.

---

[2] This is equivalent to fitting a single-feature, one-sided **stump classifier** in machine learning.

[3] See Albright (2003) and Yang (2016:ch. 4) for two rather different accounts of how this might work.