Learning Input Strictly Local Functions: Comparing Approaches with Catalan Adjectives

Alexander Shilen
Johns Hopkins University, ashilen1@jhu.edu

Colin Wilson
Johns Hopkins University, colin.wilson@jhu.edu

Recommended Citation

Learning Input Strictly Local Functions: Comparing Approaches with Catalan Adjectives

Alex Shilen
ashilen1@jhu.edu

Colin Wilson
colin.wilson@jhu.edu

Johns Hopkins University

1 Introduction

Input strictly local (ISL) functions are a class of subregular transductions that have well-understood mathematical and computational properties and that are sufficiently expressive to account for a wide variety of attested morphological and phonological patterns (e.g., Chandlee, 2014; Chandlee et al., 2014; Chandlee, 2017; Chandlee et al., 2018; Chandlee and Heinz, 2018). In this study, we compared several approaches to learning ISL functions: the ISL function learning algorithm (ISLFLA; Chandlee, 2014; Chandlee et al., 2014) and the classic OSTIA learner to which it is related (Oncina et al., 1993); the Minimal Generalization Learner (MGL; Albright and Hayes, 2002, 2003); and a novel deep neural network model presented here (DNN-ISL).

The four models were evaluated on their ability to learn the mapping from feminine singular (*fem.sg.*) to masculine singular (*masc.sg.*) surface forms of Catalan adjectives (and, separately, from provided underlying representations to *fem.sg.* and *masc.sg.* surface forms). The mappings to *masc.sg.* forms in Catalan involve several phonological modifications at the right edge of the word (e.g., Mascaró, 1976), the empirical focus of our study. The relevant processes include obstruent devoicing and strengthening (e.g., [rɔʒə] fem.sg. → [rɔtʃ] masc.sg. 'red'), post-tonic n-Deletion (e.g., [sanə] → [sa] 'healthy'), and cluster simplification (e.g., [blaŋkə] → [blaŋ] 'white'). There are opaque, counterfeeding interactions among some of the processes (e.g., [fəkundə] → [fəkun] / *[fəku] 'fertile'), consistent with the idea that the mappings are input- rather than output- determined (see Chandlee et al., 2018). A small number of apparent lexical exceptions to the typical modification pattern (e.g., [blanə] → [blan] / *[bla] 'soft') are problematic for ISL learners that assume perfect homogeneity.

Our main findings were that the DNN-ISL learner achieved high accuracy on the Catalan data, with MGL coming in a close second, while ISLFLA and OSTIA performed much worse — either failing to learn any mapping at all or predicting the correct output for less than 5% of held-out cases, even when lexical exceptions were removed from the data (see Table 1).

2 Data

The FestCat project (Bonafonte et al., 2008) provides broad transcriptions for more than 53,000 adjectival surface forms in two major dialects of Catalan. We considered the Central Catalan forms and restricted our data to the nearly 6,500 lemmas that are also attested in a subtitle lexicon (Boada et al., 2020). While our main focus was on learning, we also developed a hand-written ISL transducer for the mapping to masc.sg. forms that is highly accurate (> 98% correct), along with custom code to derive plausible underlying representations from *masc.sg.* $\sim$ *fem.sg.* pairs.

3 Models

For the purposes of this abstract, we assume familiarity with ISLFLA, OSTIA, and MGL. We verified that the implementation of MGL learns only ISL phonological rules — rules conditioned on local phonological context in the result of a morphological operation such as affixation or truncation — a connection that has not previously been made in the literature.

The deep neural network model proposed here (DNN-ISL) also applies morphological operations followed by phonological modifications, the latter being implemented with weighted constraints rather than rules. A phonological constraint as learned by DNN-ISL is defined by: a three-segment featural pattern specifying the input context to which the constraint applies; a preference for one type of modification applied to the center segment of the context (i.e., deletion, epenthesis before/after, or feature change); target output features in the case

Data set	Data split	DNN-ISL	MGL	OSTIA	ISLFLA	Baseline
All	train	.95	.79	1.0 / .84*	— / .89* (n=6)	.38
	val	.95	.79	.02 / .79*	— / .83* (n=6)	.39
	test	.95	.80	.02 / .79*	— / .83* (n=6)	.39
Regular	train	.98	.98	1.0 / 1.0*	— /1.0*	.41
	val	.98	.97	.04 / .98*	— / .96*	.41
	test	.98	.97	.03 / .98*	— / .96*	.41

Table 1: Mean model accuracy for mapping from *fem.sg.* to *masc.sg.* in 10 independent runs with random splits (20% train, 10% validation, 70% test) of all Catalan adjective data or the 'regular' subset with exceptions omitted. ISLFLA failed to learn a transducer in each run, as indicated by '—', returning "Insufficient data" (see Chandlee 2014:117). The Baseline model simply subtracted the feminine suffix -ə. *OSTIA and ISLFLA performed better when input and output forms were trimmed to their final VC*(V) sequences, as shown after the slash, but ISLFLA still failed to learn a transducer in 4/10 of the runs on the full data set. SDs were lower than .04 in most cells. Perfect performance on the training data, shown here for OSTIA only, is available to any model with sufficient memory.

of epenthesis or change; and a real-valued strength. Each constraint computes the degree to which its context matches every three-segment window in the input (i.e., it applies a novel feature based convolution operation to the input) and imposes its preferred modification in proportion to the degree of match and its strength. These preferences are summed over constraints for each input position and applied to the positions independently to derive the phonological output. The parameters of the constraints are straightforwardly interpretable and visualizable as real-valued phonological feature coefficients, modification-type logits, and strengths. The model is fully differentiable and was trained with the Adagrad optimizer on small mini-batches for 20 epochs.

4   Results

We evaluated all four models on the same training/validation/testing data, as summarized in Table 1. ISLFLA and OSTIA were unable to learn accurate mappings except when the *fem.sg.* and *masc.sg.* forms were artificially trimmed to their final VC∗(V) sequences — a strong, language-specific bias to attend to changes at the end of the word that the other models did not require. Results for larger training splits, and for mapping from URs to SRs, were similar. The errors made by DNN-ISL mostly involved underapplication of deletion (e.g., *[blaŋk]).

5   Contributions & future directions

In summary, we evaluated four learning models on an ISL phonological mapping (with a small number of exceptions) found in a large, realistic body of natural language data. The models that have proofs

of learnability and efficiency, ISLFLA and OSTIA, performed much worse than models that currently lack such theoretical guarantees but share the inductive bias for ISL patterns. The results highlight the need for further empirical and formal study of high-performing subsymbolic models such as DNN-ISL, and extension of our model to output-based patterns and learning of underlying representations. We plan to release our processed data, hand-written ISL transducer, and model implementations.

References

Adam Albright and Bruce Hayes. 2002. Modeling English past tense intuitions with minimal generalization. In *Proceedings of the ACL-02 Workshop on Morphological and Phonological Learning*, pages 58–69. Association for Computational Linguistics.

Adam Albright and Bruce Hayes. 2003. Rules vs. analogy in English past tenses: A computational/experimental study. *Cognition*, 90(2):119–161.

Roger Boada, Marc Guasch, Juan Haro, Josep Demestre, and Pilar Ferré. 2020. SUBTLEX-CAT: Subtitle word frequencies and contextual diversity for Catalan. *Behavior Research Methods*, 52(1):360–375.

Antonio Bonafonte, Jordi Adell, Ignasi Esquerra, Silvia Gallego, Asunción Moreno, and Javier Pérez. 2008. Corpus and voices for Catalan speech synthesis. In *Proceedings of the Sixth International Conference on Language Resources and Evaluation (LREC'08)*, Marrakech, Morocco. European Language Resources Association (ELRA).

Jane Chandlee. 2014. *Strictly Local Phonological Processes*. Ph.D. thesis, University of DelawarePro-Quest.

Jane Chandlee. 2017. Computational locality in morphological maps. *Morphology*, 27(4):599–641.

Jane Chandlee, Rémi Eyraud, and Jeffrey Heinz. 2014. Learning strictly local subsequential functions. *Transactions of the Association for Computational Linguistics*, 2:491–504.

Jane Chandlee and Jeffrey Heinz. 2018. Strict locality and phonological maps. *Linguistic Inquiry*, 49(1):23–59.

Jane Chandlee, Jeffrey Heinz, and Adam Jardine. 2018. Input strictly local opaque maps. *Phonology*, 35(2):171–205.

Joan Mascaró. 1976. *Catalan Phonology and the Phonological Cycle*. Ph.D. thesis, MIT, Cambridge, MA.

José Oncina, Pedro García, and Enrique Vidal. 1993. Learning subsequential transducers for pattern recognition interpretation tasks. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 15(5):448–458.